

Response to Abrusán, Shaw, and Elbourne¹

Ofra Magidor, University of Oxford

[Draft March 2016 of paper to appear in *Inquiry*]

§1 Response to Abrusán

Márta Abrusán's paper consists of a comparison between my own book and Nicholas Asher's *Lexical Meaning in Context* (Asher (2011)), which discusses a wide range of linguistic phenomena, including category mistakes. Abrusán's main focus isn't Asher's theory of category mistakes, but rather his treatment of two different phenomena: so-called cases of 'co-predication' and 'coercion'.² The suggestion, I take it, is that if (only) Asher's account of category mistakes figures as part of a unified explanation of a wide range of phenomena then it has a significant advantage over my own account. Abrusán thus raises the question whether my own theory can be extended to provide parallel explanations of these additional phenomena.

In §1.1 I briefly compare Asher's account of category mistakes to my own, and argue that Asher's account is vulnerable to many of the objections that I raise in my book against competing theories. In §1.2-1.3 I offer some remarks about coercion and co-predication. While I do not claim to have anything like a worked out positive account of these phenomena (this has not been a focus of my work), I do wish to cast doubts about Asher's (and Abrusán's) assumption that these issues are closely connected to that of category mistakes. Thus rather than directly respond to Abrusán's challenge as to whether my own theory of category mistakes can be extended to provide parallel accounts to those of Asher's, I reject the assumption that such an extension is desirable.

§1.1 category mistakes

Asher's account of category mistakes proceeds roughly as follows: start with a background hierarchy of types.³ Each lexical item is semantically assigned a (rather fine-grained) type (for example, 'apple' is assigned the type '*offspring of a plant*'⁴), but terms also inherit types that are further up the type-hierarchy (so, for example, anything of type *offspring of a plant* also counts as a member of the type *physical object*). In addition, predicative term (such as verbs, adjectives, or adverbs) trigger *type presuppositions*, i.e. ones that require the arguments to be members of a particular type.⁵ Thus for example, 'is green' might trigger the

¹ I would like to thank Cian Dorr, Paul Elbourne, and David Liebesman, and Timothy Williamson for helpful comments and discussions, as well as of course Márta Abrusán, Paul Elbourne, and James Shaw for their thoughtful engagement with my work.

² I say 'so called' because as will be clear from my remarks, these labels expose some theoretical commitments about these phenomena which I do not share.

³ Asher opts to construct this hierarchy using topos theory – an intuitionistic variant of category theory.

⁴ Asher (2011), p. 41.

⁵ I note in passing that Asher suggests that the only types required for the type presuppositions associated with category mistakes are very general types such as 'physical object', 'abstract object', or 'event'; the more

presupposition that its argument is a physical object, and ‘the apple’ can satisfy this presupposition in virtue of (indirectly) belonging to this type. Finally, when the relevant presuppositions are not satisfied, the resulting sentence suffers from some serious semantic defect.

As Abrusán rightly comments, the main similarity between my own account and Asher’s is their appeal to *presuppositions*, but the accounts are otherwise very different. In very broad terms, the two accounts differ on three main questions:

Q1: What sorts of presuppositions are responsible for category mistakes?

Asher takes the presuppositions in question to be *type presuppositions*. For example, the presupposition triggered by ‘is green’ is that the argument is a physical object. On the other hand, I argue that the presuppositions triggered are *adjectival properties*. For example the presupposition triggered by ‘is green’ is that the argument is coloured.

Q2: What determines whether a particular argument satisfies the relevant presupposition?

On Asher’s view, whether an argument satisfies a presupposition depends on *its type*, and this in turn is determined by the *lexical semantics*.⁶ On my view, by contrast, whether an argument satisfies a presupposition depends on what properties the argument *taken for granted in the conversation* to have.

Q.3: What are the effects of presupposition failures?

I rely on a pragmatic account of presupposition: when the relevant presuppositions fail the sentence is nevertheless syntactically well-formed, meaningful, and truth-valued – albeit pragmatically infelicitous. It is clear that Asher takes presuppositions failures to lead to some much more serious defect, but it is not obvious precisely which defect he intends that to be. He claims that category mistakes are “nonsensical” (Asher (2011), p.4) and that they involve predication that “cannot be interpreted and fails to result in a well formed logical form” (Asher (2011), p.7) – suggesting that category mistakes are meaningless (if not syntactically ill-formed?). At the very least he seems to accept that category mistakes do not express propositions⁷, and that they are truth-valueless.⁸

A large part of my book is devoted to arguing against the views that category mistakes are syntactically ill-formed; well-formed but meaningless; or meaningful but truth-valueless

specific types such as ‘offspring of a plant’ play a role only in accounting for phenomena other than category mistakes (Asher (2011), pp. 49-52). But this assumption is incorrect. For example, Hebrew has a special verb for putting on socks (it is a category mistake to apply this verb to shoes or gloves) and a special verb for picking grapes (it is a category mistake to apply this verb to other kinds of fruit). This means that Asher’s account would need type presuppositions with type requirements as specific as ‘sock’ or ‘grape’ (see Magidor (2013), p.35).

⁶ In Asher (2011), p. 44-48, Asher distinguishes between two kinds of semantics: ‘internal’ and ‘external’ semantics, and maintains that he takes both to be in play. It is far from clear to me how this two-tiered semantic picture is supposed to work, but at any rate as the type clashes seem to have direct implications to the ‘external’ truth-values I do not see how this distinction would help avoid the concerns I raise in this section.

⁷ For example he claims that it is hard to “accept the possibility, or even to make sense of, a competent speaker’s believing or even entertaining” category mistakes” (Asher (2011), p.5).

⁸ Although I am puzzled by one remark in which Asher classifies the sentences “Tigers are financial institutions” and “Tigers are ZF sets” (which he takes to be category mistakes) as “false, indeed necessarily false, given the widely accepted externalist semantics for natural kinds like tiger” (Asher (2011), p.49).

(whether because they fail to express a proposition altogether or because they express a partial proposition). I will not rehearse all these arguments here – it should be clear why I take Asher’s answer to third question (however we interpret it) to be unsatisfactory.⁹

However, it is worth remarking on why several lines of thought throughout the book also tell against Asher’s response to the first two questions. With respect to the first question (whether the presuppositions in question are type presuppositions), I argue in passing that type presuppositions are not adequate for accounting for the phenomenon of category mistakes.¹⁰ Consider for example the following (potential) category mistakes:

- (1) *This male is pregnant
- (2) *This idea is prime
- (3) *This number is green
- (4) *This event is curved

In order to account for their infelicity the defender of type presuppositions would need to postulate that the predicates in question trigger the presuppositions that their arguments are (1) female; (2) a number; (3-4) a concrete object; Yet such presuppositional hypotheses conflict with the observation that all of the following speeches are perfectly felicitous:

- (5) *Doctor*: you know, we now have technology which enables males to also become pregnant. This male, for example, is pregnant.
- (6) *Mathematician*: you know, not only numbers but also polynomials are prime or composite. This polynomial, for example, is prime.

