

MEANING TRANSFER REVISITED

David Liebesman
University of Calgary

Ofra Magidor
University of Oxford

§1. Introduction

A waiter notices that a diner who ordered a ham sandwich has just stormed out of the restaurant, and calls to his colleague:

- (1) The ham sandwich left without paying!

Some diners are discussing what each person ate, and one of them elaborates:

- (2) Bill was a ham sandwich.

In these scenarios, (1) is used to communicate that the customer who ordered the ham sandwich left without paying, and (2) is used to communicate that Bill ate a ham sandwich. However, on their most straightforward interpretations, (1) and (2) do not semantically express these propositions. These sentences exemplify *meaning transfer*.¹

For now, we won't try to characterise what meaning transfer is more precisely, but we'll assume that the phenomenon exhibited by the above cases is recognisable, and that the envisioned utterances of (1) and (2) are uncontroversial instances of this phenomenon.

Slightly more controversially, we will assume that (in the scenarios described) it is the predicate which undergoes meaning transfer in (2), while in (1) it is the subject-term.² We will return to clarify and support these assumptions later.

As attested in the literature, meaning transfer interests a variety of theorists across a variety of disciplines.³ However, a number of key issues about its nature remain underexplored. It deserves more attention than it has received.

One reason to be interested in meaning transfer is that it seems closely related to a number of phenomena that have received ample discussion and had substantial impact on the way that philosophers and linguists think about the nature of language and communication: metaphor, loose speech, and conversational implicature. Perhaps meaning transfer is so close to these phenomena that it can't teach us anything new. However, as we'll argue in §3, it is distinct from these and has the potential to impart distinct lessons.

A second reason to be interested in meaning transfer, which we discuss in §5, is that it is an important test case for understanding the semantics/pragmatics interface. There is a strong pre-theoretic inclination to treat meaning transfer as a purely pragmatic phenomenon. However, as Stanley (2005: 227) notes, meaning transfer (which he calls 'deferred meaning') has several properties that we tend to think track semantic content: "Not only is there is a strong intuition that the deferred meaning is part of the intuitive truth-conditions, but the deferred meaning enters into certain linguistic processes, such as anaphora and ellipsis."⁴

A third reason to be interested in meaning transfer is that it is relevant for assessing ontological commitment. Given the possibility of meaning transfer, we cannot conclude from a true occurrence of 'a P is Q' that a P is Q: (1) is true in the scenario described despite the fact that no sandwich left.⁵ So, meaning transfer presents a challenge for those who think that there is a tight link between ontological commitment and natural language truth-conditions. Of course, the discrepancy is not particularly interesting in cases where the transfer is easily identifiable, but the possibility of meaning transfer plays a more significant role in other cases. Consider the debate on *copredication*. Copredication involves true utterances of the form 'a is Q and R' where, it is often argued, nothing could be both Q and R (cf. Chomsky (2000), Pietroski (2005), and Collins (2017)). For example, we can truly say 'Lunch was delicious but lasted hours', even though it might seem that only food can be delicious and only events last hours. In other work (Liebesman & Magidor (2017); Liebesman & Magidor (*forthcoming*)), we have argued that in typical cases of copredication, one should take these ontological commitments at face value (i.e., in this case, there are entities that are both delicious and last hours). However, as Nunberg (1995: 122) stresses, no such commitments would follow if it turns out that either 'Q' or 'R' have undergone meaning transfer.

We begin in §2 by reviewing some of the key linguistic properties of meaning transfer. We then defend some answers to largely unexplored questions about its nature. The first question is whether meaning transfer can be assimilated to other figures of speech. There's a temptation, for instance, to think that meaning transfer is just boring metaphor, or alternatively, that it is just a standard case of conversational implicature. We argue in §3 that this is wrong: meaning transfer cannot be assimilated to other such phenomena. The second question is what sorts of propositions are expressed in cases of meaning transfer, and how these propositions come about. In §4, we argue that meaning transfer is not merely ellipsis and, rather, there is a meaning transfer operator of variable semantic

type that can take different sub-sentential scopes. The third question, which we discuss in §5, is whether meaning transfer is semantic or pragmatic, and what the phenomenon can teach about the semantics/pragmatics distinction. We conclude, in §6, by discussing how meaning transfer may bear on other linguistic phenomena.

§2. The Properties of Meaning Transfer

Meaning transfer exhibits a number of distinctive linguistic properties. The thread running through these is that different types of linguistic behaviour—many of which are often taken to be syntactic and/or semantic—track transferred meaning rather than non-transferred meaning. As will become clear, we can use these properties to distinguish meaning transfer from related phenomena as well as identify its locus. Many of the properties on the list are familiar from the literature, especially Nunberg's work.⁶ Here we refine them and respond to a number of worries from Ward (2004).

Before discussing the linguistic properties of meaning transfer, we need to distinguish two tasks. One task is to identify the properties of meaning transfer, given that it has occurred. Another task is to identify the conditions under which meaning transfer can occur. We're focused squarely on the former. Interesting as the latter is, we suspect that it is highly complex, as the possibility of meaning transfer is sensitive to a wide variety of social factors. That topic is worth (at least!) a paper of its own.⁷

§2.1 Number, Gender, and Animacy marking

When meaning transfer occurs, number agreement typically tracks the transferred meaning rather than the syntactic number of the expression in question. Imagine, for instance, that three customers are sitting at a restaurant counter. One is wearing a tie and shorts, a second is wearing a full suit, and the third is wearing shorts and a t-shirt. The waiter exclaims:

(3) The tie and shorts is leaving!

Ordinarily, a plural definite description like 'the tie and shorts' would combine with a plural copula. However, in the envisioned scenario, the meaning of 'the tie and shorts' is transferred to *the wearer of the tie and shorts*, and the singularity of that meaning ensures that we use the singular form of the copula.

This phenomenon generalizes beyond plural descriptions. Imagine that we're at the beach and we spot a person with particularly toned abs that is about to leave. We may exclaim:

- (4) Hot abs is leaving!⁸

Similar remarks apply to gender marking: gender agreement typically tracks transferred meaning rather than non-transferred meaning.⁹ Consider, for example, the Hebrew translation of ‘The pasta left without paying’, uttered in a context where it communicates the claim that the (male) customer who ordered pasta left without paying. The sentence would be translated as follows:

- (5) Ha-pasta yatza bli leshalem.
The-pasta_f left_m without paying.

The subject term ‘pasta’ is marked as feminine in Hebrew, and ordinarily we would expect the verb to have the same gender marking. But in this case (where we are talking about a male customer) the gender marking of the verb is changed to masculine.

By contrast, the gender marking of ‘Ofra was a soup’ in Hebrew goes with the literal gender marking of ‘Ofra’ (feminine), as we would expect given that the locus of transfer is the predicate. Contrast the following two Hebrew sentences:

- (6) Ofra haita marak
Ofra was_f soup_m
(7) Ma she-Ofra ahla haya marak
What Ofra ate was_m soup_m

In (6), where the subject term (‘Ofra’) is marked as feminine, the verb (‘was’) is also marked as feminine. But in (7), the verb is marked as masculine – the grammatical gender of ‘soup’. This suggests that, in (6), the subject term ‘Ofra’ isn’t transferred to mean ‘what Ofra ate’; if it had been, we would have expected the verb to receive the same gender marking as in (7). Rather, it is the predicate term (‘was soup’) which is the locus of meaning transfer in this case.

Like gender and number marking, animacy marking typically tracks transferred meaning. For instance, whether we use ‘who’ or ‘which’ in a particular sentence tracks the (perhaps perceived) animacy of the transferred meaning, rather than the non-transferred meaning. Imagine that two customers ate ham sandwiches. The one in the blue shirt left without paying, while the one with a green shirt didn’t.

- (8) The ham sandwich who was wearing green left without paying.
(9) # The ham sandwich which was wearing green left without paying.

Our observation in this subsection has been that in cases of meaning transfer, number, gender, and animacy marking typically track transferred rather than the non-transferred meaning.¹⁰ It is important not to confuse this with a stronger

thesis: *that it is always possible to generate meaning transfer interpretations of sentences that (otherwise) exhibit syntactic feature disagreement*. Crucially, we do not subscribe to this stronger thesis. Consider, for example, the following case.

Two groups of children are given the task of creating a drawing – one group of a dragon, and the second, of a dinosaur. The children drawing the dragon are fighting with each other. We could attempt to refer to the children in the first group using the singular noun phrase ‘the dragon’ as follows:

(10)? The dragon are fighting again!

Anecdotal evidence suggests that while some speakers find (10) acceptable, others do not. But insofar (10) is unacceptable, it merely shows that we have not successfully generated a case of meaning transfer. One plausible hypothesis is that number-mismatch is always an obstacle to interpretability and that while sometimes meaning transfer can overcome the obstacle, it cannot always, and when it can is due to a variety of contingent factors.¹¹ Number mismatches aside, recall that the question of when meaning transfer is possible is a pragmatic question that we do not tackle in this paper. Crucially, however, whether or not (10) is acceptable, the following sentence in which the number agreement tracks the syntactic number of the subject term (rather than the relevant transferred meaning) is definitely not an acceptable way to express the relevant transferred meaning:

(11) #The dragon is fighting again!

This is just as we would expect given the weaker thesis we do subscribe to, namely that (typically) *when* meaning transfer occurs, number, gender, and animacy marking track transferred rather than untransferred meaning.

§2.2 Subsequent Predication, Ellipsis, and Anaphora

When a subject term undergoes meaning transfer, the transferred meaning is available as an argument of subsequent predicates while the non-transferred meaning is not. This is demonstrated by adding conjunctions that target either the transferred or non-transferred meanings. Compare (12) and (13):

(12) The ham sandwich left without paying and seemed unsatisfied

(13) # The ham sandwich left without paying and was delicious.

(13) has only the bizarre reading that a particular person both left without paying and was delicious; it doesn’t have the reading that a particular person left without paying, and *the sandwich they ordered* was delicious.

As Barrios (2013) notes, appositives and relative clauses can be used to probe the available subjects of predication as well. Furthermore, as demonstrated in (14), such clauses provide additional evidence that animacy tracks transferred meaning.

- (14) The ham sandwich, who is an annoying guy, left without paying.
- (15) The ham sandwich, an annoying guy, left without paying.

Just as transferred meanings are available as the subjects of subsequent predicates, while non-transferred meanings are not, only transferred meanings are readily available for ellipsis. When an elided predicate depends on an earlier occurrence of a predicate with a transferred meaning, the transferred meaning, rather than the untransferred meaning, is the value of the elided predicate. For example, using the elided predicate in (16) with its transferred meaning (person who ordered a ham sandwich) is entirely acceptable, but attempting to use it in (17), with its untransferred meaning (is a certain type of food) is not:

- (16) Bill is a ham sandwich, as is Joe.
- (17) # Bill is a ham sandwich, as is the thing on this plate.

Similar observations apply to cases of anaphoric reference: only transferred meanings are readily available for subsequent anaphoric reference.

- (18) The ham sandwich left without paying; he was a jerk.
- (19) The ham sandwich left without paying; #it was delicious

Ward (2004) challenges the claim that transferred meaning is available for anaphoric reference while untransferred meaning isn't.¹² He maintains that in some contexts anaphoric reference to the non-transferred meaning is available. Imagine that there is just one filet-mignon orderer with just one filet mignon. Some report that (20) has a salient reading on which it expresses that the filet-mignon orderer says that their steak (the filet mignon) is delicious.¹³

- (20) The filet mignon says that it is delicious.

Even granting that (20) has the alleged reading, the claim that only transferred meanings are readily available for anaphoric reference is not undermined. We know independently that a wide variety of factors can make an object a salient value for a pronoun. It is not at all surprising, then, that in some contexts a non-transferred meaning may be a salient value. After all, we've provided context for (20) that makes the steak itself salient. Furthermore, the interpretations of pronouns can be affected by linguistic context, and the predicate 'delicious', which occurs in (20), usually ascribes a property to food. This further suggests a resolution of the pronoun on which it designates food rather than a person. To

put the point more bluntly, it is plausible that ‘it’ in (20) is deictic rather than anaphoric.¹⁴

A more careful articulation of the interaction between meaning transfer and reference is as follows: when phrase *p* has a transferred meaning, only the transferred meaning is available for subsequent reference *solely in virtue of its association with p*. The idea is that when a non-transferred meaning is the value of a subsequent pronoun, its salience must be explained by something beyond the fact that it is *p*’s non-transferred meaning. That is the exact sort of explanation we just provided for (20). Furthermore, the revised principle yields the prediction that in a wide range of cases without enriched context (including linguistic context) we won’t be able to refer to the non-transferred meaning. Consider, for example, the following variant of (20):

(21) The fillet mignon says that it is boring.

