



Category mistakes are meaningful

Ofra Magidor

► To cite this version:

Ofra Magidor. Category mistakes are meaningful. *Linguistics and Philosophy*, 2010, 32 (6), pp.553-581. 10.1007/s10988-010-9067-0 . hal-00567228

HAL Id: hal-00567228

<https://hal.science/hal-00567228>

Submitted on 19 Feb 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Category mistakes are meaningful

Ofra Magidor

Published online: 19 February 2010
© Springer Science+Business Media B.V. 2010

Abstract Category mistakes are sentences such as ‘Colourless green ideas sleep furiously’ or ‘The theory of relativity is eating breakfast’. Such sentences are highly anomalous, and this has led a large number of linguists and philosophers to conclude that they are meaningless (call this ‘the meaninglessness view’). In this paper I argue that the meaninglessness view is incorrect and category mistakes are meaningful. I provide four arguments against the meaninglessness view: in Sect. 2, an argument concerning compositionality with respect to category mistakes; in Sect. 3 an argument concerning synonymy facts of category mistakes; in Sect. 4 concerning embeddings of category mistakes in propositional attitude ascriptions; and in Sect. 5 concerning the uses of category mistakes in metaphors. Having presented these arguments, in Sect. 6 I briefly discuss some of the positive motivations for accepting the meaninglessness view and argue that they are unconvincing. I conclude that the meaninglessness view ought to be rejected.

Keywords Category mistakes · Selectional restrictions · Selectional violations · Compositionality · Semantics · Foundations of semantics · Montague Grammar · Type theoretic semantics · Colorless green ideas sleep furiously · Meaning · Meaningfulness · Meaninglessness · Nonsense

Introduction

My focus in this paper is on sentences of a certain kind that are often categorised as nonsense: category mistakes. Category mistakes are sentences such as ‘Colourless green ideas sleep furiously’, ‘The number two is blue’, or ‘The theory of relativity is eating breakfast’.¹

¹ For the purposes of this paper I will take the concept of a category mistake to be delimited phenomenologically, as sentences which are anomalous in a similar sort of way to the above-mentioned example sentences. The question of whether there is a more rigorous or substantive way to define the notion of a category mistake lies beyond the scope of this paper.

O. Magidor (✉)
Balliol College, University of Oxford, Oxford OX1 3BJ, UK
e-mail: ofra.magidor@philosophy.ox.ac.uk

Category mistakes are obviously highly anomalous sentences, and this fact has led a large number of philosophers and linguists over the years to claim that they are meaningless (call this ‘*the meaninglessness view*’). Indeed, although the meaninglessness view does not enjoy these days quite the vast popularity that it did for a large part of the twentieth century, endorsements of the view keep resurfacing.

Thus for example, in a recent article discussing an argument for the distinctness of a material thing and its matter which relies on predicational differences, Kit Fine says: ‘It is worth emphasizing, in this connection, that these differences lie not merely in the *correct* but also in the *meaningful* application of the predicates. A chair can meaningfully be said to be *comfortable* or *uncomfortable*, though not the wood from which it is made; a statue can meaningfully be said to be *Romanesque* or *Romanesque*, though not the clay or the alloy itself; and one can meaningfully be said to *spend* a penny or a dollar, though not some metal or paper’.² In a similar style of argument, Helen Steward argues that processes are distinct from events, because ‘“processes have properties which it would be inappropriate to ascribe to events, and vice versa. For example [...] the humming of my computer in the process sense can be persistent; but *it does not really make sense* to think of an event as persistent’.³ Timothy Williamson maintains that ‘“To apply the concept of vagueness to anything other than a representation may be treated as a category mistake... [this claim] makes ‘The facts are precise’ and ‘The facts are vague’ [...] equally meaningless’.⁴ And more directly, Beall and Van Fraassen explain in their contemporary book that ‘[T]here are different ways in which a (declarative) sentence might properly be called ‘meaningless’. Perhaps the best example involves so-called category mistakes’.⁵

Moreover, it is not difficult to see the appeal of the meaninglessness view. First and most importantly, given that category mistakes are highly anomalous (more so than standardly false sentence) it is very intuitive to classify them as meaningless. In fact, the meaninglessness view appears to have the advantage of explaining not only why category mistakes exhibit such an extreme anomaly, but also what is distinctive about their anomaly. For if category mistakes are meaningless then arguably they are the only kind of sentence which is syntactically well-formed but nevertheless meaningless, thus accounting for the distinctiveness of their anomaly.⁶ The view might also be motivated on more theoretical grounds. For example, one might think that accepting a truth-conditional theory of meaning commits them to the meaningless view because it is unclear under what conditions sentences such as ‘Two is

² Fine (2003, pp. 207–208) (His emphases).

³ Steward (1997, p. 96) (My emphasis).

⁴ Williamson (1994, pp. 249–250). It is worth noting, though, that Williamson here is only committing to the claim that *if* a sentence is classified as a category mistake it is meaningless. He is not committed to the claim that the example in question (or indeed any other example) is in fact a category mistake.

⁵ Beall and van Fraassen (2003, p. 125). For other endorsements of the meaninglessness view see for example Russell (1908), Ryle (1938), Strawson (1952), Chomsky (1957), Routley (1966), Drange (1966), Hodges (2001, pp. 7–8), Diamond (2001), and Sorensen (2001, p. 89).