⁹ In [REF] Abrusán remarks (without going into detail) that she thinks the criticisms I raised against type-theoretic versions of the meaninglessness view in Magidor (2013), Chapter 3, §2.2 do not apply to Asher’s account. (In a brief, my argument there involved considerations along the following lines: if ‘prime’ is taken to be of type $\langle number, t \rangle$, then - to take a Montagovian view of singular terms - ‘two’ would have to be of type $\langle \langle number, t \rangle, t \rangle$ in order to account for the acceptability of ‘Two is prime’. However, that would falsely predict that, e.g., ‘Two is interesting’ is a category mistakes).

I think Abrusán is right that at least the main strand of my criticism in that section does not apply to Asher’s view, but it is worth clarifying why this is so. It has to do with two features of Asher’s account. First, while on standard type-theoretic accounts one can read-off the type of a term from its standard semantic value (which specify which arguments it accepts), on Asher’s system types need to be explicitly represented in the semantics as an extra feature. Thus for example in standard systems one can infer that a semantic-value such as *λe.e is green* is of type $\langle e, t \rangle$. However, once the type-theoretic system is used to account for category mistakes one must jettison this feature. For example, ‘prime’ and ‘factorable’ both accept the same types of arguments (numbers) but they must belong to different types if we are to account for the fact that ‘This number is very factorable’ is felicitous but ‘This number is very prime’ is not (because the type-presupposition of ‘very’ must admit one but not the other).

The second feature is simply that Asher’s account isn’t *purely* type-theoretic: that is, it utilises considerable more machinery than a hierarchy of types and constraints on type-matching. Thus for example in order to account for the fact that ‘The book is interesting’ and ‘The number two is interesting’ are both felicitous, Asher requires that the subject-terms in such sentences belong to dot-types and thus the relevant arguments do not straightforwardly satisfy the type presuppositions of the predicate ‘interesting’. Rather the subject undergoes further modulation before it is fed into the predicates (see the discussion in §1.2 of this paper).

Needless to say, both these features mean that Asher’s account is far more complicated than standard type-theoretic accounts. Moreover, I do not see how Asher’s account addresses my argument at the end of Ch. 3 §2.2 which points out that in so far as the theory is a version of the meaninglessness (rather than the MBT) view, it cannot account for category mistakes such as ‘That is green’ (where ‘that’ demonstrates the number two).

¹⁰ See Magidor (2013), pp. 142-144.

(7) *Physicist*: you know, even though photons aren't concrete objects we discovered they do have colours. This photon, for example, is green.

(8) *Physicist*: you know, not only concrete objects but also space-time itself has a curvature value. In fact, it turns out that space-time is curved.

Note in particular that on the type-theoretic story, the felicity of (5)-(8) cannot be explained via some mechanism of presupposition accommodation. After all, in these contexts it is clearly still assumed the arguments in question are (1) not female; (2) not a number; and (3-4) not concrete objects.

The adjectival presuppositions hypotheses of my own theory, on the other hand, have no problem reconciling the infelicity of (1)-(4) with the felicity of (5)-(8). For example, in the context of (6) (but not of (2)) we take for granted that the subject is prime or composite, and in the context of (7) (but not of (3)) we take for granted that the subject is coloured.

With respect to the second question (whether which presuppositions an argument satisfies is determined by the lexicon), there are several problems in supposing that the relevant types are encoded into the lexicon. For a start, a series of arguments which I raise against the syntactic approach tell against this kind of lexical semantic approach as well.¹¹ Thus for example, in so far as we think (9) is a category mistake, so is (10):

(9) This male is pregnant

(10) This person which I proved that John was mistaken in believing was female is pregnant.

But it is hard to imagine any plausible way of encoding types into the lexicon that would deliver the result that the phrase 'person I proved that John was mistaken in believing was female' would be assigned the type 'male'. It is also not clear how lexically encoded types are going to explain the role of context sensitivity in generating category mistake. Recall, for example, that 'That is green' or 'The thing I am thinking of is green' are category mistakes when the demonstrative or the definite description (are assumed to) pick out a number, but not when they (are assumed to) pick out a table. Relatedly, the lexical view cannot account for the way in which the phenomenology of category mistakes is relative to what participants in the conversation take for granted. Thus, for example, while in some contexts 'This woman fathered my children' would exhibit the phenomenology of category mistakes, in other contexts (ones where speakers have higher awareness of transgender women), no such phenomenology would arise.

My pragmatic presuppositional account also has the advantage of making use of the mechanism of presupposition accommodation, and the related notion of taking for granted for the purposes of conversation. This allows the account to handle some of the more intricate behaviours of category mistakes. For example, even in contexts where participants start out believing that numbers are not coloured, the presupposition can be accommodated for the purposes of conversation: an assertion of 'Two is green' can be felicitous in the context of a philosophy lecture-series which starts out with the speaker remarking that she will defend the

¹¹ See Magidor (2013), Chapter 2, §5-7.

controversial view that numbers are coloured, and in particular that two is green. And as I show, the degree of ease of accommodation nicely explains the gradability of category mistakes: ‘This rock is thinking about the theory of relativity’ isn’t as bad as ‘This number is thinking about the theory of relativity’.¹²

The lexical theory of types suffers from another flaw¹³: it ties together the infelicity of category mistakes with their truth-value. But this ignores the fact that when participants in conversation take for granted false claims, sentences which clearly exhibit the relevant phenomenology might nevertheless be not only truth-valued, but outright *true*. Thus, for example, it is plausible that prior to the discovery of the theory of relativity, ‘Space-time is curved’ would have been a fairly paradigmatic case of a category mistake. However, the sentence was nevertheless true rather than truth-valueless or false as lexical theories predict.

Asher’s theory of category mistakes is thus far from immune to the many worries I raise against competing accounts.

§1.2 co-predication

Consider the word ‘book’. A commonly accepted view is that this word is ambiguous: in one sense, it describes a *physical object*, while in another it describes an *informational object*. This ambiguity claim is most strongly supported by a difference in how we count or individuate books. If I own three copies of ‘War and Peace’ and no other books, there is a sense in which it is true to say that I own three books (the ‘physical object’ sense), and a sense in which it’s true to say that I only own one book (the ‘informational object’ sense).

However, the ambiguity hypothesis faces a challenge with so called cases of ‘co-predication’. Consider, for example, situations where one truly utters sentences of the following sort:

(11) John picked up the book and mastered it.