Without significant contextual supplementation, it is much harder to read ‘it’ in (21) as referring to a steak. The contrast between (20) and (21) is due to the fact that the predicate ‘delicious’, by contrast with ‘boring’, provides a contextual cue that we are talking about a piece of food. Similarly, note that we can make the steak sufficiently contextually salient without ever referring to it (directly or indirectly). Thus, consider a waiter that looks at the newly designed menu for the restaurant and says:

(22) ‘Fillet Mignon’ is printed in such an odd font. Good thing all the customers think it’s delicious.¹⁵

That transferred meaning is readily available for anaphora and ellipsis, while non-transferred meaning isn’t, is a specific instance of a more general phenomenon: that transferred meaning is readily available for resolving context-sensitive elements while non-transferred meaning is not. This can be demonstrated by considering the possessive, and using a simplification of an example from Recanati (2004: ch. 4). There are two warriors: one with a lion on their shield and another with a bear on their shield. These warriors are tasked with killing two animals: a lion and a bear. As it turns out, the bear-shielded warrior kills the lion, and vice-versa. Observing the carnage, with swords sticking out of the slain animals, we may utter (23).

(23) Bring me the lion’s sword.

(23) has two available interpretations. On the untransferred interpretation of ‘the lion’, it requests the sword that is embedded in the slain lion. On the transferred interpretation, it requests the sword that was wielded by the lion-shielded warrior (and is now embedded in the slain bear). The important observation for our purposes is that the resolution of the possessive is determined by whether the

definite description has a transferred or untransferred meaning. When ‘the lion’ has a transferred meaning, the possessive must express the *wielded by* relation, and cannot the *embedded in* relation (unless we enrich the context).

§2.3 Presupposition and Uniqueness

Presupposition tracks transferred meaning. To give just two examples, consider (24) and (25).

- (24) I know that the ham sandwich left without paying.
- (25) The shirt and tie stopped eating.

The presupposition of (24) is that the ham sandwich eater left without paying, and the presupposition of (25) is that the person wearing the shirt and tie was eating before. The fact that presupposition tracks transferred meaning is consistent across triggers.¹⁶

Whether or not it is ultimately treated as presuppositional, the uniqueness associated with singular definite descriptions tracks transferred meaning. Consider a case in which one diner consumes five ham sandwiches, while another diner consumes five tofu sandwiches. In that case, many find (1) acceptable despite the fact that there is neither a single, nor a single most salient ham sandwich.¹⁷ This is easily accounted for if what’s required is that there is a unique ham sandwich orderer, which there can be even when there are multiple ham sandwiches.

§2.4 Using the Properties to Identify the Locus of Meaning Transfer

We’ve been assuming that the locus of transfer in (1) is ‘the ham sandwich’, while the locus in (2) is ‘was a ham sandwich’. We can now support this claim.

- (1) The ham sandwich left without paying!
- (2) Bill was a ham sandwich.

Begin with (2) and consider two hypotheses. On the first hypothesis, the meaning of ‘Bill’ is deferred to mean *Bill’s dish*. On the second hypothesis, the meaning of ‘ham sandwich’ is deferred to mean *ham sandwich eater*.

To see that the second hypothesis is correct, note that on the first hypothesis we’d predict that Bill’s dish is available for subsequent anaphoric reference. However, this is clearly not the case, as evidenced by the unacceptability of (26).

- (26) #Bill was a ham sandwich, and it was delicious!

The other properties of meaning transfer confirm this. For instance, note that (27) is perfectly acceptable while (28) is not (and the discussion of (6) and (7) above provides a similar argument utilising gender marking):

(27) Bill, who wore a green shirt, was a ham sandwich.

(28) #Bill, which was toasted, was a ham sandwich.

Now consider (1). Again, there are two hypotheses: that the meaning of ‘the ham sandwich’ is transferred, or that the meaning of ‘left without paying’ is transferred. If the latter is correct, then the sandwich itself would be available for subsequent anaphoric reference. However, as we’ve already seen, it is not.

(29) #The ham sandwich left without paying and it was delicious.

So, ‘the ham sandwich’ must be transferred. However, this hypothesis itself has various forms. Is the entire definite description transferred, or merely one of the general terms embedded in it? In §4.3 we will demonstrate that meaning transfer can take different scopes, so both of these options should be available. In this particular case the data from number marking suggests that, at least on one salient interpretation, the transfer is below the definite article. That’s because, as we noted in §2.3, the uniqueness constraint on the definite is arguably satisfied even when there are multiple ham sandwiches, as long as there is a single ham sandwich orderer. Given all of this, the best hypothesis is that ‘ham sandwich’ is transferred to mean *ham sandwich orderer*, while ‘the’ has its ordinary non-transferred meaning. (Of course, one could claim that ‘sandwich’ has a transferred meaning while ‘ham’ doesn’t, but we fail to see how to secure the intended interpretation in a non-*ad-hoc* manner given this hypothesis.)

§3. Against Reducing Meaning Transfer to Other Linguistic Phenomena

A natural thought is to try to reduce meaning transfer to another figure of speech. In this section, we argue against the most salient reductions. In particular, we argue that meaning transfer cannot be reduced to loose speech, metaphor, or property inheritance (wherein one entity instantiates a property in virtue of inheriting it from another; details and examples to come), and we demonstrate that meaning transfer is not generally an instance of conversational implicature.

§3.1 Loose Speech

Familiar ordinary assertions don’t require perfect precision. Consider an ordinary utterance of (30) (Lasersohn 1999: 522):

(30) Mary arrived at three o'clock.

In most contexts, (30) is perfectly appropriate even if Mary arrived at 3:01. Such lack of precision is exemplified by a variety of utterances containing a variety of linguistic items. Number words, measure expressions, quantificational expressions, and scalar adjectives are perhaps the most familiar. The general idea is that when a sentence like (30) is uttered loosely, there is a group of propositions that are close-enough to the strictest interpretation of (30) such that the utterance is appropriate just in case any of these propositions is true.

One thought, then, is that just as we expand the extension of 'three o'clock' in (30) to include some other times, we expand the extension of 'the ham sandwich' in (1) to include the ham-sandwich orderer. On this view, meaning transfer is just another instance of loose speech.

Even before delving into the details, it is hard to see how to make this idea work. After all, the envisioned utterance of (1) communicates a proposition that is in no sense 'nearby' the proposition semantically expressed by (1). In its strictest interpretation, the latter is about foodstuffs while the former is about a person. We can bolster this general impression by considering, first, extant leading theories of loose speech and, second, the linguistic hallmarks of loose speech. We will show that meaning transfer isn't covered by such theories and doesn't exhibit the hallmarks.

In our brief gloss of the nature of loose speech, we were non-committal about whether it is a semantic or pragmatic phenomenon. Extant theories of loose speech are divided on this. On Lasersohn's (1999) view, loose speech is pragmatic. On Krifka's (2007) view, loose speech is semantic. We'll focus on Lasersohn's view, but our arguments generalize to Krifka's.

According to Lasersohn's view, (30) is false if Mary arrived at 3:01. However, it is close-enough to the truth in order to be acceptable. Closeness to the truth, in turn, is determined by the *pragmatic halo* associated with the utterance of (30). This halo is determined compositionally. Each sub-sentential constituent of (30) is associated with a range of values that are "close enough . . . that the difference between them . . . is ignored in context" (Lasersohn 1999: 528). We can imagine, for instance, that in the context envisioned, five minutes doesn't make a difference. In that case, the pragmatic halo of 'three o'clock' will be the set of times between 2:55 and 3:05, inclusive. We then compose these halo values with each other to determine the pragmatic halo of the entire sentence. Given that the proposition that Mary arrived at 3:01 is in the halo, the utterance of (30) is acceptable.

Lasersohn's theory cannot be applied to cases of meaning transfer. The reason is straightforward: transferred meanings are not close enough to non-transferred meanings to be included in the halo. It is easy to construct a scenario in which the difference between 3:00 and 3:01 is pragmatically irrelevant, but it is nearly impossible to construct a scenario in which the difference between a ham sandwich and a person is pragmatically irrelevant!¹⁸

Thus far, we've argued from a relatively abstract perspective that meaning transfer cannot be assimilated to loose speech. We will now focus on two linguistic hallmarks of loose speech and show that they do not apply to meaning transfer.

Consider, first, the fact that we can regulate loose speech with a range of different terms like 'exactly'.

(31) Mary arrived at exactly three o'clock.

Given the presence of 'exactly', (31) no longer expresses (semantically or pragmatically) that Mary arrived close to three. Rather, the speaker expresses a falsehood if Mary arrived at 3:01. Such terms, which Lasersohn calls 'slack regulators', do not affect meaning transfer. Consider, for instance, a scenario in which Bill ordered a mac n' cheese, albeit a slightly unconventional version of the dish. Two pedantic foodies have the following conversation:

(32) A: Bill was a mac n' cheese.
B: No way! That wasn't a real mac n' cheese.
A: What?! Bill was *exactly* a mac n' cheese!

Assuming that the dish Bill ordered really does satisfy the criteria for being a mac n' cheese, A's second utterance in (32) communicates something true. This is contrary to the prediction that we would get if meaning transfer were an instance of loose speech, in which case A's utterance could only be interpreted as saying that Bill was literally a bit of food.

Consider, second, Solt's (2014) observation that loose interpretations are not available when the relevant term appears in a comparative construction.

(33) Mary arrived later than three o'clock.

If "three o'clock" is interpreted loosely, then (33) would express a falsehood when Mary arrives at 3:01. However, it seems to only express a truth in that scenario. Loose interpretation, in other words, is blocked by the comparative. Meaning transfer, by contrast, is not blocked by the comparative.

(34) Mary tipped more than the ham sandwich.

The natural reading of (34) is one on which Mary tipped more than the ham-sandwich eater, not one on which Mary tipped more than the inanimate object eaten. However, if the comparative blocked meaning transfer, only the latter would be available.

§3.2 Metaphor

One figure of speech that seems hard to distinguish from meaning transfer is metaphor. Compare metaphorical utterances of (35) and (36) with the meaning-transfer readings of (2) and (1):

- (35) Juliet is a shining star. (Romeo describing his lover Juliet.)
- (36) The shining star has arrived at my doorstep. (Romeo describing the event of his lover, Juliet, arriving at his doorstep).

- (2) Bill was a ham sandwich
- (1) The ham sandwich left without paying.

There seem to be close similarities between these two figures of speech. In each, we have an utterance that is used to communicate a different proposition than we get from straightforward literal semantic interpretation.¹⁹ There are further structural similarities. Speaking roughly, in (35) and (2) the locus of figurative speech is the predicate. The proposition communicated is one which assigns to the subject a property that is not literally expressed by the predicate, but nevertheless somehow associated with it. In (36) and (1) the locus of figurative interpretation is the subject term: the definite description is used to pick out an object that does not literally satisfy the description. Indeed, the close similarity between the two figures of speech raises the worry that these are not really two separate phenomena. As we put it earlier, one may worry that instances of meaning transfer are just boring metaphors. We'll now argue that this worry is unfounded. Our argument does not hinge on a general definition of either phenomenon. Rather, we demonstrate that paradigmatic cases of metaphor and meaning transfer exhibit a range of distinct linguistic behaviours.

A first difference between metaphor and meaning transfer is that metaphorical meanings can often be extended and elaborated on.²⁰ Thus it would perfectly acceptable to extend (35) as follows:

- (37) Juliet is a shining star. When she shines brightly, her light shows me the way. She is also warm and burning with fire.

On the other hand, we cannot extend meaning transfer sentences in a similar way. For example, we cannot extend (2) as follows:

- (38) Bill was a ham sandwich. #He was crunchy on the outside, but soft on the inside.

Admittedly, occurrences of meaning transfer can be extended in a slightly different manner. For example, if a restaurant has two kinds of ham sandwiches on the menu, one sweet-and-sour and one spicy, the waiter might extend (2) by

saying, 'Bill was a ham sandwich. He was a spicy ham sandwich', or, 'Bill was a ham sandwich. A spicy one.' However, in these cases the waiter is predicating further things of the sandwich, not of Bill (i.e. the waiter cannot use an elaboration such as 'He is spicy' in this context).²¹

Indeed, a way to bring out the contrast is to compare two uses of the same sentence, each used to express a different figure of speech. Thus consider the following two uses of (35):

(Metaphor) Context: Romeo describing his lover Juliet.

Utterance: Juliet is a shining star. When she shines brightly, her light shows me the way. She is also warm and burning with fire.

(Meaning transfer) Context: A father is attending the school play, where each child is wearing a shirt with a certain object printed on it, and his daughter is wearing a shirt with a star on it. Explaining to the parent sitting next to him who his daughter is, he utters the following.

Utterance: Juliet is a shining star. She is the blue star with the sparkles.

A second difference between metaphor and meaning transfer is their amenability to covering several senses with a single expression. Let us first consider cases where a single expression is used to predicate a property of one subject term literally, and of another only figuratively. While it is perhaps more strained than the initial metaphor, we can arguably accept an utterances of (39) and (40), intending to communicate that the sun is literally a bright star, while Juliet is metaphorically a bright star.

(39) Both the sun and Juliet are bright stars.

(40) The sun is a bright star. So is Juliet.

In contrast, it is extremely hard to imagine a context in which the following are used to express that the thing on the plate is literally a ham sandwich, while Bill ordered a ham sandwich.

(41) #Both the thing on this plate and Bill are ham sandwiches.