⁶ That is assuming that category mistakes are syntactically well-formed. This assumption is contested in Chomsky (1965).

green' might be true. Alternatively, someone who favours conceptual role semantics might assume that category mistakes have no conceptual role and thus, according to the theory, meaningless (I return to this issue in Sect. 6). And finally, even if one is not tempted by claim that category mistakes are in fact meaningless in English or other familiar natural languages, I suspect many are still compelled by a weaker version of the meaninglessness view, one that maintains that there could easily have been a language in which category mistakes are meaningless.⁷

My aim in this paper is to argue that, in spite of its appeal, the meaninglessness view is incorrect and category mistakes are meaningful (call this '*the meaningfulness view*'). I will provide four arguments against the meaninglessness view: in Sect. 2, an argument (or collection of arguments) concerning compositionality with respect to category mistakes; in Sect. 3 an argument concerning synonymy facts of category mistakes; in Sect. 4 concerning embeddings of category mistakes in propositional attitude ascriptions; and in Sect. 5 concerning the uses of category mistakes in metaphors.⁸ Having presented these arguments, in Sect. 6 I briefly discuss some of the positive motivations for accepting the meaninglessness view and argue that they are unconvincing. I conclude that the meaninglessness view ought to be rejected.⁹

A couple of clarificatory notes are in order. First, when I say that category mistakes are meaningful, I mean that they are *literally meaningful*. Consider for example the category mistake 'The theory of relativity eats Newtonian mechanics'. It is clear that speakers can use this sentence and interpret it pragmatically or figuratively, e.g. to communicate the claim that the theory of relativity is a better theory than Newtonian mechanics. Nevertheless, it is clear that the aforementioned content is not the sentence's literal meaning (if it has one). Thus the availability of such pragmatic interpretations is not in itself sufficient to prove that category mistakes are meaningful in the sense that I am concerned with. (How precisely we ought to draw the distinction between literal and non-literal meanings is a notoriously difficult question, and I will not address it here. For current purposes I will simply follow the widely accepted tradition in the philosophy of language that assumes that there is such a distinction to be had).

Second, my arguments will not assume any particular theory of meaning. I have intentionally tried to remain as neutral as possible with respect to the preferred

⁷ An example of someone who explicitly endorses this combination is McDaniel (MS). My arguments will focus on the stronger version of the meaninglessness view, but much of the discussion can serve to undermine the weaker version of the view as well (this is especially so with respect to the arguments in Sect. 2, and depending on the expressive powers of the language in question, the other arguments might carry over as well).

⁸ Many of the traditional defences of the meaningfulness view proceed via the claim that category mistakes are true or false: it is argued that since category mistakes have (bivalent) truth-values they must be meaningful (see for example Lambert 1968 and Haack 1971). But as emerges from the debate concerning this line of defence, it is hard to justify the premise that category mistakes have truth-values without already assuming that they are meaningful (see Routley 1969). I have thus chosen to focus on arguments in favour of the meaningfulness view which do not rely on the claim that they have truth-values.

⁹ Another recent paper which defends the meaningfulness view is Camp (2004). Some of the considerations discussed in this paper (especially in Sects. 5 and 6) are also raised in Camp's paper, but the bulk of the arguments in the two papers are quite different.

theory of meaning, so as to let my conclusion be as general as possible.¹⁰ Of course, this does not imply that I cannot help myself to any claims about meaning. My arguments assume both some intuitive claims about meaning (for example that certain pairs of sentences have the same meaning) and to some more theoretical claims about meaning (for example, certain formulations of the principle of compositionality). The important point is that the claims on which I rely are by and large widely accepted, and moreover are likely to be accepted (at least *prima facie*) even by proponents of the meaninglessness view. Of course, if my arguments are successful, proponents of the meaninglessness view might ultimately choose to reject some of these assumptions rather than abandoning their view. But this sort of move is available in response to *any* valid argument ('one person's *modus ponens* is another's *modus tollens*') and its availability does not show that my argument was all along question begging. Relatedly, it is worth noting that when I say that my arguments will not *assume* any specific theory of meaning, this does not commit me to the claim that my conclusion is *compatible* with any theory of meaning. If some theory of meaning appears to accept the background assumptions I rely on, but nevertheless maintains that category mistakes are meaningless, then that theory suffers from an internal inconsistency that ought to be resolved.

As I hope emerges from the discussion in the paper, the question of whether or not category mistakes are meaningful is not only intrinsically interesting. It also bears on foundational questions such as the limits of compositionality (can we accept not only weak versions of compositionality which claim that the meaning of a sentence is composed out of the meanings of its parts, but also stronger versions which claim that any collection of meaningful parts combined in a syntactically appropriate way compose a meaningful sentence?); the relation between syntax and semantics (is being a syntactically well-formed sentence a sufficient condition for meaningfulness?);¹¹ and the theory of meaning (should we reject certain theories of meaning if it turns out that they deem sentences such as 'Two is green' to be meaningless?). The argument in this paper points us in the direction of an affirmative answer to these questions.

2 The argument(s) from compositionality

2.1 Atomic category mistakes

Perhaps the first argument that springs to mind in favour of the meaningfulness view is the argument from compositionality. Speakers of natural languages have the capacity to understand indefinitely many new sentences. This suggests that meaning

¹⁰ Note that the argument in Sect. 2.2 constitutes no departure from this methodology. I do assume at the outset that type-theoretic semantics is the correct semantic framework. Rather, I argue that such semantics are the best resort for a proponent of the meaninglessness view, but that the view fails even if one takes on board such semantics.