The problem is that, according to one common hypothesis, one can only pick up physical books and one can only master informational books, thus the conjunction is false on either interpretation of the definite description. In fact, on Asher’s way of construing things, matters are even worse: on his view ‘master’ *presupposes* that its argument is informational object, while ‘pick up’ presupposes that it is a physical object, so on either reading the sentence would constitute a category mistake, which – by Asher’s lights – means that the sentence suffers from an even deeper semantic defect than being false.

Without going into too much detail, Asher’s own account of sentences such as (11) goes roughly like this: The noun ‘book’ is not ambiguous. It has only one meaning, on which it belongs to neither the type ‘physical object’ (‘P’) nor the type ‘informational object’ (‘I’). Rather, it belongs to a special kind of type (‘dot-type’) which we can think of as a sort of ‘product’ of the two types: P•I (read ‘P dot I’). Now as Asher recognises this in itself does not address the problem with (11): since ‘book’ belongs to *neither* constituent type, it would now seem that the argument in (11) violates *both* the presupposition of ‘pick up’ and of ‘master’,

¹² See Magidor (2013), p. 152.

¹³ On this issue, see also the discussion in §2.2 of this paper.

so rather than being felicitous, the sentence is predicted to suffer from a double category mistake. To address this problem, Asher postulates a special structure to the sentence whereby instead of applying the predicates to the subject-term, each predicate is applied to a modulated version of it. This modulation occurs via a functor which maps the book to one of two “aspects” - a “physical aspect” in the case of ‘pick up’ and an “informational aspect” in the case of ‘master’. Underlying this proposal is not only a lot of non-trivial syntactic structure, but also some rather extravagant metaphysical assumptions. Asher explains for example that “aspects are, metaphysically speaking, a bare particular combined with some property or some property instance that it has”¹⁴; He maintains that dot-objects such as books are ‘wholly physical (under one aspect) and wholly informational (under another aspect). When we speak or think of books as physical, there is no “other part” of the book that is left out and that is informational’¹⁵. Moreover, since ‘book’ is not ambiguous on Asher’s theory, he also needs to explain why counting statements such as ‘I own three books’ apparently receive two incompatible readings. In response, Asher maintains that “counting requires a criterion of identity and individuation, which is part of the doctrine of relative identity”.¹⁶ Even this brief sketch should suffice to show why Asher’s account suffers from serious drawbacks: I do not wish to extend my own theory along similar lines.

Let me instead outline at least one alternative account of sentences such as (11). According to this account (call it ‘the simple account’) there is no ‘co-predication’ involved in (11). The noun ‘book’ is (as we assumed at the outset) ambiguous between the informational sense and the physical sense. However, as opposed to what Asher and Abrusán assume, it is *not* a category mistake to apply the predicate ‘master’ to a physical book (physical books can be strictly and literally mastered – even if they are mastered in virtue of an agent mastering the information that the physical book embodies). Nor is it a category mistake to apply the predicate ‘pick up’ to an informational book (informational books can be strictly and literally picked up, even if they are picked up in virtue of someone picking up a physical realisation of them).¹⁷ Sentence (11) thus does not involve any mid-sentence shift in meaning between the two senses of ‘book’. Rather the sentence is ambiguous between two readings, but this ambiguity is not apparent because the sentence has the same truth-value on both readings.¹⁸

Let me briefly address two potential objections to the simple account.

First objection: consider the following sentence.

¹⁴ Asher (2011), p. 149

¹⁵ See Asher (2011), p. 150. This is close to a direct quote except that for ease of exposition I changed the example in Asher’s text so that my discussion can focus on the same example throughout.

¹⁶ Asher (2011), p. 150.

¹⁷ Compare this to Liebesman’s view of generics (see Liebesman (2011)). According to Liebesman generics such as ‘Dogs bark’ strictly and literally apply the predicate ‘bark’ to the kind *Dog* (even though the kind barks in virtue of the barking behaviour of its members).

¹⁸ Why then postulate two readings at all? Precisely because counting statements such as ‘I own three books’ appears to receive different readings with conflicting truth-values.

Perhaps one could develop an alternative account on which ‘book’ in (11) is not ambiguous, and the difference in counting statements is handled by Liebesman’s proposal that counting is based on a relevant measure rather than on identity (see Liebesman (2015)). But the theory would still need to account for apparently true readings such as (12) and (13) below as well as for apparently true readings of ‘This is the same book as that one’ (when one is pointing to two copies of *Moby Dick*).

(12) Jane picked up every book in the library.

Suppose the library contains exactly four (physical) books – two copies of *War and Peace*, and two copies of *Things Fall Apart*. Suppose that Jane has picked up two physical books: one of each (informational) book. It seems that in these circumstances an utterance of (12) would be false. But, the objection goes, according to the simple theory, in addition to this false reading there should be an additional *true* reading of the sentence (one according to which ‘book’ means informational-book).

Response: I maintain that the sentence does in fact have an additional true reading. There are all sorts of reasons that certain readings of ambiguous sentences are not easy to access. One obvious reason in the current case is that in most contexts we have little interest in the question of how many informational books an agent has picked up. But consider, for example, the following context: a certain school library is trying to encourage its students to read more. However, since most students find reading very daunting they decide to start by encouraging students to simply *pick up* books (in the hope that after they are picked up, students might be tempted to read the books). Moreover, the school does not see much point to a student’s picking up two copies of the same (informational) book. To that end, they instate a weekly competition, in which each student gets a point for every (physical) book they pick up, as long as the student hasn’t yet picked up another copy of the same (informational) book. Now, suppose that Jane picked up one copy from each (informational) book in the library. It strikes me that in those circumstances, we *can* access the true reading of (12) (consider: ‘Jane is bound to win the competition this week – she picked up every book in the library!’).

Finally, note in presenting example (12) I have intentionally focused on a case where the reading involving quantification over informational-books is relatively hard to access, and thus the sentence might be thought to present a challenge to the simple theory. It is worth noting, however, that in other cases such second readings are much more salient. Thus consider an utterance of the following:¹⁹

(13) Last week, I picked up those three books from the local library.

This sentence seems to fairly clearly have a reading on which the books picked up are those very copies, as well as a reading on which the speaker picked up three different copies of those same informational books.

Second objection: According to the simple theory, it is not a category mistake to apply predicates such as ‘read’ or ‘master’ to physical objects such as physical books. But this suggests that applying the predicate to *other* physical objects should be felicitous as well. This prediction, however, seems wrong. For example, the following are arguably category mistakes:

(14)* Jane mastered the table,

(15)* Jane read the chair.

¹⁹ Thanks to David Liebesman for this example.