(42) #The thing on the plate is a sandwich. So is Bill.

Next, consider the case of predicating the same expression metaphorically of two different subjects, but intending a very different metaphorical interpretation in each case. Again, while perhaps requiring some contextual priming, we can imagine acceptable utterances of (43).

(43) Both Juliet and Achilles are bright stars, each in their own way.

By contrast, imagine a lunch meeting at the conference of the association of sandwich makers. In this context, it is salient both that each member has ordered a particular sandwich and that each member specialises in producing a particular sandwich. Despite these salience features, it is very hard to hear a reading of (44), which is intended to communicate that John ordered a ham sandwich and that Jane is a producer of ham sandwiches:

(44) *Both John and Jane are ham sandwiches, each in their own way.

A third difference between metaphor and meaning transfer is that the agreement behaviour exhibited by meaning transfer (as outlined in §2) is not obligatory for metaphor. Recall that, in cases of meaning transfer, only the transferred meaning is available for subsequent anaphoric reference, and that this manifests itself in person agreement: when the transferred meaning is animate, we must use a personal pronoun. Metaphor does not share this property. For example, someone can metaphorically describe their tantruming toddler daughter as follows:

(45) The wild beast is at it again. It is tearing apart the living room apart.

(45) is acceptable even though we are using ‘The wild beast’ to metaphorically describe a person. In that case, the subsequent anaphoric pronoun ‘it’ can be used to pick out that same person as long as we continue to engage in the metaphor, i.e. to treat them as a wild beast. By contrast, ‘it’ in (46) cannot acceptably pick out the ham-sandwich orderer.

(46) * The ham sandwich left without paying; it knocked over a stool on the way out.

Note that, in the case of metaphor, there is flexibility. In (45), we can use the impersonal ‘it’ or the person-marked ‘she’. This flexibility is related to our earlier observation that metaphor can be extended in a way that meaning transfer cannot. When we extend the metaphor from the first clause of (45), we continue to treat a person as a wild beast, and so the impersonal pronoun is acceptable. Continuation is, of course, not obligatory; that’s why there is an acceptable variant of (45) that utilizes person-marked pronouns. Meaning transfer, by contrast, cannot be extended in the same way; and that’s why the only available anaphoric pronouns are those that agree with the transferred meaning of their antecedent.

§3.3 Property Inheritance

We’ve pushed the table against the wall. Our assertion of (47) expresses a truth.

(47) The table is touching the wall.

Not every part of the table is touching the wall (consider an arbitrary portion in its centre). How is it, then, that (47) expresses a truth? One plausible view is that ‘is touching a wall’ designates a property that can be instantiated in different ways. It is instantiated by the edge of the table because there’s no space between any of the edge and the wall. It is also instantiated by the table because it has a part—the edge—such that there is no space between any of that part and the wall. To use our favoured terminology, the table *inherits* the property of touching the wall from one of its parts: the edge. There is just one property that both the table and its edge instantiate in different ways.

Once identified, it becomes plausible that property inheritance is widespread. Here’s another potential example:

(48) The apple is red.

The apple itself has the property of *being red* in virtue of its skin having that property.

Yet another potential example: in other work (Liebesman and Magidor 2017, Liebesman 2011) we argue that types often inherit properties from their tokens and that this can be used to explain some otherwise puzzling constructions. Thus for example, it might be true that *War and Peace* (a “type” book) has the property of being on the shelf, in virtue of the fact that one of its physical instantiations is on the shelf.

In cases of property inheritance, a predicate designates a property that can be instantiated in different ways. A hallmark of property inheritance, then, is that we can ascribe that property to an object that instantiates it in one way, and then, via anaphora or elision, ascribe the very same property to an object that instantiates it in a different way. This is illustrated in (49) and (50)

(49) The table is touching the wall. So is the edge of this chair.

(50) The apple is red. So is this pen.

Property inheritance is relevant to our discussion because it presents an alternative to meaning transfer. We’ve been defending the view that (1) is an instance of meaning transfer and that, in particular, the meaning of ‘the ham sandwich’ is transferred, such that it expresses *the ham sandwich orderer*. Property inheritance, however, provides us with two other salient hypotheses. The first is that ‘left without paying’ expresses a property that can be instantiated by people, but also by sandwiches in virtue of their orderers leaving without paying. The second is that ‘ham sandwich’ expresses a property that can be instantiated by sandwiches, but also by people in virtue of ordering sandwiches.

The challenge is to justify treating (1) as a case of meaning transfer, rather than as a case of property inheritance. The properties of meaning transfer that

we identified in §2 allow us to meet this challenge. Begin with the first hypothesis: that ‘left without paying’ expresses a property that can be inherited by sandwiches (in virtue of being ordered by a person that left without paying). In that case, ‘the ham sandwich’ in (1) would designate a ham sandwich. However, we’ve already demonstrated that this is false. After all, if ‘the ham sandwich’ designates the ham sandwich, and does not undergo meaning transfer, then the ham sandwich itself would be available for subsequent anaphoric reference, but it is not, as evidenced by (19).

(19) The ham sandwich left without paying; #it was delicious.

The second hypothesis is that ‘ham sandwich’ expresses a property that can be instantiated both by sandwiches and by people. If this hypothesis is correct, we should be able to use ‘ham sandwich’ to ascribe that single property to both things at once. However, we cannot, as evidenced by (41).

(41) #Both the thing on this plate and Bill are ham sandwiches.

Stepping back, we can clearly distinguish property inheritance and meaning transfer using the sub-sentential effects of meaning transfer we identified in §2. So, we have no reason to think that familiar instances of meaning transfer can be reduced to property inheritance.

§3.4 Conversational Implicature

Another initially appealing hypothesis is that meaning transfer is an instance of conversational implicature.²² The proposal is clear: when a waiter utters (1), the proposition literally expressed by the sentence is trivially false and thus violates Grice’s maxim of quality.²³ Participants in the conversation therefore assume that the speaker is intending to communicate a proposition other than the one literally expressed (in this case, the proposition that the customer who ordered a ham sandwich left without paying).

At a first pass, one might hope to test this hypothesis using standard tests for conversational implicature. Two main ways to distinguish conversational implicature from semantic content are cancelability and reinforcingability.²⁴ Consider an uncontroversial case of conversational implicature:

(51) Bill: Is Jane coming to dinner?

Jill: Jane is working until late today.

Jill’s utterance in (51) merely conversationally implicates, but does not semantically entail, that Jane is not coming to dinner. One way to see this is that

the implicature is *cancellable* – it is felicitous to follow Jill’s utterance with the negation of the alleged implicature, as in (52). Another is that it is *reinforcible*. It is felicitous to follow Jill’s utterance by stating the implicature explicitly, as in (53):

(52) Jill: Jane is working late today, but she will come to dinner after work.

(53) Jill: Jane is working late today, so no – she is not coming to dinner.

This contrasts with a case where Jill responds to the question with the explicit ‘No, she is not coming to dinner’ in which case an attempted cancellation (‘but she is coming to dinner’) would be contradictory and an attempted reinforcement (‘so she’s not coming to dinner’) would be redundant.

What do these tests suggest with respect to meaning transfer? It might seem that meaning transfer is not cancelable, as (54) is highly infelicitous:

(54) The ham sandwich left without paying. #But it’s not the customer who ordered the ham sandwich who left without paying; rather, it is the sandwich itself.

However, one can respond that this is not because the cancellation clause contradicts the original semantic content, but rather because the semantic content of that clause is itself an infelicitous category mistake, and so cannot be the intended meaning. This line of thought is supported by the fact that, if we choose an example involving a predicate that can naturally apply to both sandwiches and people, the cancellation clause seems entirely acceptable:

(55) The ham sandwich is quite interesting. It’s not the customer who ordered the ham sandwich who is interesting, but rather the sandwich itself.

Exactly parallel remarks apply to reinforcibility. While (56) is infelicitous, the apparent reinforcement in (57) is entirely acceptable:

(56) The ham sandwich left without paying. #That is, the customer who ordered the ham sandwich left without paying, not the sandwich itself.

(57) The ham sandwich is quite interesting. That is, the customer who ordered the ham sandwich is quite interesting, not the sandwich itself.

So far, it seems that, as long as we focus on the right examples, transferred meaning is both cancellable and reinforcible. However, Sadock (1978) persuasively argues that neither test is of much use in cases where the relevant expression may be *semantically ambiguous* (whether lexically or structurally). This is because, in any case of ambiguity, we can use cancellation to set aside an unintended semantic meaning (as in (58)) and use reinforcement to highlight the intended meaning (as in (59)):

- (58) Visiting relatives can be annoying. That is, it is not annoying for you to visit relatives- it's the relatives that visit you that can be annoying.
- (59) Visiting relatives can be annoying. That is, the relatives that visit you can be annoying.

This problem is central to the case we are interested in here: we are attempting to discern whether transferred meaning is a type of semantic content (loosely construed). However, no one who holds that hypothesis is going to require meaning transfer to be obligatory (at least where no category mistakes are in sight). Thus, even if the transferred interpretation of 'The ham sandwich is quite interesting' is one available semantic content, clearly the interpretation on which meaning transfer is not in play and the sentence simply concerns food is also available. Just as with (58) and (59), (55) can be used to clarify that we are intending the sentence to be read with the non-transferred interpretation and (57) to clarify that we are intending it to be read with the other, transferred interpretation. So far we have reached an impasse: the results of the cancellability and reinforcibility tests are inconclusive.

However, an entirely different set of considerations tell against the hypothesis that meaning transfer is derived via implicature.²⁵ Conversational implicature operates by first calculating the proposition that the sentence semantically expresses as a whole, and only then deriving the implicature. But as we have seen in §2, meaning transfer has many sub-sentential effects, including gender marking, number marking, and availability for anaphoric reference.²⁶ This is an instance of a more general worry, stressed by Recanati (1993) and others: that Grice's account cannot explain cases in which the output of a pragmatic inference is embedded within a larger clause.²⁷

Recall, for example, the case where a waiter describes a single customer who stormed out of the restaurant by saying, 'The shorts and tie is leaving without paying!'. If this were a case of standard conversational implicature, it would be very hard to explain the use of the unconventional number marking. If, in the first instance, the waiter simply expressed the proposition that some items of clothing left without paying, then that proposition should have been expressed as 'The shorts and tie *are* leaving without paying'. Similar remarks hold for other gender and number marking cases we discussed.²⁸

A similar point applies to anaphoric reference. If, in the first instance, we are expressing the proposition literally expressed by the sentence, then components of that proposition should be available for anaphoric reference. Consider the classic Gricean example of the professor writing a reference letter for her student, and uttering, 'Jack has excellent handwriting' to implicate that Jack is not a very good philosopher. In this context, it is possible to continue the utterance with further anaphoric reference to the handwriting ('It is cursive and neatly on the line'). By contrast, as we have seen, it is not possible to continue 'The ham sandwich left without paying' with anaphoric reference to the sandwich ('It has lots of mustard on it').

There are numerous reactions to this problem,²⁹ but the one that interests us extends a broadly Gricean notion of implicature to sub-sentential constituents (Simons 2009, 2017). Her idea is that Gricean reasoning that begins at the level of the proposition expressed by an utterance can then go on to have local effects concerning the interpretation of constituents of the sentence uttered.

How does this affect the relationship between meaning transfer and implicature? Our argument thus far was that the sub-sentential effects of meaning transfer show that it cannot be derived via implicature. However, if implicatures can themselves be sub-sentential, this reasoning falters. The question then becomes whether sub-sentential views of implicature can accommodate meaning transfer.

Debates concerning this question remain unresolved and we won't weigh in here.³⁰ However, even if some subsequent developments in sub-sentential implicature would allow a treatment of meaning transfer, we would not take these to challenge the core claim of this section. These theories are sufficiently significant departures from Grice's original view that our main conclusion in this subsection stands. Meaning transfer cannot be derived via implicature, on a traditional understanding of the notion.

§4. The Implementation of Meaning transfer

Having argued that meaning transfer cannot be reduced to familiar linguistic mechanisms, we're left with a question: what is the linguistic implementation of meaning transfer? When we use sentences such as (1) to express the proposition that the ham sandwich-orderer left without paying, what is the structure and makeup of this proposition, and how do we succeed in expressing it?

This is an enormous question with multiple facets not all of which we will address. But we make some headway as follows. We begin in §4.1 by discussing one hypothesis concerning the nature of meaning transfer: *the ellipsis view*. According to the ellipsis view, meaning transfer is primarily a syntactic phenomenon, where instances of meaning transfer are simply elliptical for their explicit counterparts. That is, a meaning transfer occurrence of (1) has the same syntactic and semantic structure as an ordinary occurrence of 'The ham sandwich-orderer left without paying'. We argue that the ellipsis view ought to be rejected. The upshot is that meaning transfer is achieved (just as the name suggests) via some operation on *meanings*. In §4.2, we hypothesize a meaning transfer operator which we label *MT*, though we are largely neutral on its nature, including whether it is to be understood as semantic, pragmatic, or meta-semantic.³¹ We argue that *MT* is a highly flexible operator: it is context-sensitive, operates at a sub-sentential level, and has no fixed semantic type. In §4.3, we defend a more specific claim concerning *MT*: that it is a non-shiftable and non-bindable operator. We conclude that this provides at least tentative support for the claim that *MT* is *content-invisible*: the proposition expressed by utterances involving meaning transfer contains only

the output transferred meaning (i.e. the result of applying *MT* to its argument), with no constituent of this proposition corresponding to *MT* itself.