¹¹ For an argument that syntactic well-formedness isn't a necessary condition for meaningfulness see Magidor (2009).

must be compositional: the meaning of a sentence is composed from the meanings of its parts, and speakers grasp the meanings of sentences via grasping the meaning of their parts. Moreover, some particular principles for how such composition works seem to suggest themselves.

Take for example simple subject-predicate sentences of the form 'Fa' (where 'a' is a singular term, 'F' is a predicate, and 'Fa' is the grammatically correct sentence that has 'a' as its subject and 'F' as its predicate). Putting category mistakes to the side for a moment, the following principle seems highly plausible:

(Principle 1) If *S* is a generally competent speaker of a language *L* and *S* understands the terms 'a' and 'F' of *L*, then *S* understands the sentence 'Fa'.

Principle 1 can be taken either as a quasi-empirical claim about which subjects actually understand sentences such as 'Fa', or as a constitutive condition on what it is to understand terms such as 'F' and 'a'.¹² Either way, the principle seems very plausible. Assuming that one is a generally competent speaker of English and that one understands the phrases 'is red' and 'the parrot' then surely one understands the sentence 'The parrot is red'. But if Principle 1 is correct, category mistakes must be meaningful. For when one understands the singular term 'the number two' and the predicate 'is green', then according to Principle 1, one also understands the sentence 'The number two is green'. But understanding a sentence requires grasping its meaning. So if one understands the sentence then it must have a meaning to be grasped, i.e. it must be meaningful.

Principle 1, then, entails very straightforwardly that category mistakes are meaningful. One might worry, though, that the principle begs the question against the meaninglessness view. Proponents of the meaninglessness view might insist that the principle be replaced with a more qualified version such as the following:

(Principle 2) If *S* is a generally competent speaker of a language *L* and *S* understands the terms 'a' and 'F' of *L*, then *S* understands the sentence 'Fa', *if this sentence is a meaningful sentence of L*.

The idea behind Principle 2 is that if 'Fa' can be understood at all, then understanding both 'F' and 'a' is sufficient for grasping the meaning of the sentence 'Fa'. This still leaves open the possibility of claiming that no one can understand sentences such as 'The number two is green' because this sentence is meaningless. In other words, the principle of compositionality is interpreted as entailing only that the meaning of meaningful complex expressions is composed of the meanings of their simpler constituents, not that the meanings of simpler expressions can always be combined to compose complex meanings.

The shift to the weaker principle 2 is thus a reasonable response, but it is not one that comes without a price. The problem is that this response carries with it the

¹² The latter approach is taken by Evans when he introduces the generality constraint (see Evans 1982, pp. 100–105). Interestingly, Evans qualifies the generality constraint so that it does not apply to category mistakes (see footnote 17, p. 101).

burden of finding a semantic framework which on the one hand will explain the mechanisms that enable compositionality when things ‘go right’, and on the other hand will explain what ‘goes wrong’ in the case of category mistakes. In the next sub-section I shall describe what I take to be the most promising way to address this challenge, and argue that this attempt fails.

2.2 Type-theoretic semantics to the rescue?

The proponents of the meaninglessness view are presented with the challenge of finding a semantic framework which both explains how the meaning of ‘harmless’ sentences can be successfully derived in a compositional manner, and explains what ‘goes wrong’ in the case of category mistakes.

It seems to me that there is one popular semantic framework that is best suited for addressing this challenge: functional type-theoretic semantics such as Montague Grammar and its variants. On this framework the semantic values of predicates are taken to be functions from some type of entity to truth-values. Working within this framework, the proponent of the meaninglessness view might make the following proposal.

(Proposal 1): Take the semantic value of a predicate ‘F’ to be function the domain of which includes an object a only if ‘Fa’ is not a category mistake. For example, take the semantic value ‘green’ to be a function from concrete objects to truth-values, of ‘is eating’ to be a function from animate objects to truth-values, and so forth.

Given this proposal, it seems that on the one hand we can explain the mechanism that enables compositionality in the case of ‘harmless’ sentences such as ‘My chair is green’: the mechanism is simply functional application. On the other hand, we can also explain what ‘goes wrong’ in the case of category mistakes such as ‘Two is green’: since the number two is an abstract object not in the domain of the function denoted by ‘green’, functional application fails in this case. The challenge I presented above thus seems to be met.

The problem with this suggestion is that Proposal 1 forces us into an implausible semantic theory. I think it only seems plausible when we consider a very limited fragment of language, but once we try to apply this idea to semantics for larger fragments of language, serious problems arise. To see why, it will be helpful to fix on one type-theoretic semantic theory. I will fix on what is the most prominent theory of this sort, namely Montague Grammar.¹³

In Montague Grammar proper names do not directly denote individuals, but rather they denote the functional analogue of second-order properties, namely functions from first order-properties to truth-values.¹⁴ The semantic value of ‘Two’ is thus taken to be $\lambda X.X(\text{two})$ and the semantic value of ‘prime’ is taken to be

¹³ Montague (1974). For simplicity, I will ignore whenever possible the intensional aspects of Montague Grammar. I will return below to the question of whether moving to other versions of type-theoretic semantics might help my opponent.