Response: Rather than telling against the simple account, this case provides yet another reason to prefer the adjectival-presuppositions that I postulate in the case of category mistakes over Asher's type-presuppositions. On my own account 'read' would trigger (roughly) the presupposition 'either read or unread' and 'master' (roughly) the presupposition 'capable of being mastered'.²⁰ This hypothesis predicts that the predicates can felicitously apply to some physical objects (e.g. books) but not to others (e.g. tables or chairs).

No doubt these brief remarks do not amount to anything like a proper theory of sentences such as (11).²¹ But I raise them to illustrate that it is far from clear that such cases have anything in particular to do with category mistakes.

§1.3 coercion

A different phenomenon that Abrusán discusses is that of 'coercion'. Let me focus on the most interesting case, that of 'event coercion'. Consider the following two sentences:

(16) Jane finished reading the book.

(17) Jane finished the book.

According to Abrusán (following Asher), 'finished' presupposes that its (internal) argument is an event. In the case of (16), this presupposition is satisfied. But since in (17) the internal argument (the book) is *not* an event, the sentence should have resulted in a category mistake. The challenge is to explain why, contrary to this prediction, (17) is perfectly acceptable. Asher's response is to argue, first, that 'finished' does not quite presuppose that its argument is an event. Rather, the internal argument can be either an event or a functor which maps two objects (Jane and the book, in this case) onto an event (Jane's reading of the book, in this case). The material in (17) together with the context somehow generates the relevant functor as well as the objects it takes as arguments to get (17) to receive the same reading as (16).

As with the cases of co-prediction, I wish to reject the assumption that (17) involves any kind of tacit category mistake (one which is then repaired by some special mechanism). A much more plausible hypothesis seems to be that 'finish' has, in addition to the reading which takes an event as its internal argument, a second reading which takes a standard object as its internal argument (strictly and literally, with no coercion involved).²² In fact, it is not unusual for verbs to take either an event or an object as arguments. The following, for example, are all felicitous:

(18a) John imagined Mary.

(18b) John imagined seeing Mary.

²⁰ See Magidor (2013), Ch. 5, §4.

²¹ In particular, one would need to say a lot more on the metaphysics of informational objects, and under which circumstances various properties apply to them.

²² A subtly different hypothesis is that 'finish' has just one reading, which takes either an event or an object as an argument. I will not adjudicate between these two views here.

(19a) Jill loves Mary.

(19b) Jill loves hearing Mary play the drums.

(20a) John fears the snake.

(20b) John fears stepping on the snake.

I see no reason for taking (18a), (19a) and (20a) to involve, if taken at face-value, some kind of implicit category mistake.

Abrusán also discusses a related interesting issue, namely the contrast between the following two sentences:

(21) Jane finished the book.

(22) Jane stopped the book.

Although both sentences have true readings, they are not used in quite the same way. While (20) can be used to describe an event where, for example, Jane finished reading the book, (21) cannot be used to mean that Jane stopped reading the book (though it can be used to describe events where, for example, Jane stopped the book from falling).

The contrast between (21) and (22) seems to belong to the very rich and complex topic of *argument realization*: the question of which semantic-roles arguments play, and how those semantic-roles are represented in syntax.²³ To get a sense of this complex issue, consider another example, which Asher also discusses:²⁴

(23) John swept the closets.

(24) John swept the leaves.

Although superficially the sentences seem to have a very similar structure, on the most natural readings of (23) and (24) the two internal arguments play very different roles (even if it is extremely hard to characterise this difference explicitly), and indeed on their most natural readings one cannot describe the conjunction of (23) and (24) by saying that John swept the closets and the leaves. Asher, in his discussion, attributes this difference to different types associated with the lexical items ‘leaf’ and ‘closet’. But this cannot be right because each sentence *also* has a second reading, one where the semantic roles of ‘closets’ and ‘leaves’ are exactly reversed. Thus we can use (23) to describe a situation where John is in a room full of tiny closets that are cluttering the room and need to be cleared away using a broom, and (24) can be used to describe a situation where John is in charge of some extremely precious leaves, and he is required to gently sweep the dust off them.²⁵ More likely, the difference between the most natural readings of (23) and (24) has to do with a subtle difference in their

²³ See Levin and Rappaport-Hovav (2005).

²⁴ Asher (2011), pp. 32-33.

²⁵ Indeed, on those readings, the conjunction test works as we would expect: if John is sweeping the closets in the standard sense and sweeping the dust off the highly precious leaves, we can describe his action as ‘sweeping the closets and the leaves’.

syntactic structures, which in turn entails that their respective internal arguments occupy a subtly different semantic-role.

Similarly, I take it that the internal arguments in (21) and (22) play subtly different semantic roles. Moreover, these differences also seem reflected in much more overt syntactic differences between ‘finish’ and ‘stop’. For example, ‘stop’ but not ‘finish’ allows for the causative alternation, so that (25b) is grammatical, but (26b) not:

(25a) Jane stopped the book (from falling).

(25b) The book stopped.

(26a) Jane finished the book.

(26b)* The book finished.

Here too, I have not offered anything like a full account of the difference between ‘stop’ and ‘finished’. But it is worth asking whether the Asher’s account, as presented by Abrusán, offers any real explanation of the difference either. According to this account, while ‘finished’ accepts as a functor from a pair of objects to events, ‘stop’ accepts a different kind functor – one which maps one argument (in this case, the internal argument of ‘stop’) to an event. But in order to explain the difference between the two verbs, one needs not only to stipulate these rather different type-presuppositions, but also explain why the functor in the case of ‘stop’ cannot, in context, map the internal argument (the book) onto the event of Jane’s stopping the reading of the book.²⁶

It is thus far from clear that type presuppositions in particular, or the phenomenon of category mistakes in general, plays any useful role in accounting for the difference between (21) and (22).

§1.4 conclusion

Unified accounts are attractive as they make for theoretical parsimony. But while the discussion in Asher (2011) encompasses an impressive variety of phenomena, his theories are far from parsimonious. This is partially for somewhat technical reasons (for example, as I point in note 9, he needs to explicitly represent types in the semantics rather than read them off ordinary semantic-values, and the kinds of ‘accommodation’ that Asher appeals to are not the standard mechanisms of presupposition accommodation). But the account also falls short on parsimony because it appeals to extremely heavy resources (to name just a few: unpronounced syntactic material, topos theory, bare particulars, and relative identity). Most importantly, at least in the case of category mistake, I do not think Asher provides a

²⁶ Abrusán says this is because “since the dependant type lacks an agent argument, the defeasible value for an eventuality involving physical objects will be licensed, which is assumed to be a physical process. So the presupposition to be accommodated is that [the book] was involved in some physical process” [REF]. But I do not see why in context, we cannot assume the process in question is that of reading (rather than a physical process); nor does this explain why we cannot use sentences such as (22) to describe, e.g., the physical process in which Jane stops falling over the book.

successful account of the phenomenon. Alas, sometimes a wide variety of phenomena require diverse explanations.