Note that our discussion in this section does not directly tackle questions about the role of the proposition expressed by sentences involving meaning transfer in a larger linguistic theory – i.e. whether this proposition qualifies as semantic content, conventional meaning, what is said, etc. (To be clear, our use of the term ‘proposition expressed’ to designate our object of study is intended to be neutral on these questions.). We’ll revisit these issues in §5, where we will see that, while by no means settling these issues, our conclusions in §4 have significant bearing on them.

§4.1 The Ellipsis View

We start by discussing one hypothesis concerning the nature of meaning transfer: *the ellipsis view*. According to the ellipsis view, meaning transfer does not (directly) transfer meanings but is primarily a syntactic phenomenon: in cases of meaning transfer, the uttered sentence has exactly the same syntactic structure as that of its corresponding non-meaning transfer sentence. For example, the occurrence of (1) has the same syntactic structure that ‘The ham sandwich orderer left without paying’ (or some variant) has in ordinary contexts.

There are several reasons we don’t find such a view promising.³²

First, while the ellipsis view may seem natural for (1), it is much less natural in other cases. In particular, some cases would require the elliptical form to depart massively from surface form. In this way, the would-be ellipsis would differ from more familiar (and well-motivated) ellipsis as occurs in, for example, VP ellipsis and sluicing. Consider an utterance of (60) which expresses the claim that the faculty member from Magdalen college always arrives late. In this case, “the faculty member from” would have to be elided, and there is no precedent for a process that could elide this non-constituent.

(60) Magdalen college always arrives late.

Other examples require even further departures from the original structure. For instance, imagine that there is a school event in which students are holding up cards. Each card has a letter printed on it, and each letter is (in a way familiar to everyone in the audience) code for a certain shape. Suppose that a certain student is holding a card with a ‘c’, which we all know is code for ‘circle’. We can then refer to the student as follows:

(61) The circle is stealing the show!

But on the ellipsis view, (61) is presumably elliptical for something like ‘The student who is holding a card which has a letter on it which is code for circle’, in which the word ‘circle’ occupies a very different syntactic position than in (61).

Second, there are straightforward reasons to think that sentences exhibiting meaning transfer are inequivalent to their would-be elliptical counterparts. Consider a context in which we know the names of the dogs in the park but not the names of their owners, and use the dogs to identify the owners (either explicitly as in (62) or via meaning transfer as in (63)). Note that the second clause is licenced in (62) but not in (63), showing that the two initial sentences are not linguistically equivalent.³³ A similar contrast applies for the pairs (64) and (65), and (66) and (67).³⁴

(62) Fido's owner is getting angry. Spot's is too.

(63) Fido is getting angry. # Spot's is too.

(64) The customer who ordered the ham sandwich left without paying, as such customers often do.

(65) The ham sandwich left without paying, #as such customers often do.

(66) The ham sandwich orderer left without paying, as such orderers often do.

(67) The ham sandwich left without paying, #as such orderers often do.

Of course, each such pair only tells against a specific hypothesis for the alleged ellipsis, but the range of examples suggests the ellipsis view is on the wrong track.

Third, a standard constraint on natural language ellipsis is that the elided constituent has a linguistic antecedent.³⁵ However, none of the examples of meaning transfer have depended on the existence of a linguistic antecedent.³⁶

Fourth, in addition to cases such as (61), where the full sentence has a very different structure than the original, there are more extreme cases where there does not seem to be any natural way to spell out what the meaning transfer sentence is supposed to be elliptical for. To see this, consider a case discussed in Magidor (2017: 72–73). Imagine a restaurant in which many customers are eating ham sandwiches in a variety of sizes. Only two of those customers are impatient: Jill and Jane. Jill is eating a ham sandwich that is large relative to Jane's sandwich, but isn't large relative to the totality of ham sandwiches eaten around the restaurant. Jill storms out of the restaurant. Consider the following conversation between waiters A and B:

(68) A: The impatient ham sandwich left without paying!

B: Which impatient ham sandwich? There were two of them . . .

A: The large impatient ham sandwich!

Although A's choice of expression is somewhat convoluted, it is not too hard to understand what A is expressing. However, note that spelling this out as one explicit English sentence is quite tricky (we would need something like: 'The impatient customer who was eating a ham sandwich that was large relative to

the sandwiches that were ordered by impatient customers eating ham sandwiches left without paying' – but this is by no means a natural construal).

In light of these considerations, we conclude that the ellipsis view ought to be rejected.

§4.2 The Flexibility of *MT*

The propositions expressed in instances of meaning transfer are clearly somehow connected to the words explicitly used in the sentence uttered. With the ellipsis view rejected, the remaining hypothesis is that this proposition is somehow generated as a function of the original *meanings*. To that end, we postulate an operator – which we label *MT* - taking us from original meanings to their transferred interpretation. As a starting point, we use the term 'operator' loosely, remaining neutral on a number of key questions: should this "operator" be thought of as semantic or meta-semantic; is there any element corresponding to *MT* in either the syntactic structure of or in the proposition expressed by a sentence involving meaning transfer; and what sorts of arguments does this operator take and what sorts of outputs does it have?

Being less neutral, we now want to argue that *MT* is a flexible operator in at least the following three senses. First, *MT* is clearly context sensitive: it contributes different functions in different contexts. For example, while we have been focusing on contexts in which (2) expresses the proposition that Bill was a ham sandwich orderer, we have seen that in other contexts (e.g. that of the sandwich makers association conference), the sentence might express the proposition that Bill was a manufacturer of ham sandwiches.³⁷ Second, we maintain (§4.2.1) that *MT* can operate at a sub-sentential level, with the ability to take different scopes and affect numerous sub-sentential phenomena. This is in contrast to the simpler hypothesis, according to which *MT* always contributes a function from propositions to propositions. Third, we maintain (§4.2.2) that *MT* has no fixed semantic type: it receives as arguments, and outputs as values, elements of a variety of different semantic types.

§4.2.1 Scope and Sub-sentential Effects Like us, Nunberg (1995: 119) argues that *MT* can take multiple different sub-sentential scopes. (69) demonstrates the point:

(69) The attractive sandwich left without paying.

In the following cases, the most salient proposition expressed by (69) requires *MT* to scope over 'attractive sandwich' and 'sandwich' (but under 'attractive'), respectively.³⁸

First, imagine that the kitchen has produced several sandwiches, and one of them is attractive while the rest are unattractive. The person who ordered the attractive sandwich left in a huff. In this case, the most salient reading of (69) transfers ‘attractive sandwich’ to *orderer of an attractive sandwich*.

Second, imagine that several people ordered sandwiches, and these sandwiches are all unattractive. One of the orderers is attractive, while the rest are unattractive. The attractive sandwich-orderer left in a huff. In this case, the most salient reading of (69) transfers ‘sandwich’ to *sandwich-orderer*. Since there are no attractive sandwiches, *MT* cannot scope over ‘attractive’; shifting ‘attractive sandwich’ to *orderer of an attractive sandwich* would yield a non-designating description.³⁹

Admittedly, such examples don’t conclusively establish that *MT* takes sub-sentential scope. After all, one could identify different propositional functions for each of these cases. However, such a hypothesis misses the crucial observation, that the contribution to the propositions expressed in these cases can be sourced to particular constituents. Furthermore, as we observed extensively in §2, *MT* can affect the linguistic properties of particular sub-sentential constituents, like gender/number/animacy marking, presupposition, and anaphoric availability. Together, these features make a strong case for the sub-sentential availability of *MT*.

§4.2.2 Flexible Typing What is the semantic type of the functions contributed by *MT*? In the examples we’ve been discussing, *MT* seems to be type $\langle\langle e,t\rangle, \langle e,t\rangle\rangle$, e.g. when it transfers *ham sandwich* to *ham sandwich orderer*. However, we’ve also seen examples in which it seems to be $\langle e,e\rangle$, for instance, when we use ‘Spot’ to designate Spot’s owner.

(70) Spot is looking kind of glum today.

There’s a complication with this example: the view that names are actually type $\langle e,t\rangle$ has been become increasingly popular.⁴⁰ However, the general observation is that meaning transfer is a massively flexible phenomenon that can be used on just about any linguistic item, given enough context.⁴¹

In all of the examples we’ve considered thus far, *MT* picks out a contextually specified function that has the same domain and range. Must it?⁴² It is perhaps not surprising that most familiar examples have this property. After all, familiar instances of meaning transfer generally use sentences that are semantically interpretable without transfer. In such cases, an easy way to ensure interpretability is to preserve types. However, there are some putative examples that plausibly lack this property. (None of these examples is uncontroversial.) First, imagine that two logic students are working hard, one on the existential quantifier and the other on the universal. We can then use ‘some’ and ‘every’, with transferred senses, to refer to the students themselves.

(71) Some is working hard today, but not as hard as every.

Though the type of proper names is controversial, it is uncontroversial that they are not type $\langle\langle et \rangle, \langle et \rangle, t \rangle\rangle$, the type of unrestricted quantifiers. So, this is a potential instance of cross-type meaning transfer.⁴³ Second, consider some spies who agree to use a single code word to express the proposition that the enemy has exposed their plot. It not obvious whether this is a case of meaning transfer, but if it is, it involves transfer from a word type to a sentence type.

§4.3 Shifting and binding

Consider the context-sensitive term ‘local’. In the simplest cases, ‘local’ is used to denote the location provided by the context of utterance.⁴⁴ Thus, if Jill is currently in London and says, ‘I am attending a local conference’, she (typically) expresses the proposition that she is attending a conference in London. However, as has been widely observed, the value of ‘local’ can be shifted by modal, temporal, and intensional operators, and be bound by quantifiers. Thus, imagine that Jill (in London) utters the following:

(72) When I lived in Paris, the Eiffel Tower was local.

(73) Had I moved to Paris, the Eiffel Tower would have been local.

Clearly, ‘local’ in these sentences is picking out a location in Paris: (72) doesn’t express that when Jane lived in Paris, the Eiffel Tower was in London, and (73) doesn’t express that had Jane moved to Paris, the Eiffel Tower would have been moved to London. A natural explanation for why ‘local’ receives these readings is that the temporal and modal operators shift the value of ‘local’. Admittedly, there is an alternative hypothesis, according to which the context of utterance in (72) and (73) (especially the linguistic context, and the mention of ‘Paris’ in this case) makes Paris, rather than London, the contextually supplied location. However, one reason this alternative hypothesis does not look very promising is that other, non-shiftable indexicals, such as ‘here’, still seem to receive the ‘London’ location in this very context, as we can see from the following case:

(74) Had I moved to Paris, the Eiffel Tower would have been local, but it would have been harder to see my family here.

Moreover, even if we can account for the alleged shiftability of ‘local’ by taking the contextually supplied location to be non-standard, we cannot, in the same way, account for cases of binding where the indexical takes a range of different ‘values’ for the same utterance. Thus consider:

- (75) Both Jane and Kate went to a local conference.
- (76) Every reporter went to their local bar.

Note that there is a true and felicitous reading of (75), even when Jane and Kate are each in a different location, and similarly for (76), where every reporter is in a different city.

Interestingly, *MT* is neither shiftable nor bindable in these ways. To see this, recall two ways to transfer the meaning of ‘is a ham sandwich’. The first occurs in a context where we are waiters discussing what each person ordered. There, ‘is a ham sandwich’ can be transferred to mean *is a person who ordered a ham sandwich*. The second occurs in a context where we are at a food industry conference. There, ‘is a ham sandwich’ can be transferred to mean *is a person who makes ham sandwiches*. Suppose that we are currently in the context of waiters distributing food in a restaurant. If *MT* is shiftable, we should be able to shift from the first sort of transfer (ham sandwich orderer), to the second (ham sandwich maker). However, we can’t, as demonstrated by (77) and (78).

- (77) Jill is a cheese sandwich, but when she worked in the food industry, she was a ham sandwich.
- (78) Jill is a cheese sandwich, but had she worked in the food industry, she would have been a ham sandwich.

(77) does not get a reading where it expresses that Jill ordered a ham sandwich, and that when she worked in the food industry she manufactured cheese sandwiches. Nor does (78) get the analogous modal reading.

The same observation applies to binding. Consider the following:

- (79) Both Jill and Jane are ham sandwiches.