¹⁴ This treatment of proper names is motivated by the syntactic similarities between proper names and quantified terms such as ‘Every man’ or ‘A woman’. Since the latter are given semantic values of type $\langle\langle e, t \rangle, t \rangle$, so are the former.

$\lambda x.\text{prime}(x)$, where ‘ x ’ ranges over individuals (i.e. it is a variable of type e) and ‘ X ’ ranges over functions from individuals to truth-values (i.e. it is a variable of type $\langle e, t \rangle$). The meaning of sentences such as ‘Two is prime’ is then yielded by applying the function that ‘Two’ denotes to the function that ‘prime’ denotes so that we get $\lambda X.X(\text{two})(\lambda x.\text{prime}(x))$ which reduces to the desired $\text{prime}(\text{two})$.

Now consider Proposal 1. Since ‘ x is prime’ is a category mistake if and only if x is not a number (or some such similar proposal), then according to Proposal 1 the function that ‘prime’ denotes is ‘ $\lambda x_{\text{num}}.\text{prime}(x_{\text{num}})$ ’, where x_{num} is a variable of a type (call it ‘num’) which contains all and only numbers. Thus ‘prime’ denotes a function of type $\langle \text{num}, t \rangle$, i.e. a function from numbers to truth-values. But since the semantic value of ‘Two’ is a function that only takes as arguments functions of type $\langle e, t \rangle$, i.e. functions whose domain includes all individuals, we now face the absurd situation that functional application breaks down in the case of ‘Two is prime’.

One might propose that in the light of this problem we ought to revise the semantic value of ‘Two’ so as to belong to the type $\langle \langle \text{num}, t \rangle, t \rangle$, i.e. so that its domain consists of functions from numbers to truth-values. This, however, will not do. The reason is that some predicates which are appropriately applicable to ‘Two’ will have to receive a semantic value whose domain includes both numbers and other kinds of objects (examples of such predicates are ‘interesting’, ‘is an object’, and ‘is thought about’). But on the revised proposal we cannot combine such functions with the semantic value of ‘Two’ and thus we reach another absurd conclusion, namely that ‘Two is interesting’ or ‘Two is thought about’ are meaningless.

Several responses to this argument could be suggested by the proponent of the meaningfulness view. First, she might argue, the Montague Grammar view of singular terms as denoting second-order properties is not universally accepted—other versions of type theoretic semantics treat singular terms as denoting entities of type e , i.e. individuals.¹⁵ This point will not help, though, because the above argument can be generalised to other expressions which receive a less controversial treatment. Consider for example the word ‘very’. It is natural to think that ‘very’ receives an interpretation of type $\langle \langle e, t \rangle, \langle e, t \rangle \rangle$, i.e. it denotes a function from first-order properties to first-order properties. For example, the expression ‘very green’ involves applying the function that ‘very’ denotes to the property that ‘green’ denotes which yields a new property—that of being very green. But according to Proposal 1, ‘green’ does not denote a function from individuals to truth-values (i.e. an entity of type $\langle e, t \rangle$) but rather a function from concrete objects to truth-values (call this an entity of type ‘ $\langle c, t \rangle$ ’). It follows that the function ‘green’ denotes is not in the domain of the function ‘very’ denotes which by the lights of the proposed view entails that ‘very green’ is meaningless. One could claim that ‘very’ is ambiguous, receiving one meaning of type $\langle \langle e, t \rangle, \langle e, t \rangle \rangle$ and one of type $\langle \langle c, t \rangle, \langle c, t \rangle \rangle$. But as soon as we consider the fact that Proposal 1 is formulated so as to block all subject-predicate category mistakes it becomes apparent that this ambiguity suggestion

¹⁵ This approach is taken, for example, by Heim and Kratzer (1998).

is unfeasible. For according to the intended generalisation, there will be some first-order predicates that denote functions from individuals to truth-values (e.g. 'is interesting'), some from concrete objects to individuals (e.g. 'is green'), some from humans to individuals (e.g. 'likes dancing'), some from numbers to individuals (e.g. 'is greater than 5'), and so forth. This will entail that 'very' is a *massively* ambiguous word, which seems like a very implausible suggestion. (And of course parallel arguments would show that many other words would have to be massively ambiguous).

A second response might be to claim that words like 'two' have semantic values which belong to a special kind of type, which contains functions with a gerrymandered domain, one that includes functions from individuals to truth-values (such as the function denoted by 'is interesting'), functions from numbers to truth-values (such as the function denoted by 'is prime'), and so forth. On this view, the semantic value of 'two' is the function $\lambda X.X(\text{two})$, where X is a variable which belongs to the special gerrymandered type described above. (Similarly, one might argue that the semantic value of 'very' belongs to a special kind of type containing functions with a gerrymandered domain which includes entities of type $\langle e, t \rangle$, $\langle c, t \rangle$, and various other types). The problem with this proposal is that it constitutes a radical divergence from the type-theoretic semantics under consideration. A large part of the attraction of type-theoretic semantics is its simplicity. One starts with a very small number of basic types (often only two or three), and derives all other types by the sole rule that if a and b are types then $\langle a, b \rangle$ is a type (containing functions from entities of type a to entities of type b). Moreover, this is a tight link between the syntactic category of an expression and the type of its semantic value. This simplicity is completely lost on the suggested response. In particular, note that this suggestion will more or less require that every single word receive a semantic value of a unique type. Consider for example the type of the semantic value of 'two'. The domain of the functions it contains should clearly include functions from prime numbers to truth-values, but also functions from even numbers to truth-values. But then 'two' will be the only (or almost only) word which receives a semantic value of this type, because the number two is the only object which is both prime and even. It therefore turns out that the variable X above will have to belong to a type which only plays a role in the interpretation of the word 'two'. But this would be a big blow to the simplicity and universal character of type theoretic semantics.