§2. Response to Shaw

James Shaw’s paper focuses on my rejection of the ‘MBT view’: the claim that category mistakes are meaningful but truth-valueless. He raises concerns about two of my arguments: a general argument I present against the MBT’s view’s commitment to truth-value gaps²⁷, and an objection I raise against the claim that the infelicity of category mistakes provides a good motivation for adopting the MBT view.²⁸ I discuss each of these in turn in §2.1 and §2.2 below.

§2.1 Shaw on MBT’s commitment to gaps

§2.1.1 The attractiveness of the necessitated Tarskian biconditionals

In §2 of chapter 4 of the book I present a general argument against the MBT view. My argument proceeds in two stages: in stage one, I show that defenders of the MBT view cannot rely exclusively on a model where category mistakes are truth-valueless in virtue of failing to express a proposition. Rather, cases of ‘contingent category mistakes’ (sentences such as ‘The thing I am thinking of is green’) force defenders of the view to accept that (at least some) category mistakes do express propositions – albeit ones that are either truth-valueless, or at least *possibly* truth-valueless. In the second stage, I offer a modal extension of an argument due to Williamson, in order to show that accepting such ‘partial propositions’ (propositions that are truth-valueless relative to some possible worlds) comes at a high theoretical price.²⁹ I show that one can derive a contradiction from the assumption that a given proposition is partial, using a very weak propositional modal logic coupled and the highly attractive necessitated versions of the Tarskian biconditionals:³⁰

(NT) $\forall P$, Necessarily, it is true that P if and only if P.

(NF) $\forall P$, Necessarily, it is false that P if and only if $\neg P$.

Shaw accepts the first stage of my argument (the claim that defenders of the MBT are committed to the existence partial propositions) but objects its second stage. He starts by pointing out that, despite their initial attractiveness, defenders of the MBT view can ultimately opt to reject (or at least fail to accept) the biconditionals or alternatively eschew even the weak logic utilised in the argument. This is a point I fully acknowledge in my

²⁷ Magidor (2013), Chapter 4, §2.

²⁸ Magidor (2013), Chapter 4, §3.1.

²⁹ Incidentally, Shaw seems to misattribute my modal extension to Williamson.

³⁰ The quantification here should be understood as quantification into sentence position. As I understand such quantification, only sentences of the form ‘It is true that p if and only if p’ where ‘p’ expresses a proposition count as instances of the universal quantifier in NT (similarly for NF). This strikes me as the most plausible way of understanding such quantification (after all, we don’t standardly take ‘Fa’ to be an instance of ‘ $\forall xFx$ ’ where ‘a’ doesn’t refer). Moreover, since the argument is directed against MBT-supporters who accept that category mistakes express propositions this assumption is innocuous in context.

discussion: the problem, as I note, is that either of these moves comes at a high theoretical price. At the very least, my argument reveals that offering a semantics based on partial propositions is not (as the linguistics literature often suggests) merely a matter of jotting down some clauses for how to compositionally compute truth-values in the face of truth-value gaps, but rather come with much deeper foundational challenges.

However, Shaw proceeds to press his point more resolutely: he suggests that eschewing (NT) and (NF) does not even come at a particularly high price. The principles, he maintains, are not particularly attractive. Rather, they only seem attractive when we restrict ourselves to mundane instances such as the following:

(1) Necessarily, it is true that it is raining if and only if it is raining.

However, he suggests, when we consider precisely those cases that are pertinent to the discussion in question, the biconditionals do not seem particularly plausible. Examples such as (2) strike us as infelicitous rather than straightforwardly true:

(2) * Necessarily, it true that the number two is green if and only if the number true is green.

Shaw's objection suffers from a number of difficulties. First his line of argument misleadingly suggests a picture according to which (NT) and (NF) seem attractive only because we consider many of their mundane instances, and then – rashly, on this view – generalise (perhaps by analogy to observing a large number of black ravens and then rashly concluding that all ravens are black, foolishly ignoring the few albino ones...). But this is misconstruing the epistemic role of examples such as (1) in motivating the biconditionals. A quick reflection on a very small number of examples suffices to draw out the attractiveness of (NT) and (NF). The fact that one can immediately recognise the truth of (1) without any knowledge about the truth-value of either of the two statements flanking the biconditional or indeed without possessing any knowledge of meteorology, strongly suggests that the statement is true in virtue of very general features pertaining to the nature of truth.

Second, while I share Shaw's judgement that (2) is infelicitous, this judgement does little to undermine (NT), because this infelicity is entirely predicted by a range of views that fully accept the biconditionals. Thus, for example, those who maintain that category mistakes do not express propositions at all can accept (NT) while maintaining that (2) is infelicitous and truth-valueless (on this view, of course, (2) isn't an instance of (NT) – see n.30 above). More importantly, my own pragmatic account of category mistakes predicts precisely that (2) is infelicitous (despite being true): the biconditional generates the presupposition that the number two is coloured, a presupposition that in most contexts we do not accept and cannot accommodate.³¹

³¹ Probably anticipating this response, Shaw claims that (2) sounds considerably more infelicitous than other cases of conditionals where the antecedent generates a false presupposition ([Ref]). I'm not sure I share his judgement about the particular example he presents, but at any rate it is worth noting that the presupposition failures associated with category mistakes often sound worse than other cases of presupposition failure (for

Indeed, reflecting on this last point suggests that the kind of example that Shaw should have considered here is precisely the case that is at the centre of the second stage of my argument. Consider a context in which the thing I am thinking of is the table next to me (and it is common ground that I am thinking about the table). Now consider an utterance of (3):

- (3) Necessarily, it true that the thing I am thinking of is green if and only if the thing I am thinking of is green.

Since Shaw maintains that the embedded sentence in (3) expresses a partial proposition, his view predicts that (3) should be just as unattractive as (2). My own view, on the other hand, predicts that since (3) is both true and does not suffer from any infelicity, it is entirely acceptable. It seems to me that the data here supports my own prediction here over Shaw's.

§2.1.2 Contingent semantic paradoxes

Shaw goes on to discuss what is no doubt the most difficult challenge for the view that denies partial propositions: whether this view can handle the semantic paradoxes, and in particular contingent semantic paradoxes. This is a topic that I touched on very briefly in the book³² but Shaw presses the point by raising an interesting argument from the semantic paradox literature involving propositional attitude reports.³³

Before I turn to discuss the argument, let me make one thing clear at the outset: providing a full response to this case requires a complete theory the semantic paradoxes – something I obviously do not purport to offer. (Of course, it should be noted that merely accepting partial propositions doesn't in itself give us anything like a complete theory of the semantic paradoxes either!). However, I do want to say what stance, in broad terms, those who reject partial propositions should take on this sort of case.