It is very hard to get a reading of (79) where it means that Jill (who is at a restaurant) ordered a ham sandwich and Jane (who works in the food industry) manufactures ham sandwiches.⁴⁵

The lack of shiftability of the *value* of *MT* should not be confused with a lack of shiftability of its *output*. Recall the use of ‘Magdalen College always arrives at the meeting late’, to mean ‘The faculty member working in Magdalen College always arrives at the meeting late’. In this case, the proper name ‘Magdalen College’ has plausibly been transferred to a definite description-type meaning, and there are clearly at least some uses on which this meaning is read as non-rigid and shiftable. For example, we can get a true reading of ‘Magdalen college has always arrived to the meeting last’ where we quantify over occasions when different people have been the person working in Magdalen College, and of ‘Jack believes that Magdalen college is the laziest person on the faculty’ where Jack

believes that Jane is the laziest person on the faculty, but falsely believes that Jane works in Magdalen college.⁴⁶

It is particularly interesting to contrast the non-shiftability of *MT* with the behaviour of a similar (but distinct) figure of speech: metaphor. One prominent view, defended by Stern (2000), takes metaphor to involve an operator similar to *MT*, which Stern labels ‘*Mthat*’. Just as with *MT*, Stern takes *Mthat* to be capable of scoping over any phrase (complex or simple) and of receiving different interpretations in different contexts. Stern himself maintains that *Mthat* is not shiftable. However, both Stern’s own discussion and subsequent discussion of his proposal (see especially Camp (2005)) suggest that, unlike *MT*, *Mthat* is shiftable and bindable in all the ways described above. Borrowing an example from Stern (2000: 207), consider (80):⁴⁷

(80) This television is no Rolls Royce. Maybe if Rolls Royce were to manufacture the worst cars ever this television would have been a Rolls Royce.

This sentence can have a true reading on which it means that the television does not have the properties that are actually metaphorically associated with ‘Rolls Royce’ (being high quality), but had we associated different properties with Rolls Royce, the television would have had *those* properties. And to look at an example adapted from Camp (2005):

(81) Juliet and Achilles are both a shining star, though in their own different ways.

This utterance can have a reading on which it means that Juliet is a star in the sense that she brightens her surroundings while Achilles is a star in the sense that he harms those who get too close. It is not our role in this paper to offer any particular theory of metaphor, but it’s clear that, if metaphorical interpretation involves a metaphorical operator such as *Mthat*, that operator, unlike the meaning transfer operator *MT*, is shiftable and bindable.

§4.3.1 Content Invisibility We have argued that *MT* is neither shiftable nor bindable. In §5 we will see some of the theoretical import of this claim to debates concerning the semantics/pragmatics distinction. We conclude by tentatively suggesting another potential implication: that *MT* is *content invisible*.⁴⁸ Assume that sentences express structured propositions. Content invisible expressions are those for which there is no corresponding element of the structured proposition corresponding to that expression. An analogy from the case of definite descriptions can help explain what we mean here. Suppose we accept that the definite article is of type $\langle\langle e,t \rangle, e \rangle$, i.e. a function from properties to individuals. In a framework of structured propositions, this still leaves open two possibilities. Consider a use of ‘the dog’ in a particular context where the unique salient dog is Emma. According to one hypothesis, the definite article is *content visible*: the

proposition expressed by ‘the dog’ contains both the function expressed by ‘the’ and the property of being a dog (or some such similar entities). According to an alternative hypothesis, the definite article is *content invisible*: the proposition expressed by ‘the dog’ consists of neither of these entities: its content is merely Emma herself. The claim that *MT* (whether or not it is represented in syntax) is content invisible corresponds to the latter hypothesis: the proposition expressed by $MT(\varphi)$ contains neither an entity corresponding to the meaning transfer function nor an entity corresponding to its argument, but the result of applying the function to the argument.⁴⁹

If content invisibility is correct, that would provide a simple explanation of the lack of ability to bind or shift *MT*: binding and shifting are typically thought to involve a variable-like element in the semantic content of the sentence that is shifted or bound. Admittedly, the lack of shifting and binding are not conclusive proof of content invisibility. For example, it is extremely implausible that the indexical ‘here’ is content invisible: it is hard to see how ‘here’ can affect content in the way it does without contributing a corresponding element (a location) to the propositions expressed (the fact that unlike *MT* it does not denote a function is pertinent here). Yet, it is easy to see that ‘here’ is neither shiftable nor bindable. For example, if Jill is in London, she cannot use ‘here’ in either of the following utterances to talk about Paris:

(82) When I lived in Paris, the Eiffel Tower used to be here.

(83) Had I moved to Paris, the Eiffel Tower would have been here.

And similarly, ‘here’ cannot receive covarying interpretations in the following:

(84) Every reporter went to a bar here.

Thus in principle, it’s possible that like ‘here’, *MT* is a content *visible*, yet unshiftable indexical. Having said that, it should be noted that the only uncontroversial examples in the literature of non-shiftable (and visible) indexicals involve ‘pure’ indexicals: indexicals such as ‘here’, ‘now’, ‘I’, and ‘today’ which arguably have a very simple linguistic rule. If *MT* is not content invisible, then, it would provide an exception to the rule.^{50, 51}

§5. Meaning Transfer and the Semantics/Pragmatics Distinction

Our discussion thus far raises a question: is meaning transfer semantic or pragmatic? There are many different views on how to draw the semantics/pragmatics distinction, and, to a certain extent, the question is terminological. However, there are some substantive issues that arise concerning how our view of meaning transfer interacts with the debate concerning the semantics/pragmatics distinction.⁵²

§5.1 Three Ways of Drawing the Distinction

As Stanley (2000) notes, there are at least three prominent ways to draw the distinction. According to the first, semantics concerns context-invariant aspects of meaning (call this ‘semantics as context invariance’). On this view, the semantic content of ‘That is a dog’ does not, on its own, determine a proposition, because ‘that’ does not have a context-invariant referent. According to the second view, semantics can involve some forms of context-sensitivity, but only those that are linguistically controlled. On this view, the semantic content of an utterance is a proposition that is derived by assigning contents to the various elements of the utterance’s logical form, and combining these according to rules of composition. In some cases, an element of logical form is an indexical expression, and extra-linguistic context is used to provide content for that element. For example, the logical form of ‘She is tall’ contains the indexical expression ‘she’, and if ‘she’ is used to pick out Jill, then the utterance semantically expresses the proposition that Jill is tall. (Call this conception ‘semantics as linguistically controlled content’). According to the third, semantics is concerned with assigning truth-conditional content to utterances. On this conception, when a sentence is uttered in context, it expresses a content which is used evaluate whether the utterance is true or false and which the speaker can be taken to be committed to. On this view, semantics is concerned precisely with the assignment of this content (call this ‘semantics as truth-conditions’).⁵³

Where does meaning transfer fall, on each of these three conceptions? It is uncontroversial that the interpretation of meaning transfer is sensitive to extra-linguistic context, and thus according to semantics as context invariance, meaning transfer is pragmatic. What about semantics as linguistically controlled content? In §4, we argued that *MT* (the meaning transfer operator) is context-sensitive and gave evidence that it is content invisible. As we have noted (fn. 49 above), though, these claims can be modelled in two very different ways. According to the first, *MT* is present at logical form (although only the output of *MT* as applied to a constituent is a component of the proposition expressed). According to the second, while the transferred meaning is represented in the proposition expressed, *MT* is not present at the level of logical form. Rather, meaning transfer makes it to the relevant proposition expressed directly, via an optional, sub-sentential ‘enhancement’ process (what Recanati (2003; 2004) calls a ‘primary pragmatic process’). Assuming semantics as linguistically controlled content, meaning transfer is semantic if the first model is correct, and pragmatic if the second is correct. (We ourselves wish to remain neutral between these two models of meaning transfer, but will return to the debate between them below.) Finally, what about semantics as truth-conditions? If our only notion of truth-conditions is that of linguistically controlled content, the question of whether meaning transfer is pragmatic or semantic collapses into the question of how the phenomenon ought to be classified according to semantics as linguistically controlled content. If, however, we have an independent notion of truth-conditions, then the

question remains open. Perhaps more than any other phenomenon that might fall on either side of this divide, meaning transfer has a strong claim to figure in the intuitive truth-conditions of the relevant utterances. Thus, even Stanley, who ultimately rejects (for reasons we will turn to shortly) the claim that meaning transfer figures in truth-conditions concedes that “Of all the constructions . . . the central worry for the semanticist is . . . the case of deferred reference. Not only is there is a strong intuition that the deferred meaning is part of the intuitive truth-conditions, but the deferred meaning enters into certain linguistic processes, such as anaphora and ellipsis” Stanley (2005: 227).⁵⁴

§5.2.1 A dilemma Thus far, we have merely suggested how meaning transfer would be classified according to each of a variety of conceptions of the semantics/pragmatics distinction. The phenomenon of meaning transfer does, however, raise some significant issues in the debate. To see this, we present three claims one might be tempted to hold. Surprisingly, in light of our observations in the previous sections, these claims are inconsistent:

- (i) Transferred meaning figures in the truth-conditions of the relevant utterances. (For example, when, in the relevant context, the waiter utters ‘The ham sandwich left without paying!’, the content they express is true if and only if a particular person—rather than a piece of food—left without paying.)
- (ii) Truth-conditions always coincide with linguistically controlled content.
- (iii) One must only posit the existence of covert variables in logical form when one has evidence of bound readings of these variables.

As noted above, and as Stanley himself emphasizes (2005), claim (i) is highly intuitive. Stanley defends claim (ii) in a series of papers.⁵⁵ As is by now familiar, if logical forms were to include only overt structure, then this claim would be obviously false. In a typical context, if one utters ‘It’s raining’ while in London, the truth-conditions of one’s utterance require that it’s raining in London (or perhaps in a more specific location), rather than that it is raining somewhere or other. And yet, there is no obvious overt element of this logical form corresponding to the location of London. Defenders of claim (ii) maintain that logical forms can contain covert or unpronounced material, specifically, covert indexicals, i.e. variables that can take different values in different contexts (in this case, a location variable which receives the value London in context).⁵⁶ However, objectors to (ii) might rightly complain that positing covert structure for every case of alleged discrepancy between truth-conditions and linguistically controlled content seems *ad hoc*, and effectively makes thesis (ii) an unfalsifiable empirical hypothesis. This is where thesis (iii) comes in: if we only posit covert structure in those cases where we have independent evidence of its existence, then the reliance on such structures cannot be accused of being ad-hoc. And, as

Stanley and Szabó (2000) argue, the central criterion for positing covert variables is the availability of readings on which those variables are bound.⁵⁷ Consider, for example, ‘Everywhere John goes, it rains’. In the relevant context, the sentence expresses the claim that, for every location *l* to which John goes, it rains in *l*. The crucial point is not that ‘it rains’ is interpreted as restricted to a particular location, but rather that the location in question varies in accordance with the quantifier. Since variable binding is (at least partially) a syntactic phenomenon, there needs to be a variable in logical form representing the location.

Thus each of (i)-(iii) might seem *prima facie* attractive. However, in light of our discussion of meaning transfer, these three claims are inconsistent. If transferred meaning figures in truth-conditions, then, for (ii) to be correct, there has to be a covert element of logical form corresponding to the meaning transfer operator. However, as we have argued in §4, if there is a meaning transfer operator, it is not available for binding and hence (iii) is violated.

§5.2.2 Stanley’s response Stanley’s proposal (2005:227–231) is to reject (i). We find his arguments for this rejection unpersuasive. Central to his arguments are reasons for maintaining that meaning transfer is not semantic in the sense of semantics as linguistically controlled content. Even if these arguments were successful, it is not clear why that would motivate rejecting (i) rather than (ii). If transferred meaning is part of intuitive truth-conditions, but not part of linguistically controlled content, then so much the worse for (ii)! (Indeed, one of Stanley’s chief opponents here is Recanti, who argues precisely that meaning transfer – as well as other phenomena – proves there is a gap between linguistically controlled content and intuitive truth-conditions, necessitating a pragmatic intrusion into truth-conditions). Moreover, we find his arguments in favour of the claim that transferred meaning isn’t part of linguistically controlled content unpersuasive.

Stanley’s first argument is that the semantic-value (understood as linguistically controlled content) of an expression has to be constrained by its conventional meaning, but “[i]f deferred reference were semantic, the denotation of ‘the ham sandwich’ would be something of which the predicate ‘ham sandwich’ were not true,” (2005:229). The problem with this argument is that it ignores the very possibility of additional covert structure: what matters isn’t the conventional meaning of ‘ham sandwich’ on its own, but rather the conventional meaning of the (context-sensitive) meaning transfer operator, as applied to the argument (‘ham sandwich’). Indeed, this point precisely parallels the observation that, although the conventional meaning of ‘It’s raining’ doesn’t include any mention of London, the content of ‘It’s raining in London’ may still be semantic in his sense.