A final response, which strikes me as a lot more promising than the previous one, is to retain the original type-theoretic idea that expressions of the same syntactic category receive semantic values of the same type, but slightly modify the definition of types. On the revised definition, if a and b are types, then $[a, b]$ is a type denoting the collection of all *partial* functions from entities of type a to entities of type b . According to the revised proposal, the semantic *values* of predicates are just as in Proposal 1, but their semantic *type* is $[e, t]$. Similarly, the semantic value of 'two' is a function of type $[[e, t], t]$, the semantic value of 'very' is a function of type $[[e, t], [e, t]]$, and so forth.¹⁶ On this suggestion, merely knowing the types of the semantic values of 'two' and of 'green' is not sufficient for determining whether or not an

¹⁶ Thanks to Jason Stanley for this suggestion.

attempted interpretation of 'Two is green' will involve some breakdown in functional application. However, knowing the particular semantic values that 'two' and 'green' in fact receive is sufficient in order to determine that the partial function that 'green' denotes is not in the domain of the partial function that 'two' denotes.

I will not argue here that this proposal is incorrect.¹⁷ Rather, I would like to argue that even if it were correct, the best way to interpret it is as saying that category mistakes such as 'Two is green' are truth-valueless rather than meaningless. Compare this with Heim and Kratzer's discussion of empty definite descriptions.¹⁸ According to Heim and Kratzer, the semantic value of the determiner 'the' is a function of type $\langle\langle e, t \rangle, e \rangle$, although strictly speaking the semantic value of 'the' is only a *partial function* from entities of type $\langle e, t \rangle$ to truth-values.¹⁹ A function f of type $\langle e, t \rangle$ will only be in the domain of the semantic value of 'the' if it yields the value 'true' for one and only one individual. Thus for example, the function denoted by the predicate 'queen of France' is not in the domain of the function denoted by 'the', and thus the expression 'the queen of France' involves a failure of functional application. But clearly, the expression 'the queen of France' is meaningful and suffers at worst from denotation failure (not from meaningfulness failure). Consequently the sentence 'The queen of France is rich' suffers at worst from failure to express a proposition or to possess a truth-value, but is clearly nevertheless meaningful.²⁰

One might argue that this case is different than that of 'The number two is green' because while in the case of the category mistake one only needs to know the meanings of the words in order to find out that there is a failure of functional application, in the case of 'the queen of France' one needs to know further contingent or empirical facts about the world. But as Heim and Kratzer rightly argue, this is a bad criterion for meaningfulness. It is implausible to think that the sentence 'The greatest prime number is odd' is any less meaningful than 'The queen of France is rich'. But one does not need to know any empirical or contingent facts to discover that the function that (according to the Heim and Kratzer's theory) 'greatest prime number' denotes is not in the domain of the function that 'the' denotes. As Heim and Kratzer put it: "In the case of an uninterpretable structure, information about the *type* of each subtree is sufficient to decide that the structure receives no denotation. To detect presupposition failure, by contrast, we must know more about the denotations of certain subtrees than the mere semantic types".²¹

As further support for this claim, consider the sentence 'That is green'. It is plausible to think that as uttered in a context where 'that' refers to the number two, 'that' will receive the same semantic value that 'two' receives, and hence an attempted interpretation of 'That is green' will involve the same breakdown in functional application that 'Two is green' is said to involve. But it would be absurd

¹⁷ Although I think it is incorrect, and argue for that in detail in Magidor (2007, Chap. 4).

¹⁸ Heim and Kratzer (1998, pp. 81–82).

¹⁹ Using my terminology the semantic value of 'the' is of type $[[e, t], e]$ or at least $[\langle e, t \rangle, e]$.

²⁰ The claim that (at least some) sentences containing empty definite descriptions are meaningful is rarely contested. For example, it is the one clear point of agreement between Russell and Strawson in this issue (see Russell 1905, p. 484 and Strawson 1950, p. 321).

²¹ Ibid, p. 82.

to suppose that the possibility of such breakdowns entails that ‘That is green’ is a meaningless sentence. At worst, in the relevant contexts, ‘That is green’ fails to express a proposition or to possess a truth-value.

A similar point holds when we consider sentences such as ‘The thing I am thinking of is green’, as evaluated relative to a possible world in which the thing I am thinking of is the number two. At least according to some theories of definite descriptions (roughly, ones of a Fregean orientation) the definite article ‘the’ as evaluated in a possible world will denote a function from properties to individuals, or rather on its Montague grammar version—a function from properties to second-order properties.²² Thus as evaluated relative to a possible world w^* in which the thing I am thinking of is the number two, ‘The thing I am thinking of’ will denote the function $\lambda X.X_{w^*}(\text{two})$, and the interpretation of ‘The thing I am thinking of is green’ will involve exactly the same breakdown in functional application as is involved with ‘Two is green’. But it would be absurd to conclude from this that ‘The thing I am thinking of is green’ is meaningless.