The argument concerning propositional attitude ascriptions proceeds as follows. Suppose that Lia thinks that Jane is thinking that 21 is a prime number (Lia is right that this is what Jane is thinking and knows that 21 is not prime). But Lia also thinks that Jane is the tallest person in the room (in fact, here Lia is wrong: Lia herself is the tallest person in the room). Moreover, putting these beliefs together, Lia endorses the sentence 'What the tallest person in room is currently thinking is not true'. It is tempting to think in this kind of situation the following ascription is true:³⁴

- (4) Lia is thinking that what the tallest person in the room is currently thinking is not true.

example 'Relativity drinks procrastination' is arguably much more infelicitous than 'The king of France is bald').

³² Magidor (2013), p. 88, n.15.

³³ See for example Field (2008), pp.132-133.

³⁴ Of course, at least one condition for this ascription to be true is that Lia is only thinking one thing at the relevant time. Even if we restrict the description to thoughts she is having by virtue of endorsing the relevant sentence this assumption is far from trivial – the very threat of paradox might entail that she is thinking more than one thing (cf. Prior's discussion of the Liar in Prior (1971), ch.6). For the remainder of the discussion, I will assume that Lia is thinking at most one thing in the context.

However, plausibly, for a propositional attitude ascription to be true requires its complement clause to successfully express a proposition. It follows that (5) expresses a proposition - call it p :

(5) What the tallest person in the room is currently thinking is not true.

The problem, however, is that if Lia is the tallest person in the room and the thing that Lia is currently thinking is p , then p cannot receive a truth-value on pain of paradox - thus entailing that p is a partial proposition and hence that partial propositions exist.

As Shaw notes, one way for deniers of partial propositions to resist this arguments is to maintain that (4) can be true in this context even if the complement clause fails to express a proposition. For example, certain versions of the Fregean view of attitudes could maintain that (4) expresses a relation between Jane and the *sense* of the embedded sentence, while denying that the embedded sentence expresses a *content*.³⁵ A different (and perhaps more promising?) way to sever the link between the truth of the ascription (4) and a commitment to partial propositions is to maintain that while the truth of (4) requires there to be *some* particular proposition which Lia stands in the thinking relation to, that proposition need not be one that is semantically expressed by (5). For example, a certain kind of contextualism about attitudes can maintain that it is a matter for context to resolve which proposition is reported as being thought by Lia in (4). Perhaps, for example, in this context, Lia is described as thinking that what the person that *she thinks is tallest* in the room is thinking is not true (on this view, the very threat of paradox can play a part in the contextual resolution).

But each of these responses comes with a commitment to fairly unattractive theories of propositional attitude ascriptions. Rather, I think the most promising avenue for a denier of partial propositions is to deny that (4) is true. To understand this position better, it is helpful to start with a slightly different case (one closer to the two cases of paradox I do briefly discuss in the book).

Suppose that the actual world is just as is in the case above, except that it really is *Jane* (rather than Lia) who is the tallest person in the room. I think in this case we should certainly accept that (4) is true, and indeed that (5) succeeds in expression a proposition – call it p^* . Now, let us ask ourselves what we should think about a close possible world w in which it is *Lia* (rather than Jane) who is the tallest person in the room. My claim is that on pain of paradox, in w , Lia cannot be thinking p^* . Of course, Lia can still have lots of closely connected thoughts (e.g. the thought that Jane is the tallest person in the room and what Jane is thinking is not true, or the thought that what the person that Lia believes is the tallest person in the room is thinking is not true). But if Lia is currently thinking any proposition in w , it is not p^* .

³⁵ One might be tempted to think that this feature can also be achieved on a sententialist view of propositional attitude ascriptions (see Davidson (1968)). But even if the Davidsonian view does not technically require such complement clauses to ‘express a proposition’, it does require them to be contentful which, on my view, will require them to conform to the Taraskian bicodntionals. (Cf the discussion in Magidor (2013), pp. 65-66).

This is certainly a surprising result, but it is in the nature of paradox to lead to surprising conclusions. Moreover, note that this result is not so different from other, less remarkable, cases which arguably involve contingently available thoughts. For examples, consider the proposition that Ofra is currently writing a paper (call it p^{**}). Many philosophers maintain that I only contingently exist, and many of those maintain that relative to worlds in which I do not exist no one can think p^{**} .³⁶ Indeed, we can imagine a case where in the actual world, George thinks that Ofra is writing a paper (i.e. George is thinking p^{**}). Yet, there is another possible world w in which I do not exist, and hence in which George cannot think p^{**} . Moreover, in some such worlds George will be an *internal duplicate* of actual-George (in such worlds, George might have misleading evidence about a philosopher called ‘Ofra’ writing a paper, etc.), and while it might be tempting to think that such internal duplication requires George to have the same thoughts in both worlds, on the contingentist supposition laid out, this cannot be right.³⁷

Returning to Shaw’s original case, it is clear what I wish to say. We are tempted to accept (4) because there are very similar worlds to the actual one (worlds where Jane is tallest, and perhaps Lia is an internal duplicate of actual-Lia) where utterances of (4) would have been true, and where they would have ascribed to Lia a thought such as p^* . Nevertheless, we should resist the temptation to infer from this that (4) is actually true.

§2.2.2 The infelicity argument

The role of §3 of Chapter 4 was to discuss a range of positive arguments in support of the MBT view and argue that they are not ultimately compelling. The first of these is the ‘infelicity argument’ – the claim that we should adopt the MBT view because it provides an attractive account of the infelicity of category mistakes. (Indeed, as I have laid things out in the book, the various views presented were taken to be primarily responses to the question of what accounts for the infelicity of category mistakes.)

As Shaw recognises, my main reason to reject this argument is that I do not think the MBT view is necessary to account for the infelicity: my own pragmatic account provides, I argue, a compelling explanation of this infelicity which does not require category mistakes to be truth-valueless. However, I also note that the MBT’s account of infelicity is at a disadvantage because of the following feature: whether sentences exhibit the relevant phenomenology of infelicity often depends on what participants in the conversation *assume* or *believe*. This in turn entails that (on the standard way of developing the view) the MBT view often predicts a divergence between whether a sentence is infelicitous in the relevant way and whether it is truth-valueless. This can be seen in the following examples:

³⁶ Admittedly, one disanalogy with the paradoxical case is that here contingentists would presumably argue that in such worlds p^{**} does not exist, whereas in my example p^* exists but is not available to be thought by Lia in the relevant world. But see n.37 below for an example without this disanalogy.