Stanley’s second argument is that treating meaning transfer as semantic would entirely trivialise semantics: “since...virtually any word can have a

deferred meaning, it follows that any word could in principle acquire any meaning, *via* a sense-transfer function. So, the resulting semantic theory is one according to which semantic content is unconstrained by conventional meaning. The semantic content of the word ‘house’ could be the property of being a dog – the only thing that would prevent it from acquiring this semantic content is pragmatic facts about a context,” (2005: 230). But this complaint isn’t correct either. On the envisioned view, the semantic content of the word ‘house’ is *not* the property of being a dog. Rather, the semantic content of the complex expression ‘*MT*(house)’, can, in some (rather exotic) contexts, pick out the property of being a dog – but only if the logical form of the particular utterance contains *MT*, and the context allows the operator to take a function which maps the property of being a house to that of being a dog.⁵⁸

Finally, recognising that theoretical arguments for the claim that transferred meaning isn’t part of linguistically controlled content do not, in themselves, help with his overall project of equating intuitive truth-conditions with linguistically controlled content, Stanley attempts to show that, despite the initial intuitions, transferred meaning might not be part of intuitive truth-conditions after all. He points out that, in response to a waiter’s utterance of ‘The ham sandwich is getting annoyed’, one could felicitously respond with ‘That’s absurd; sandwiches do not get annoyed!’, and alleges that it is, therefore, the ‘literal’ rather than the ‘transferred’ meaning that figures in the truth-conditions. The problem with this argument, however, is that it severely overgenerates.⁵⁹ Suppose that Jane is clearly tall for an ordinary woman, but medium-sized for a basketball player. X could utter ‘Jane isn’t tall’ (intending to say that she isn’t tall for a basketball player), to which Y could felicitously respond, ‘That’s absurd; of course she’s tall!’ (meaning *tall for an ordinary woman*). Similarly, if X said ‘I bet John went to the bank’ (meaning he went to the river bank), Y can felicitously respond; ‘That’s absurd; he told us just this morning the bank is closed today!’ (meaning *money bank*). As we have noted in our discussion of conversational implicature (§3.4), denial can often be used to target one of several possible semantic interpretations.

We conclude that, if transferred meaning is (at least) part of the intuitive truth-conditions, the phenomenon shows that one of two substantive conclusions ought to be accepted: either (as Recanti and others have argued) intuitive truth-conditions can diverge from linguistically controlled content and are subject to pragmatic intrusion; or there are covert indexicals which are not subject to binding.⁶⁰

§6. Beyond Sandwiches

We’ll conclude by mentioning a number of additional linguistic phenomena that meaning transfer may impact. Our discussions here are suggestive rather than conclusive; adequate discussion of each case requires a paper of its own. Nonetheless, they will illustrate the broader implications of meaning transfer.

§6.1 Systematic Polysemy and Generalized Meaning Transfer

Nouns like ‘lamb’ can be used as both mass and count. Consider (85), which is ambiguous.⁶¹

(85) Mary had a little lamb.

On one reading, (85) is true if Mary was served a small lambchop. On the other, (85) is true if she had a small pet lamb. Common nouns like ‘lamb’ are hardly the only terms that have multiple semantically distinct occurrences. While no single case is uncontroversial, it is plausible that definite descriptions, names, gerunds, and infinitives all do.⁶² Names, for instance, plausibly occur as both referential and predicative.⁶³

Focusing on ‘lamb’, a natural question arises: what is the relationship between these semantically distinct occurrences? One answer is that ‘lamb’ is ambiguous, and the lexicon contains both the count and mass sense. A second is that there is just one lexical entry for ‘lamb’ and whether the count or mass use is selected for depends on particular features of context, with linguistic context playing a central role (cf. Pelletier 2012). A third is that the lexicon contains just one of these meanings and there is a mechanism for deriving the other.⁶⁴

Given our discussion, a natural version of the third option arises: that the lexicon merely contains the count sense of ‘lamb’ and there is an oft-salient meaning of *MT* that allows us to derive the mass sense.⁶⁵ There is a clear difference between the *MT*-account of count-to-mass conversion and the paradigmatic cases of meaning transfer we’ve been discussing thus far. Paradigmatic cases rely on very specific features of contexts to make a value for *MT* salient. Count-to-mass conversion, by contrast, seems readily available across a wide variety of contexts without special priming. However, this is hardly a worry. Just as there are generalized and particularized conversational implicatures, we’d predict that there are generalized and particularized meaning transfers, where the former involves values for *MT* that aren’t dependent on specific features of individual contexts.

Of course, count-to-mass conversion is just one instance of systematic polysemy, the phenomenon wherein a type of polysemy seems to be available across a wide range of lexical items and languages. If an *MT*-based account of count-to-mass conversion is promising, it will be natural to extend this to a variety of cases of systematic polysemy. Whether an *MT*-based account is correct, however, will depend on the details of the particular cases.⁶⁶

To stress again, our aim here is not to defend an *MT*-based account of these sorts of semantic shifts, but merely to point out that this account is worth serious consideration when it comes to any sort of systematic polysemy.

§6.2 Remote Complex Demonstratives

Complex demonstratives are sometimes felicitously uttered even when the item intuitively demonstrated is not physically present in the location of the utterance (call these ‘remote complex demonstratives’). Consider, for example, the following (uttered while pointing to a painting of Philippa Foot):

(86) This woman is a famous philosopher!

Or, suppose we are standing outside a field, which we all know is typically occupied by a certain cow, Angus. Today, however, Angus was taken to the vet and the field is empty. Still one can point at the field and utter:

(87) That cow produces a lot of milk!

Some have found such utterances puzzling; one might have expected that the entity being demonstrated needs to be present at the location of the utterance. A natural hypothesis is that this is a case of meaning transfer. Without entering disputes about the exact semantics of complex demonstratives, perhaps, in (86), we are, strictly speaking, demonstrating a painting, which then undergoes meaning transfer to a person, while, in (87), we are, strictly speaking, demonstrating a field, which then undergoes meaning transfer to a cow.⁶⁷

However, it is worth noting that there is a natural alternative account of remote complex demonstratives, one that does not involve Meaning Transfer. (We ourselves are sympathetic to this latter view, though we do not argue for it here.) The view is simply that an object need not be present in order to be (directly) demonstrated. In order to be demonstrated, an object needs to be sufficiently salient (in a particular way). One common source of salience is being directly perceived by the conversational participants, but there can be other sources (e.g. being depicted by a salient painting, or being known to occupy a salient field).⁶⁸

§6.3 Property Inheritance, Dead Meaning Transfer, and Copredication

In §3.3, we argued that meaning transfer cannot be reduced to property inheritance. Our arguments there departed from the observation that property inheritance and meaning transfer have distinct linguistic properties.

Different as they are, there may be close relationships between property inheritance and meaning transfer. In particular, if an instance of meaning transfer is common-enough, it may be absorbed into the lexicon, becoming a case of property inheritance. In fact, this has plausibly occurred with a famous example from discussions of meaning transfer (which, as we’ll make clear, we think is not an example of meaning transfer):

(88) Noa is parked out back.

Though (88) is often put forth as a paradigmatic instance of meaning transfer, our discussion thus far gives us reason to think that it is not meaning transfer at all. To see this, begin by noting that, if (88) exhibits meaning transfer, it must be the predicate ‘parked out back’, rather than ‘Noa’, which has its meaning transferred. This is demonstrated in (89) and (90).

(89) Noa is parked out back. He has to walk a long way.

(90) Noa is parked out back and he’s regretting it.

On the view that (88) exemplifies meaning transfer, ‘parked out back’ is transferred to mean *has a vehicle that is parked out back*. However, as Barrios (2013: 283) notes, there is some evidence against this view. Many (though not all) report that (91) is acceptable.

(91) Noa is parked out back, as is Mark’s car.

However, if (88) is an instance of meaning transfer, then (91) should be unequivocally odd, as it attempts to ascribe the property of *having a vehicle parked out back* to Noa’s car.⁶⁹

A straightforward explanation of the divergent intuitions regarding (91) is that the meaning transfer is in the process of dying, so some speakers have access to a lexicalized meaning for ‘parked out back’ that applies to people while others don’t. We can call this ‘dead meaning transfer’, as it is on par with dead metaphor.

In fact, we can distinguish two ways for an instance of meaning transfer to die. Focus on ‘parked out back’. The meaning transfer may die because there is a new item in the lexicon that is true of people who possess vehicles that are parked out back. Or, the meaning transfer may die because the meaning of ‘parked out back’ is expanded, so that it can be instantiated in different ways: the familiar way that applies to cars; and (once the death is complete) by people when they inherit the property from their vehicles. The fact that (91) is acceptable to some suggests that the second sort of death is operative in this case, and ‘parked out back’ can give rise to property inheritance rather than meaning transfer.

Of course, dying meaning transfer is hardly limited to ‘parked out back’. In fact, some of the cases we discussed in §6.1 may now be instances of dead meaning transfer. ‘Chicken’ for instance, is plausibly used often enough in both its count and mass sense that both are incorporated into the lexicon. This doesn’t mean that we’d want to incorporate all mass-to-count conversion into the lexicon—our arguments from §6.1 still stand—but some such conversions may be ubiquitous enough to give rise to their own lexical entries (or to give rise to novel expanded lexical entries).

The bridge from meaning transfer to property inheritance is also reflected in their relationship to copredication. Copredication occurs when an intuitively true sentence seems to ascribe two incompatible properties to a single argument. The most common reaction to such sentences is to take the properties to be genuinely incompatible and either abandon truth-conditional semantics (Chomsky 2001), or modify one's semantics and metaphysics enormously to accommodate the examples (Asher 2011; Gotham 2016). In Liebesman and Magidor (2017), we pursue a different approach. We hold that the purportedly incompatible properties are, in fact, compatible. On our view, this compatibility is often explained by invoking property inheritance. While two entities may instantiate the same property, they often instantiate it in different ways: one inherits it while the other doesn't. In that paper, we don't focus on cases in which one of the predicates undergo meaning transfer, but, as Nunberg (1995) stresses, this is one plausible way to explain-away the apparent incompatibility.

Imagine, for instance, that a waiter is complaining that ham sandwich orders are always jerks. Wishing to cite a counterinstance, his colleague may utter the following:

(92) That's not always true! Bill is a ham sandwich and a perfect gentleman.

In this case, the apparent copredication puzzle is immediately dissolved when we note that the meaning of 'ham sandwich' is transferred.

Whether meaning transfer can be used to explain away an instance of copredication will depend on the details of that particular instance. What is important to note here is that both meaning transfer and property inheritance are tools that can be used to explain cases of copredication without any revisionary semantics or metaphysics.⁷⁰

Notes

1. Meaning transfer has been discussed extensively by Nunberg. See his (1978), (1979), (1993), (1995), and (2002). As will be clear in our discussion, his insightful work has informed our own in myriad ways.
2. When we say that the locus of the meaning transfer is the subject-term, we mean that it occurs *somewhere* with the expression 'The ham sandwich'. This leaves open both the possibility that the meaning transfer scopes over the entire description and the possibility that it takes a narrower scope. We leave this question open, except in specific discussions of scope (see e.g. §2.4 below).
3. Philosophers of language and semanticists have been interested due to the connections between meaning transfer and the semantics/pragmatics distinction, as well as indexicality, e.g. Stanley (2005), Recanati (1993), and (2010), Fauconnier (1994), as well as all of Nunberg's work cited in fn. 1. Psychologists and cognitive scientists have studied processing in such examples, e.g. Schumacher (2014) and McElree, et. al. (2006). Cognitively-oriented linguists have been

interested in meaning transfer as an example of polysemy and metonymy. See Falkum and Vicente (2015) for a survey article on the former, and the introduction to Barcelona (2012) for an overview of the latter. Philosophers, Linguists, and Computer Scientists have all been interested in coercion. See, e.g. Clark and Clark (1979), Copestake and Briscoe (1995), and Pustejovsky (1995).

4. Note that while many authors seem to use the terms ‘deferred meaning’ or ‘deferred reference’ to pick out the same phenomenon we call ‘meaning transfer’, others use them to distinguish a different phenomenon. Thus for example, Recanti (1993), §16.4 argues that cases where the index of a context sensitive term diverges from its referent (e.g. when pointing at a painting of Murdoch and saying ‘That’s a famous philosopher!’) are a distinct phenomenon for meaning transfer, and (as far as we understand) he reserves the term ‘deferred reference’ as a label for the former phenomenon. We are very sympathetic to Recanti’s claim that these are different phenomena (cf. our discussion in §6.2) but will not follow this use of the terminology.
5. Of course, if both ‘P’ and ‘Q’ are transferred (or untransferred), both in the original utterance and when they are disquoted, then the ontological commitments will match. More cautiously put, the possibility of meaning transfer shows that we can’t safely conclude from a true utterance of ‘P is Q’ that it is true that P is Q, where the second occurrence of ‘P’ and ‘Q’ are literal.
6. See, for instance, Nunberg (1995: 110; 115–16).
7. To appreciate this point, compare a similar question for metaphor. What are the conditions under which a metaphorical reading arises? That’s a difficult question that will have a highly complex answer. Fortunately, we can inquire after the linguistic properties of metaphor without answering it.
8. Some have reacted to this example by claiming that ‘Hot abs’ is an *ad hoc* name introduced in the context in order to designate the person with the hot abs. This proposal has not been developed in the literature, so we can’t engage with it in a thoroughgoing way, but, more importantly, it may not conflict with the view that (4) exemplifies meaning transfer. Meaning transfer is plausibly one device for creating *ad hoc* names.
9. One potential exception might be the gender marking of the definite article in French and Italian. For example, suppose that the male customer who ordered a niçoise salad has complained about the sauce. Several informants have reported that they find the following instance of meaning transfer acceptable: ‘La salade niçoise s’est plainte de l’assaisonnement’. While the noun (‘salade niçoise’) is marked as feminine, the verb (‘s’est plainte’) receives masculine gender marking (just as we predict, given that the transferred meaning picks out a male customer). Interestingly, however, these informants nevertheless expected the gender marking of the definite article to be the female ‘la’ rather than the male ‘le’. On the other hand, we found mixed reactions about parallel examples in Italian. Some speakers reported that they would be happy to describe a male customer who ate pasta and left without paying using ‘Il pasta è andato via senza pagare’ with both the verb and the definite article marked as masculine like the customer, even though the noun ‘pasta’ is feminine. Others preferred ‘La pasta è andato via senza pagare’ with the article marked as feminine and matching ‘pasta’ but the verb marked as masculine. Clearly more cross-linguistic research is needed.