Again, it does not help to point out that as opposed to the case of ‘Two is green’ the semantic value of ‘The thing I am thinking of is green’ involves no breakdown of functional application as evaluated relative to some worlds. For suppose the thing I am actually thinking about is the number two, then ‘The thing I am actually thinking of is green’ will involve (on the proposed view) the relevant kind of breakdown in functional application as evaluated relative to *every* possible world, but is nevertheless meaningful.²³ Nor will it help to point out that one needs to know some empirical fact to see the alleged failures in functional application. For example, assuming that ‘x is prime’ is a category mistake when ‘x’ denotes a non-natural number, the sentence ‘The ratio between the circumference of a circle and its diameter is prime’ will involve (on the proposed view) the relevant kind of breakdown in functional application as evaluated relative to any world. But although one would not need to know any empirical facts to discover this, the sentence is nonetheless meaningful.²⁴

I conclude that if the final proposal is accepted, category mistakes are perhaps truth-valueless, but not meaningless. As it happens, I also hold the view that category mistakes possess truth-values. But my aim at the moment is only to argue for the more modest claim that they are meaningful. In particular, I have argued that the attempt to use type-theoretic semantics to explain how category mistakes can be meaningless fails. That is, on most versions the attempt leads to absurd conclusions and on its most plausible version it at best entails that category mistakes are truth-valueless, not that they are meaningless. The challenge to the meaninglessness view thus remains unanswered.

²² In general, ‘The’ will denote the function $\lambda w.\lambda Y.\lambda X.X_w(\text{tx}.Y_w(x))$, with Y and X being variable of type $\langle s, \langle e, t \rangle \rangle$, and thus as evaluated in a world w^* ‘The’ will denote the function $\lambda Y.\lambda X.X_{w^*}(\text{tx}.Y_{w^*}(x))$.

²³ Nor will it help to note that ‘The thing I am actually thinking of is green’ will involve no breakdown as uttered in a different context of utterance: if ‘ w_0 ’ denotes the actual world then ‘The thing I am thinking of in w_0 is green’ will involve the relevant kind of breakdown relative to any pair of context of utterance and context of evaluation (fixing the time and the speaker).

²⁴ I am assuming here that the natural numbers are a sub-set of the rational numbers, and thus that at least some ratios can be natural numbers and the sentence ‘The ratio is prime’ is not always a category mistake.

2.3 Conjunctions and quantifier phrases

The argument from compositionality is not restricted to the relations between subjects and predicates. Take for example conjunctions. The following principle seems highly plausible:

(Principle 3) If ‘p’ and ‘q’ are meaningful declarative sentences, then ‘p and q’ is a meaningful sentence.

But now consider the following example:

- (1) That is green and that is prime.

Assuming the two occurrences of ‘that’ are co-referential, (1) seems to be a category mistake. After all, an utterance of (1) where it is clear that the two occurrences of ‘that’ are co-referential is anomalous in just the same way as an utterance of ‘That is a green prime’. So by the phenomenological criterion employed in this paper since the latter is a category mistake, so is the former.

But it is obvious that each of the conjuncts ‘That is green’ and ‘That is prime’ is meaningful. It follows by Principle 3 that the conjunction is meaningful, and hence that category mistakes are meaningful. The only way out for the proponent of the meaninglessness view seems to be to deny the plausible Principle 3.

Similar problems arise when we consider category mistakes containing quantifier phrases. Consider the following example:

- (3) Something is a green prime.

This sentence seems to be a category mistake. The following argument suggests that it is also meaningful:

- (4) ‘green(x)’ and ‘prime(x)’ are both meaningful.

- (5) ‘green(x) ∧ prime(x)’ is meaningful.

- (6) ‘ $\exists x(\text{green}(x) \wedge \text{prime}(x))$ ’ is meaningful.

- (7) ‘Something is a green prime’ is meaningful.

Let us go through this argument, starting with (4): assuming that ‘green(x)’ expresses a genuine constituent in some meaningful sentence (for example ‘ $\exists x(\text{green}(x))$ ’), then since the meaning of a (meaningful) sentence is composed out of the meaning of its constituents, ‘green(x)’ must be meaningful. (5) follows from (4) by Principle 3.²⁵ The move to (6) depends on another compositionality principle:

²⁵ One might wish to argue that this is not so, because Principle 3 requires the two conjuncts to be declarative sentences and open formulas are not declarative sentences. However, even if this is so, it should be noted that the only reason I restricted Principle 3 to declarative sentences was to ensure that ‘p and q’ was well-formed. One can therefore replace Principle 3 with the principle that whenever ‘p’ and ‘q’ are meaningful, and ‘p and q’ is well-formed, ‘p and q’ is meaningful. Since ‘green(x) ∧ prime(x)’ is well-formed, this suffices to show that it is meaningful.

(Principle 4) If $\varphi(x)$ is meaningful then $\exists x(\varphi(x))$ is meaningful.

This seems like a very plausible principle. The proposition expressed by $\exists x(\varphi(x))$ relative to an assignment function f to the variables, is such that it is true if and only if there is an assignment f' to the variables that differs from f perhaps only in its assignment to x , and relative to f' , $\varphi(x)$ expresses a true proposition. This truth-condition is most naturally matched up with the following falsity-condition: the proposition is *false* if and only if *it is not the case* that there is an assignment f' to the variables that differs from f perhaps only in its assignment to x , and relative to f' , $\varphi(x)$ expresses a true proposition.