³⁷ Another example of “unthinkability” (one which does not involve the controversial commitment to contingentism) is the following: in the actual world George is thinking that Tigers have stripes, but in another world w , George is an internal duplicate of actual-George who has had no casual contact with the natural kind Tiger. Thus in w George cannot be thinking the same proposition.

- (i) Context: Participants in the conversation take for granted that the thing Jane is thinking of is a table but in fact she is thinking of the number two.
Utterance: ‘What Jane is thinking of is green’
Data: the utterance is felicitous
MBT’s claim: the utterance is truth-valueless.
- (ii) Context: Participants in the conversation take for granted that the thing Jane is thinking of is the number two, but in fact she is thinking of is a table.
Utterance: ‘The thing Jane is thinking of is green’
Data: the utterance is infelicitous
MBT’s claim: the utterance is truth-valued.
- (iii) Context: participants in the conversation take for granted that men aren’t able to be pregnant, but in fact they are wrong as transgender men can be pregnant.
Utterance: ‘This man is pregnant’
Data: The utterance is infelicitous
MBT’s claim: the utterance is truth-valued.

In response, Shaw follows a natural proposal (one that that I myself suggest in a footnote), namely that the MBT view can account for infelicity in the following way: a sentence is infelicitous just in case speakers *think* (perhaps falsely) that it expresses a truth-valueless proposition. Moreover, Shaw complains that my argument here is inconsistent with some of my other commitments. After all, I accept that contingent semantic paradoxes or sentences with an empty demonstrative fail to express a proposition even when we mistakenly assume that the relevant utterance is not paradoxical or the demonstrative in question is successful. In such cases, the utterances would be perfectly felicitous despite suffering from expression failure: thus by my own commitments infelicity and “semantic defectiveness” can sometimes come apart.

Shaw’s complaint however misconstrues the dialectical role of my argument. I did not claim that MBT’s commitments were *inconsistent* with the infelicity data. Rather, my point was that in so far as the MBT’s primary motivation was to provide an explanation of this infelicity data, it does not provide the simplest or best explanation. For one thing, once the infelicity is accounted for in terms of speakers *thinking* (rightly or wrongly) that the relevant sentences are truth-valueless it is not even clear why this explanation requires there to be any actual truth-value gaps. For another, at least cases such as (iii) seem to suggest that this sort of account of infelicity constitutes a kind of error-theory about category mistakes. After all, in cases such as (iii) the speakers do not make any mistake about what the ‘this man’ or ‘is pregnant’ means or refers to in context, yet they still mistakenly take the sentence to suffer from the relevant semantic defect. If one compares my own account which directly predicts the infelicity/felicity data in examples such as (i)-(iii) with the MBT’s account which explains the data not only by requiring a revisionary semantics but also by relying on a sort of error-theory, it seems to me that my own theory provides the better explanation.

This also explains why my argument is not inconsistent with my views on empty demonstrative or the semantic paradox. Claiming that there is a semantic defect in paradoxical sentences is not motivated first and foremost as an attempt to explain linguistic infelicity. That the Liar is a semantically defective sentence is a conclusion we reach not

because the sentence sounds funny, but rather because (non-trivial) theoretical reflection reveals that trying to assign it a truth-value results in contradiction. And even in the case of the empty demonstrative, the argument for expression-failure is far more theoretically loaded than merely accounting for the oddness of utterances of ‘That is green’ when one is pointing to thin air. (Among other things, the argument presumably involves claims about the compositional structure of such sentences, which in turn entail that in order for them to express a proposition each term in the sentence should be contentful...).

My claim is thus not that a divergence of infelicity data and truth-value gaps is in principle untenable. Rather, my claim is that given the divergence, truth-value gaps do not provide a very straightforward explanation of infelicity. Coupled with the theoretical challenges for accepting partial propositions, the MBT does not provide the best account of category mistakes.

§3 Response to Elbourne

Paul Elbourne’s paper offers an interesting novel argument *in favour* of my account of category mistakes. While I obviously welcome this proposal (and find this sort of constructive contribution to be a refreshing addition to the standard critical modes of philosophical engagement), I think it is somewhat less straightforward than he suggests for my view to account for the cases he discusses. Nevertheless, as I explain below, I do agree that his argument ultimately renders support for my account.

The examples Elbourne discusses involve multi-sentence discourses which seem to exhibit very similar infelicity phenomenology to that of classic, single-sentence category mistakes. Consider the following examples:

- (1) * The thing John just mentioned is green. The thing he mentioned is the number two.
- (2) * The number two won our annual contest. Our annual contest is a drinking contest.

Before proceeding, let me make two preliminary comments about Elbourne’s examples. First, one of his key examples involves the following discourse:

- (3) The number two has the property that John just mentioned. The property that John just mentioned is the property of being green.

While my intuitions are not entirely clear, I am not sure this is the best example for the phenomenon he wants to draw out, simply because it is not clear to me that the single-sentence analogue of this example exhibits the kind of infelicity associated with category mistakes. While ‘The number two has the property of being green’ is not in itself acceptable, it is not obvious that it passes the usual projection tests (‘The number two does not have the property of being green’ at least strikes me as much better than ‘The number two isn’t green’). And even setting aside the potential category mistakes, sentences of this form not entirely natural: ‘My table has the property of being green’ isn’t a common locution in non-philosophical English. I thus prefer to avoid discussing this particular example and focus on cases such as (1) and (2).

The second point is that (1) and (2) suffer from additional noise which results from the fact that the second sentence in each example directly contradicts the first (issues to do with category mistakes aside, there is no possible world relative to which both are true at the same time). Interestingly, however, we get the exactly the same infelicity data with the following sentences, which lack this problematic feature:

(1b) * The thing John just mentioned isn't green. John just mentioned the number two.

(2b) * The number two didn't win our annual contest. Our annual contest is a drinking contest.

As Elbourne correctly argues, both the syntactic approach and the meaninglessness view of category mistakes would have a difficult time accounting for the infelicity of such examples: after all, each discourse consists of two sentences, each of which on its own is clearly both syntactically well-formed and meaningful.

That leaves in play the MBT view and my own pragmatic approach. Elbourne begins by convincingly arguing that the MBT's only plausible option is to maintain that the first sentence (and only the first sentence) in each pair is truth-valueless. This already raises an issue (related to the one discussed in §2.2 above), concerning the fact that the initial utterance of the first sentence in each discourse is, even if truth-valueless, not in itself infelicitous. But following the proposal discussed in §2.2, we can grant the MBT view the following analysis of the case: when speakers initially hear the utterance of the first sentence in these examples, they (perhaps falsely) assume it is true. However, after hearing the second sentence, they go back and revise their initial assessment, now rendering that sentence truth-valueless, and thus in hindsight infelicitous.