10. It is worth noting that one feature that might not track transferred meaning is the proximity of demonstratives used to pick out the trigger of transfer. Thus for example, we can say ‘This [pointing to a nearby sandwich] ham sandwich has just stepped outside, but that [pointing to a further away sandwich] ham sandwich is standing right behind you’. There are several potential explanations of what is going on here: in addition to the hypothesis that proximity does not transferred meaning, two alternative hypotheses are (i) that proximity is a delicate matter, so that whether an object counts as ‘proximal’ depends on more than its physical distance from the speaker (and hence the customers outside can count as the more proximal in this context) (ii) that meaning transfer scopes over the entire demonstrative phrase here (in which case we wouldn’t expect the features of the demonstrative to track the transferred meaning). See also our discussion in §6.2 below.
11. If this is correct, we’d predict that some uses in some languages will provide examples of meaning transfer with singular subjects and plural predicates. A potential example is that in British English it is often acceptable to use singular definite descriptions in subject position with plural predicates, e.g. ‘The hospital are useless at keeping the rooms sterile’ or ‘The college are considering me for a fellowship’ to express a proposition about the hospital employees/college employees.
12. Fauconnier (1994): 169 (n. 3) attributes this kind of example to Encrevé.
13. We can give similar examples for ellipsis. For instance, “Bill is a rib-eye steak, as is that delicious-looking thing I’ve been eyeing ever since I walked into the kitchen.” In what follows, we’ll focus solely on anaphora, though all of our reasoning generalizes to ellipsis.
14. Thanks to Rob Stainton for this way of putting the point. We take it that the idea that ‘it’ here is a deictic rather than anaphoric pronouns is what Ward has in mind when he considers the possibility that the ‘it’ here is a “pronoun of laziness” (Ward (2004): 270), though we ourselves would not use that term in this context. Ward himself dismisses this proposal on grounds of parsimony (*ibid.* n. 10).
15. Further support for the claim that the piece of food isn’t available for anaphoric reference comes in an example from Fauconnier (1994): 6. Fauconnier notes that “*The mushroom omelette was eating itself/himself with chopsticks’ is infelicitous (it cannot receive the reading that the mushroom omelette orderer was eating the omelette with chopsticks).
16. Of course, gender and number marking themselves might be presuppositional in which case it is not surprising that they track transferred meaning.
17. A similar point applies to definite descriptions with embedded quantifier phrases. Imagine that three customers shared two ham sandwiches between them. ‘The three ham sandwiches left without paying’ (on the transferred reading) is acceptable, because there is a unique trio of customers who ate ham sandwiches (even though there wasn’t a unique trio of sandwiches).
18. Similar remarks apply to Krifka’s view, which would require some sort of scale that renders irrelevant the difference between a person and the sandwich they order. (By contrast, it is easy to generate a scale that renders irrelevant the difference between 3:00 and 3:01: the scale of 15-minute temporal increments does just that.)

19. That is, setting aside non-cognitivist views of metaphor on which there is no proposition metaphorically communicated (see Davidson (1978)).
20. See Tirell (1989) on the central role that extendibility plays in the theory of metaphor.
21. Another thing to note is that there are readings of (38) that are acceptable, but these are precisely readings in which the first sentence is read as a metaphor rather than an instance of meaning transfer.
22. Indeed, the *Stanford Encyclopaedia of Philosophy* entry on conversational implicature (Davis (2014)) classifies metonymy, together with other figures of speech, as an instance of conversational implicature.
23. According to some theorists, the sentence, when taken literally, is meaningless rather than false, because it is a category mistake. Note, though, that since implicatures are generated by the proposition literally expressed, the view that category mistakes are meaningless conflicts with the hypothesis that meaning transfer is derived via implicature (see Magidor (2017)).
24. Note that we are not using ‘semantic content’ in a particularly precise or theoretically loaded sense here. We merely use the term to give a contrast to implicature. We’ll discuss the relevance of meaning transfer to the semantics/pragmatics distinction in more detail in §5.
25. Recanti (1993): 264–5 provides yet another argument for thinking that meaning transfer isn’t an instance of implicature. According to Recanti at no point in deriving the interpretation of sentences such as (1) does the speaker need entertain the “absurd” proposition about a piece of food walking out the door. We ourselves are more sceptical that the presence of such Gricean derivations are always accessible to speakers in this way, and hence won’t rely on this argument.
26. This is an instance of what Simons (2009) calls the composition problem for intrusive implicatures.
27. This sort of reasoning is stressed by, e.g. Sperber and Wilson (1986), Carston (2002), and Recanati (2003b).
28. Wearing (2006) gives somewhat similar arguments against treating metaphorical interpretation as derived via conversational implicature. In particular. She focuses on the fact that metaphorical interpretations can be freely embedded.
29. Chierchia (2004), Recanati (2004), Geurts (2011), and Simons (2009) represent some of the differing reactions.
30. See, e.g. the exchange between Simons (2017) and Recanati (2017a).
31. Recanati (2010) introduces a *mod* operator to account for modulation which he takes to subsume meaning transfer, among other things. On his view, *MT* could be eliminated meaning transfer subsumed under of *mod*.
32. Our discussion here is influenced by Stainton’s (2006) criticism of the elliptical view of sub-sentential assertion, as well as discussions with Rob Stainton about generalizing his arguments. Our comments in this section also draw heavily from discussion with Michael Caie who considered this view in commentary on our paper at the 2019 Eastern APA.
33. This example comes from Michael Caie’s APA commentary.
34. Thanks to Rob Stainton for this example.
35. See Stainton (2006: Chapter 6) for detailed discussion and motivation.

36. Stanley (2000) notes that, in principle, we can make linguistic antecedents salient without explicitly uttering sentences containing them. We're somewhat skeptical of Stanley's claims in the first place, but even granting them, it is hard to see how they could be generalized to all of our cases of meaning transfer. Also, see Stainton (2006: Chapter 7) for responses to Stanley.
37. A difficult and interesting question, which we don't tackle, is just what constrains and determines the value of *MT* in a particular context. Nunberg (1979; 2002) investigates this issue.
38. The case of 'The large impatient ham sandwich left without paying' in the context of example (68) suggests particularly complex scope possibilities. Given that 'impatient' qualifies customers who ordered a ham sandwich, while 'large' seems to qualify sandwiches, we would presumably need at least two occurrences of *MT*: one transferring 'ham sandwich' to 'customer who ordered a ham sandwich'; and one transferring 'large' to a highly non-standard function which, takes a property *F*, not to the property of being a large *F*, but to the property of being an *F* with a relatively large food order. (Note that Magidor (2017), who proposes the example, draws a different conclusion from this case, though she also does not offer a positive theory of how the meaning transfer comes about.)
39. Examples involving numerical quantifiers also display the multiple scopal possibilities for *MT*. Consider 'The five ham sandwiches left without paying', which has two distinct readings. Imagine, first, that there are five different people that ordered ham sandwiches. In that case, it expresses that the five ham sandwich orderers left without paying. By contrast, imagine that there is a single person that ordered five ham sandwiches. In that case, it expresses that the orderer of five ham sandwiches left without paying. Since there is only one ham sandwich orderer in the second scenario, *MT* must scope over 'five ham sandwiches'—if it scoped under 'five' there wouldn't be enough ham sandwich orderers for the phrase to designate. If we assume that in the first scenario multiple people share sandwiches (or the restaurant is out of ham) then *MT* must scope under 'five', as there aren't five ham sandwiches. Such sentences are discussed by Sag (1981: 285).
40. See, e.g., Graff Fara (2014).
41. Another interesting scope possibility for *MT* is defended in Ward (2004). Ward argues that in certain uses of sentences such as 'John is the pad thai' (which he calls 'equatives'), both 'John' and 'the pad thai' receive their ordinary referents, and it's rather the copula that undergoes meaning transfer. Some support for this claim (*ibid.* 281) is given by the fact that apparently both John and the pad thai are available for further predication (consider: 'John, who drives a Rolls Royce, is the pad thai' and 'John is the pad thai, which is delicious').
42. Recanati (2010: 11) hypothesizes that his mod operator preserves semantic types. Recanati suggests that meaning transfer is accounted for using this operator, in which case meaning transfer would have to preserve semantic types as well.
43. Also, recall our example in §4.1, where 'Magdalen College always arrives late' is used to mean that *the person working in Magdalen College always arrives late*. Focus on a reading of this sentence where it might be different people working in Magdalen College on different occasions being quantified over. The fact that we have a variable reading entails that the name must be transferred to a description.

If descriptions are a different type, then we have an example of cross-type transfer. Even if descriptions are the same type as names, then we have a case of transfer from a meaning that can't give rise to variability to one that can.

44. Note that this doesn't have to be the location in which the utterance takes place. For instance, in some contexts where Dave, who is on vacation in Paris, utters, 'I love my local coffee shop', 'local' will designate a location in Calgary (his usual location) not Paris, the location of his utterance.
45. At a stretch, there is perhaps a reading of (79) where 'are ham sandwiches' is transferred to a very general property such as: 'are somehow related to ham sandwiches'. But this isn't the bound reading we are looking for.
46. Of course, this same phenomenon can occur when the input sentence also contains a description. For example, 'The bicycle always arrives last' can be used to mean that whoever commutes to work by bicycle always arrives last, even if it's a different person in each situation.
47. Stern himself doesn't take this example to show the shiftability of *Mthat*, because he thinks the antecedent of the conditional is enough to alter the actual context in play. As Camp (2005) argues, though, other non-shiftable indexicals don't appear to allow such flexibility when they appear in the consequent of a conditional, and at any rate, this response does not help with the binding data discussed below.
48. There's some precedent for content invisibility in Nunberg (1993) who argues that indexicals all exemplify deferred reference and are content-invisible.
49. Note that the content invisibility of *MT* is consistent with several views about the semantics/pragmatics interface, but not all. Consider four views. On the first, *MT* is present at the level of logical form and is present in the proposition expressed, so that the proposition expressed by (1) itself contains *MT*. On the second, *MT* is present at the level of logical form, but not in the proposition expressed. On the third, *MT* is not present at the level of logical form, but makes its way into the proposition expressed (though a process of enrichment). On the fourth, *MT* is not present at either logical form or in the proposition expressed, though it could be used to model the relation between the two (which would be a form of pragmatic modulation). Content invisibly rules out the first and third view, but are compatible with the second and fourth.
50. Cf. Stanley (2000) who maintains that there aren't any non-shiftable indexicals beyond the pure ones.
51. Another reason one might be tempted to accept content invisibility has to do with data pertaining to 'one' anaphora. Thus, for example, Hawthorne and Manley (2012: 104) argue that the property of being a student must be present in the proposition expressed by 'A certain student came to my office hours' because it can be anaphorically referred to as follows: 'A certain student came to my office hours. Another one came later'. This test seems encouraging for the content invisibility of *MT* (and its argument) because, for example, the following is infelicitous: 'The ham sandwich left without paying. *The one on my plate has mustard on it'. However, we do not want to put too much weight on this, since the data on 'one' anaphora is quite complex and hard to systematize. For example consider the contrast between following 'The ham sandwich left without paying' with (i) *The tuna sandwich one left a large tip (highly infelicitous) and