Now, it may be that in the case in question relative to any such assignment f' , 'green(x) \wedge prime(x)' expresses a false proposition. One might even try to argue that relative to any such assignment f' 'green(x) \wedge prime(x)' expresses no proposition, because any assignment would either assign x a number, in which case 'green(x)' would express no proposition, or assign x something which is not a number, in which case 'prime(x)' would express no proposition. But even so, this would not entail that ' $\exists x(\text{green}(x) \wedge \text{prime}(x))$ ' is meaningless or even that it fails to express a proposition. According to the natural falsity-condition suggested above, if relative to every such f' , *no* proposition is expressed, then it is not the case that *a true proposition* is expressed, so ' $\exists x(\text{green}(x) \wedge \text{prime}(x))$ ' is simply false. Moreover, note that it will not do to insist that the suggested falsity-condition is inadequate, and that the definition should be interpreted so that if relative to some assignment f' the open formula expresses no proposition, then the quantified sentence should be rendered meaningless. The reason is that the same phenomenon occurs in the case of perfectly innocent sentences. The sentence ' $\forall x(\text{green}(x))$ ' has truth-conditions that require us to consider some assignments which make the atomic sub-formula 'green(x)' into a category mistake (e.g. an assignment f such that $f(x)$ is the number two), and on the proposed view this entails that the open formula expresses no proposition relative to such assignments. But should not conclude from this that 'Everything is green' is meaningless.²⁶

Finally, (7) follows from (6) on the assumption that ' $\exists x(\text{green}(x) \wedge \text{prime}(x))$ ' correctly represents the logical form of 'Something is a green prime'. There are, of course, good reasons to think that it does not. But this is not a problem: it is easy to see that the argument could be generalised to other candidates for the logical form of the sentence. For example, suppose the correct logical form is one on which existential quantifier is a second-order predicate applying to a first-order predicate, in something like the following form: ' $\exists x(\lambda x.\text{prime}(x) \wedge \text{green}(x))$ '. Then all we need is to apply to (5) the plausible principles that if $\varphi(x)$ is meaningful then $\lambda x.\varphi(x)$ is meaningful, and that if $\lambda x.\varphi(x)$ is a meaningful first-order predicate, then we can

²⁶ Perhaps one could try and argue that 'Everything is green' only seems meaningful if we implicitly restrict the domain of quantification to concrete objects and that taken over an unrestricted domain this sentence would indeed be meaningless. But first, note that this move will not work if one takes the view of quantifier domain restrictions defended in Stanley (2002), and second, a parallel consideration suggests that on the proposed view 'Something is green' would also be meaningless over an unrestricted domain—which strikes me as much more counterintuitive.

meaningfully apply the existential quantifier to it. Similarly, it is easy to modify the argument to the case where one treats the existential quantifier as a binary second-order predicate so that the logical form of the sentence is taken to be something like ' $\exists x(\lambda x.\text{thing}(x), \lambda x.\text{prime}(x) \wedge \text{green}(x))$ '. Simple and plausible principles of compositionality entail that category mistakes such as (3) must be meaningful.²⁷

The arguments from compositionality are not indefeasible: they rely on compositionality principles such as Principles 3 and 4. Such principles seem extremely plausible, but in light of category mistakes the proponent of the meaninglessness view may wish to reject or at least qualify them. However, given the price of rejecting such principles the arguments from compositionality at least provide us with a strong reason to prefer the meaningfulness view.

3 The argument from synonymy

Two sentences are said to be synonymous if and only if they have the same meaning. This suggests that if two sentences are synonymous then they must be meaningful: if the two sentences have the *same* meaning, then each of them has *a* meaning. Now at least on the face of it, the English sentence 'Two is green' is synonymous with the French sentence 'Deux est vert'. This suggests that the two sentences—both category mistakes—are meaningful.

Of course, a proponent of the meaninglessness view can insist that despite intuitions to the contrary the two sentences are not synonymous, but the denial of this plausible intuition is a price that the meaninglessness view has to pay. Moreover, it is worth noting that it is not entirely obvious how to explain away this synonymy intuition. For example, one might try to argue that two sentences merely seem synonymous because for each word in the English sentence there is a corresponding word in the French sentence to which it is synonymous. This, however, will not do: the Hebrew translation, for example, of 'Two is green' is 'Shtaim yarok' which contains only two words and is not word-to-word synonymous with the English sentence. Still, it is intuitively synonymous with the English sentence (no less so than the French sentence is).

It should be acknowledged that in literary contexts one does encounter a notion of 'adequate translation' which does not assume preservation of meaning. There can be better and worse translations of Jabberwocky in spite of the fact that it contains meaningless words and a literary translation might do a better or worse job of reconstructing in the target language the ungrammatical speech of a linguistically incompetent literary character. Still the case of category mistakes seems very different: it takes no great literary ingenuity to translate 'Two is green' to other languages and it seems perfectly clear which translations are correct and which are incorrect. A proponent of the meaninglessness view might try to come up with a

²⁷ The arguments for compositionality of conjunction and quantifier phrases are particularly important if one adopts a neo-Davidsonian view of verbs. On this view, even atomic category mistake such 'The rock is thinking' will receive an analysis such as the following: ' $\exists e(\text{thinking}(e) \wedge \text{agent}(e) = \text{the rock})$ '. But each of the conjuncts here are clearly meaningful which, following the previous claims, suggests that the whole condition must be meaningful.

sophisticated explanation of this phenomenon that does not entail that category mistakes are meaningful, but the burden of proof is on her, and considerations of simplicity suggest that the best explanation for the apparent synonymy of the English and French category mistakes is simply that they are in fact synonymous.²⁸

4 The argument from propositional attitude ascriptions

Category mistakes can be used in propositional attitude ascriptions.²⁹ Consider the following:

- (8) John said that the theory of relativity is eating breakfast.
- (9) John believes that the number two is green.
- (10) John dreamt that his toothbrush was pregnant.