Elbourne assumes my own pragmatic approach has a straightforward account of the case, one which takes only the utterance of the second sentence to be directly infelicitous, without any need to go back and reassess the first utterance (I will return to this point shortly). His crucial observation is that that the two proposals make different empirical predictions concerning the processing of these two discourses: according to my pragmatic approach (but not according to the MBT view), the infelicity should arise before any extra re-processing of the first utterance occurs.³⁸ Even more strikingly, he maintains that there already are experiments involving ERPs ('event-related brain potentials') supporting the pragmatic approach's prediction over MBT's prediction. Apparently, speakers respond to infelicity (whatever its source) by heightened neurological activity, which causes measurable electromagnetic effects on the scalp. Moreover, Elbourne cites a range of such experiments that suggest that in

³⁸ The connection between semantic theorising and psychological processing is a highly complex one, and I am often weary of drawing quick conclusions based on processing results (see Magidor (2010), §2 and Magidor (2013), p.67). Nevertheless, I am sympathetic to Elbourne's remark in n. 16 [Ref]: the infelicity of category mistakes is primarily a psychological phenomenon, and I don't think this particular use of processing data is especially objectionable.

examples such as above, the infelicity occurs at exactly the same time as speakers process the second sentence, thus suggesting no re-processing is involved in generating the infelicity.³⁹

Unfortunately (for me), it is not quite as straightforward as Elbourne assumes for my own view to predict the relevant infelicities. According to my account, category mistakes arise when (certain kinds of) presuppositions are generated by an utterances, and when those presuppositions are neither taken for granted in the conversation nor easily accommodated. The problem, however, is that second sentence in each pair does not in itself generate any presuppositions (of the relevant kind), and thus the infelicity generated by these utterances cannot be exactly assimilated to my general theory of category mistakes.

Nevertheless, my view does predict that the second utterance in each pair should be infelicitous, and indeed infelicitous without any need for any re-processing. For example, in discourses (1) and (1a), the utterance of the first sentence in the discourse generates the presupposition that the thing John has just mentioned is coloured. Since at this point in the conversation participants have no particular reasons to deny this, I take it that the presupposition is accommodated (i.e. added to the common ground). This in turn means that by the time the second sentence in each discourse is processed, it directly contradicts the common ground: the second utterance in each case is thus taken for granted to be false, and hence infelicitous.

Where does this leave us? The ERP experiments that Elbourne discussed do not distinguish between different sources of infelicity (trivial falsity triggers the same kind of response as other sources of anomaly). This does mean that my own view fares well (certainly better than standard MBT views) in predicting these experimental results.

However, I tend to believe that the infelicity associated with presupposition failures has a fairly distinctive phenomenology, one that separates it from mere trivial falsity⁴⁰ (roughly: presupposition failures have the oddness of not just asserting something controversial, but rather assuming it as background). Thus *if* the second utterances in each of these discourses triggers this distinctive phenomenology, that is not quite predicted by my view (because, to repeat, the second utterances in each discourse does not suffer from the relevant kind of presupposition failure). Of course, intuitions about (let alone empirical tests of) the particular nature of infelicities are quite fragile. This is especially so in these cases, where there presumably can be an overdetermination of sources of infelicity: an initial reaction due to the second utterance contradicting the common ground, as well as perhaps, an additional infelicity due to the reprocessing of a very recent presupposition. Thus although the issue is

³⁹ As Elbourne notes ([REF]) the exact examples we are discussing have not been tested, so there is some room for doubt about the relevance of the results cited. However, the cases that have been tested do seem sufficiently similar to suggest at least a preliminary conclusion. (For the remainder of the paper I will adopt the working hypothesis that testing the particular examples in this paper will yield similar results – though it would certainly be interesting to actually test this hypothesis.). At any rate, I take the most interesting point Elbourne makes to be that the diverging predictions of the two theories are in fact empirically testable in practice.

⁴⁰ See Magidor (2013) , p. 113.

not straightforward, I do think that Elbourne's case ultimately provides a consideration in favour of my theory.⁴¹

A final remark is in order: the discussions of the MBT view above assume a standard version of the theory – one that relies on the gaps to account for infelicity and takes truth-values to depend on standard semantic-values for terms in each sentence. As I briefly note in the book, there are other possible versions of the MBT view. For example, one can have a theory of category mistakes which, just like my own, relies on presuppositional account of infelicity. Moreover, this account also relies on a Stalnakerian notion of context but nevertheless maintains that presupposition failures make for truth-value gaps.⁴² In terms of *infelicity* this view makes very similar predictions to my own view, and can thus also accommodate the ERP results that Elbourne discusses (an utterance of the first sentence in (1) generates the presupposition that the thing John mentioned is coloured, this presupposition is then accommodated, and contradicted by the second sentence in the discourse...). Elbourne's argument thus does not tell against this particular kind of an MBT view, though there are other good reasons to reject it.⁴³

References

- Asher, N. (2011), *Lexical Meaning in Context: A Web of Words*, Cambridge: Cambridge University Press.
- Davidson, D. (1968), 'On saying that', *Synthese* 19: 130-146.
- Field, H. (2008), *Saving Truth from Paradox*, Oxford: Oxford University Press.
- Liebman, D. (2011), 'Simple generics', *Nous* 45: 409-442.
- Liebman, D. (2015), 'We do not count by identity', *Australasian Journal of Philosophy* 63: 21-42.
- Levin, B. and Rappaport-Hovav, M. (2005), *Argument Realization*, Cambridge: Cambridge University Press
- Magidor, O. (2010), 'Natural Language and how we use it: psychology, pragmatics, and presupposition', *Analysis* 70: 160-74.
- Magidor, O. (2013), *Category Mistakes*, Oxford: Oxford University Press.
- Prior, A. N. (1971), *Objects of Thought*, Oxford: Clarendon Press.

⁴¹ Moreover, it should be noted that similar examples involving standard presuppositions exhibit a very similar phenomenology. Thus for example: 'John stopped smoking. He never did smoke' is odd in at least a fairly similar manner to 'John has never smoked and he stopped smoking' (cf. Elbourne [REF]). Thus whatever the explanation of the infelicity in multi-sentence discourse cases, it seems that there is no disanalogy between category mistakes and other cases of presupposition failure.

⁴² See Magidor (2013), p. 82 and pp. 130-131, n. 29.

⁴³ Its commitment to partial propositions is one reason to reject this view. Another, is that it has the odd consequence that, for example, when uttered relative to a context where speakers falsely take for granted that England lacks a queen, the sentence 'The queen of England in 2016 is called 'Elizabeth'' is truth-valueless rather than true. See Magidor (2013), n. 29.