- (ii) The tuna one left a large a tip (arguably acceptable). We are not sure how to account for all of this data.
52. Nunberg (1979) pushes a very different line: that there is no semantics/pragmatics distinction. Meaning transfer plays a key role in his argument, which departs from several considerations that we don't discuss here.
 53. While the second and third conceptions are defined in different terms they might turn out to be extensionally equivalent. After all, one might maintain that we have no independent notion of what an utterance's truth-conditions are; the only handle we have on truth-conditions is precisely those which correspond to the linguistically controlled content. However, a wide variety of authors maintain that we do have an independent grasp of an utterance's truth-conditions (often labelled 'intuitive truth-conditions'; see Stanley (2005) and Recanti (2003)). Given such an independent notion of intuitive truth-conditions, one can ask whether these truth-conditions do or do not correspond to those determined by linguistically controlled content. One prominent hypothesis (defended in, e.g., Stanley (2000; 2005)) is that the two coincide.
 54. 'Semanticism' is Stanley's label for his view that intuitive truth-conditions turn out to coincide with linguistically controlled content, essentially claim (ii) in our triad below.
 55. To cite one formulation of this claim: "My purpose in this paper is to defend the thesis that all truth-conditional effects of extra-linguistic context can be traced to logical form" (Stanley 2000: 1).
 56. Alternatively, defenders of (ii) can maintain that the verb 'to rain' is itself an indexical – but the same complaint about *ad hoc* structures will apply to this possibility.
 57. See, e.g., Stanley and Szabó (2000) and Stanley (2000).
 58. Stanley also argues that the fact that transferred meaning is available for anaphoric reference or ellipsis does not prove that it's a semantic phenomenon, because the same is true of phenomena that he takes to be clearly non-semantic, such as metaphor or irony. Discussion of this argument falls beyond the scope of this paper - for one thing, whether one finds it successful will depend on one's theory of metaphor/irony. At any rate, our current concern isn't with positive arguments for the claim that the phenomenon is semantic.
 59. Stanley does recognise that, taken on its own, this consideration would over-generate for at least some cases – such as quantifier domain restrictions, but again stresses the importance of the theoretical considerations we have already dismissed above.
 60. We note that there are other phenomena (assuming these are not subsumed as a case of meaning transfer) that might demonstrate this point. Consider various forms of flexible type-shifting (e.g. mass/count conversions, or noun/verb conversions). It is highly likely in these cases that the 'converted' meanings are part of the intuitive truth-conditions ('There is too much apple in my salad' or 'I shelved the book' are, in the relevant contexts, intuitively true). But it is hard to see any plausible evidence of a bindable indexical in these cases. (Interestingly, Stanley (2000: n.1) claims he can set aside type-shifting cases, as they involve the effects of 'linguistic context' on interpretation, while his thesis only applies to the effects of extra-linguistic context. Presumably, he thinks that, in cases of

type-shifting, the sentence would not even be well-formed without the relevant shift. But this isn't true of all cases of type-shifting. 'Some apple is in my salad' can be well-formed whether 'apple' is used as a count or as a mass noun).

61. The example comes from Quine (1960: 91).
62. See Liebesman (2015: §3) for discussion.
63. This is far from uncontroversial, though see Schoubye (2017) for a recent defence.
64. As stressed by Recanati (2017b), these options are not exhaustive. The senses of a polysemous expression may be both derived and conventional.
65. Nunberg (1995) also considers treating systematic polysemy in this manner.
66. For instance, many think that words like 'book' exhibit a token to type systematic polysemy. Our own view, defended in Liebesman and Magidor (2017), is that this is incorrect for 'book' and other such terms, because their standing meanings cover both tokens and types. This is a version of the second view we considered above for 'lamb'.
67. Adapting a discussion in Elbourne (2008), which is based on Nunberg (1993: 24), one might argue against this suggestion as follows. Consider a case where we are standing next to Field A with Field B further behind. Field A is typically occupied by Cow A, and Field B by Cow B, but both cows are currently absent (and, if you want, Cow A is located further away than Cow B). Elbourne notes that, in this context, we can utter 'This cow [gesture at Field A] is healthier than that cow [gesture at Field B]', to mean that Cow A is healthier than Cow B. The interesting thing is that the proximity features of the demonstratives seem to track the proximity of the *fields* rather than the *cows* (with the nearer field being referred to using 'this' and the further field using 'that'). Given that (as we have argued) feature agreement typically tracks the transferred meaning, this might suggest that remote complex demonstratives are not a case of Meaning Transfer.

We do not find this argument persuasive. First, note that various other features of the demonstratives *do* seem to track the cows rather than the fields. For example, if Field A typically contains two cows, it's acceptable to say, 'These cows [gesture at Field A] are healthier than that cow [gesture at Field B]' (similar data occurs regarding gender marking in Hebrew). Second, note that as we discuss in note 10, even in uncontroversial cases of meaning transfer, proximity, at least arguably, tracks the original rather than the transferred meaning. (For example, if we say 'This ham sandwich [gesture at Sandwich A] went out to smoke with that ham sandwich [gesture at sandwich B]' we are typically tracking the respective distances of the sandwiches, not of the costumers).
68. Cf. the discussion in Recanti (1993), §16.4 who also argues that such cases are not instances of meaning transfer.
69. Similar considerations tell in favour of a property inheritance (rather than meaning transfer) view of another well-known example in the literature: 'Billy's shoes were neatly tied but dirty' (Nunberg 1995: 123).
70. Thanks to Jonathan Payton, Francois Recanati, and Rob Stanton for detailed comments on drafts that led to numerous improvements. Thanks to Herman Cappelen, Cian Dorr, John Hawthorne, Adam Przepiórkowski, and Betsy Ritter for helpful discussion. Thanks to Maria Aloni, Sophie Grace, Marie Guillot, Daphna Oren-Magidor for suggesting some example sentences that feature in the paper. Thanks to Beth Barker and Michael Caie who gave public responses

to the paper in the Kline Workshop and the Eastern APA respectively. Thanks to audiences in the Eastern APA 2019; CEU Budapest; ILLC Amsterdam; University of Missouri, Columbia; University of Oslo; University of Oxford as well as support from the Leverhulme Trust, and the Social Sciences and Humanities Research Council of Canada.

References

- Asher, N. (2011): *Lexical Meaning in Context*. Cambridge: Cambridge University Press.
- Barcelona, A. (2012): *Metaphor and Metonymy at the Crossroads*: Berlin: De Gruyter.
- Barrios, E. (2013): 'Meaning shift and the purity of 'I'', *Philosophical Studies* 164: 263–288.
- Camp, E. (2005): 'Joseph Stern: *Metaphors in Context*', *Nous* 39:715–731.
- Carston, R. (2002): *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Malden, MA: Blackwell.
- Chomsky, N. (2000): *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Chierchia, G. (2004): 'Scalar implicatures, polarity phenomena and the Syntax/Pragmatics Interface.' In Adriana Belletti (ed.), *Structures and Beyond: The cartography of Syntactic Structures*, Volume 3. Oxford: Oxford University Press, 39–103.
- Clark, E. and Clark, H. (1979): 'When Nouns Surface as Verbs', *Language* 55(4): 767–811.
- Collins, J. (2017): 'The copredication Argument', *Inquiry* 60(7): 675–702.
- Copestake, A. and Briscoe, T. (1995): 'Semi-Productive Polysemy and Sense Extension', *Journal of Semantics* 12: 15–67.
- Davis, W., 'Implicature': *The Stanford Encyclopedia of Philosophy* (Fall 2014 Edition), Edward N. Zalta (ed.), URL=<<https://plato.stanford.edu/archives/fall2014/entries/implicature/>>.
- Davidson, D. (1978): 'What metaphors mean': *Critical Inquiry* 5: 31–47.
- Elbourne, P. (2008): 'Demonstratives as individual concepts', *Linguistics and Philosophy* 31: 409–466.
- Fauconnier, G. (1994): *Mental Spaces: Aspects of Meaning Construction in Natural Language*, Cambridge University Press: Cambridge.
- Geurts, B. (2010): *Quantity Implicatures*, Cambridge: Cambridge University Press.
- Gotham, M. (2016): 'Composing Criteria of Individuation in Copredication'. *Journal of Semantics* 1–39.
- Graff Fara, D. (2014): 'Names are Predicates', *The Philosophical Review* 124 (1):59–117.
- Hawthorne, J. and Manley, D. (2012): *The Reference Book*, Oxford: Oxford University Press.
- Krifka, M. (2007): 'Approximate Interpretations of Number Words', In *Cognitive Foundations of Interpretation*, G. Bouma, I. Kramer & J. Zwarts, editors, 111–26. Koninklijke Nederlandse Akademie van Wetenschappen.
- Laserson, P. (1999): 'Pragmatic Halos', *Language* 75(3): 522–51.
- Liebesman, D. (2011): 'Simple Generics', *Nous* 405: 409–442.
- Liebesman, D. (2015): 'Predication as Ascription', *Mind* 124: 517–69.
- Liebesman, D. & Magidor, O. (2017), 'Copredication and property inheritance', *Philosophical Issues* 27: 131–166.
- Liebesman, D. & Magidor, O. (forthcoming): 'Copredication, counting, and criteria of individuation: a response to Gotham', forthcoming in *Journal of Semantics*.
- Lossius Falkum, I. and Vicente, A. (2015): 'Polysemy: Current Perspectives and Approaches', *Lingua* 157: 1–15.
- Magidor, O. (2017): 'Category mistakes and figurative language', *Philosophical Studies* 154: 65–78.

- McElree, B., Frisson, S., and Pickering, M. (2006): 'Deferred Interpretations: Why Starting Dickens is Taxing but Reading Dickens Isn't', *Cognitive Science* 30: 181–92.
- Nunberg, G. (1978): *The Pragmatics of Reference*, Bloomington, IN: Indiana University Linguistics Club.
- Nunberg, G. (1979): 'The Non-Uniqueness of Semantic Solutions: Polysemy', *Linguistics and Philosophy* 3.
- Nunberg, G. (1993): 'Indexicality and Deixis', *Linguistics and Philosophy* 16:1–43.
- Nunberg, G. (1995): 'Transfers of meaning', *Journal of Semantics* 12: 109–132.
- Nunberg, G. (2002): 'The Pragmatics of Deferred Interpretation', in *The Blackwell Encyclopedia of Linguistics*, Laurence Horn and Gregory Ward, eds., Blackwell, 2002.
- Pelletier, J. (2012): 'Lexical Nouns are Both +Mass and +Count, but They are Neither +Mass nor +Count', in *Count and Mass Across Languages*. Massam, D. (ed.) Oxford: Oxford University Press, 2012.
- Pietroski, P. (2005): 'Meaning Before Truth', in *Contextualism in Philosophy*. Preyer and Peters, (eds.). Oxford: Oxford University Press, 2005.
- Pustejovsky, J. (1995): *The Generative Lexicon*, Cambridge, MA: M.I.T. Press.
- Quine, W.V.O. (1960) *Word and Object*, Cambridge, M.I.T. Press.
- Recanati, F. (1993): *Direct Reference: From Language to Thought*, Blackwell: Oxford.
- Recanati, F. (2003): 'What is said and the semantics/pragmatics distinction' in Claudia Bianchi, Carlo Penco (eds.), *The Semantics/Pragmatics distinction: Proceedings from WOC 2002*, CSLI Publications, 2003.
- Recanati, F. (2004): *Literal Meaning*, Cambridge: Cambridge University Press.
- Recanati, F. (2010): *Truth Conditional Pragmatics*: Oxford: Oxford University Press.
- Recanati, F. (2017a): 'Local Pragmatics: Reply to Mandy Simons', *Inquiry* 60(5): 493–508.
- Recanati, F. (2017b): 'Contextualism and Polysemy', *dialectica* 71: 379–397.
- Sadock, J. (1978): 'On testing for conversational implicature', In Cole, Peter (ed.), *Syntax and Semantics: Pragmatics*, Academic Press: 281–297.
- Sag, I. (1981): 'Formal Semantics and Extra-Linguistic Context', in Cole, Peter (ed.), *Radical Pragmatics*, Academic Press: 273–94.
- Schoubye, A. (2017): 'Type-Ambiguous Names.' *Mind* 126: 715–67.
- Schumacher, P. (2014): 'Content and Context in Incremental Processing: 'The Ham Sandwich' Revisited.' *Philosophical Studies* 168: 151–65.
- Simons, M. (2009): 'A Gricean View on Intrusive Implicatures', In Petrus, K., Sauerland, Uli (eds.) *New Essays on Grice*, Palgrave: 138–69.
- Simons, M. (2017): 'Local Pragmatics in a Gricean Framework', *Inquiry* 60(5): 466–492.
- Solt, S. (2014): 'An Alternative Theory of Imprecision', In *Proceedings of Semantics and Linguistic Theory (SALT) 24*, Todd Snider, Sarah D'Antonio & Mia Weigand, (eds.) 514–33. CLC Publications, Ithaca.
- Sperber, D. and Wilson, D. (1986): *Relevance*, Oxford: Blackwell.
- Stainton, R. (2006): *Words and Thoughts*, Oxford: Oxford University Press.
- Stanley, J. (2000): 'Context and logical form', *Linguistics and Philosophy* 23: 391–434.
- Stanley, J. (2005): 'Semantics in context', in Preyer and Peter (ed.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, Oxford: Oxford University Press: 221–54.
- Stanley, J. and Szabó, Z. (2000), 'On quantifier domain restriction', *Mind and Language* 15: 219–261.
- Stern, J. (2000): *Metaphor in Context*, Cambridge, MA: MIT Press.
- Tirell, L. (1989): 'Extending: The structure of metaphor', *Nous* 23: 17–34.
- Ward, G. (2004): 'Equatives and deferred reference', in *Reference: Interdisciplinary Perspectives*, edited by Jeanette Gundel and Nancy Hedberg, Oxford: Oxford University Press: 73–92.
- Wearing, C. (2006): 'Metaphor and what is said', *Mind and Language* 21: 310–32.