How exactly does this lead us to the meaningfulness view? Let us consider the following four claims:

(*M*): (8)–(10) are meaningful sentences.

(*T*): For each of (8)–(10), there is some possible circumstance in which it is true.

(*M-entailment*): If *M* is true, then the category mistakes embedded in (8)–(10) (i.e. ‘The theory of relativity is eating breakfast’, etc.) are meaningful.

(*T-entailment*): If *T* is true, then the category mistakes embedded in (8)–(10) are meaningful.

Clearly, either *M* and *M-entailment* or *T* and *T-entailment* are sufficient to establish the meaningfulness view. Since a sentence can only be true if it is meaningful, *T* entails *M* and *M-entailment* entails *T-entailment*. One can therefore either start with *M* which is the weaker of the first two claims (and hence easier to defend), and then establish *M-entailment* which is the stronger of the last two claims (and hence harder to defend). Alternatively, one can start from the stronger claim *T*, leaving us

²⁸ A final move might be to acknowledge that ‘Two is green’ is synonymous with ‘Deux est vert’ but claim that this is so simply because both sentences are meaningless, and thus they both have the same “null” meaning. It is not clear that a proponent of the meaningfulness view would want to go for this move because plausibly they will maintain that ascribing synonymy to meaningfulness sentences results in a category mistake, and hence is itself meaningless (thanks to an anonymous referee here). But at any rate, the move will not do because given the meaningfulness view it would entail that ‘Deux est vert’ is synonymous with every English category mistake, which it is not. For example, it is not synonymous with ‘The theory of relativity is eating breakfast’.

²⁹ The observation that category mistakes can be felicitously embedded in propositional attitude ascriptions is discussed in and McCawley (1970). However, McCawley’s primary target there is Chomsky’s analysis of category mistakes as syntactically ill-formed rather than the meaningfulness view. Moreover, although some things McCawley says certainly suggest he endorses the meaningfulness view, in other places he seems to accept that category mistakes are meaningless (see for example his claim that “an after-image is something that it makes no sense to speak of someone’s owning” in McCawley 1971, p. 294).

with the need to establish only the weaker *T-entailment*. In what follows I will argue that all four claims are highly plausible, leaving either course of argument open.

I shall not say much to defend *M*. It seems to me clear that we can understand (8)–(10), and so they are meaningful sentences. A die-hard proponent of the meaninglessness view might insist on denying this claim, but I am guessing that few others would.³⁰

How about *T*? It is easy to imagine a situation which would make (8) true. John simply utters the sentence ‘The theory of relativity is eating breakfast’. Now, perhaps that is not enough. For John can utter the sounds ‘blablabla’, and yet it is meaningless (and therefore it cannot be true) to say: ‘John said that blablabla’. But there are several reasons to think the case of John’s utterance of a category mistake is quite different from the case of John’s utterance of ‘blablabla’. We may adequately report John’s utterance. Moreover, we may do so using different words from the ones John originally used. For example, even if John actually uttered the French sentence ‘La théorie de la relativité est en train de prendre le petit déjeuner’, we can adequately report him using (8). We can also claim that John said something very odd (note the implication that he said something), ask John to explain why he believes what he said, and so forth.

How about (9)? One might think that this sentence cannot be true for the following reason: if one does not possess the concepts of ‘two’ or of ‘green’ then one cannot believe that the number two is green. On the other hand, the reasoning goes, anyone who does possess the concepts of ‘two’ and of ‘green’ cannot believe that the number two is green, for it is a necessary condition on possessing these concepts that one would not believe the obviously absurd claim that the number two is green. So either way, the argument goes, it is impossible for anyone to believe that the number two is green.

The problem with this argument is with the claim that it is a necessary condition on possessing the concepts of ‘two’ and ‘green’ that one does not believe that two is green. I think this claim is false. Consider the following scenario: John is a philosopher. He recently developed a new theory in the philosophy of mathematics according to which numbers are coloured, and the colour of the number two is green. For example, John may hold some naturalist position according to which the number two is the set of all pairs of physical objects in the world. In addition John might hold that if most such pairs have a certain colour, then the set—and therefore the corresponding number—have this colour. (Compare this to a more popular philosophical position suggested by Penelope Maddy: a set of physical object has a spatial location and it is located wherever its members are located).³¹ Thus, if it happens to be the case that most pairs of physical objects are green then, according to John’s theory, this makes the number two green. Suppose that following some empirical investigation John concludes that it is in fact the case that most pairs of objects are green and so, following his theory, John comes to believe that the

³⁰ Note that proponents of the meaninglessness view might at least be tempted to accept sentences such as ‘It is nonsense to say that the number two is green’—such sentences are good enough for my purpose in this section.

³¹ See Maddy (1980).

number two is green (or, to not yet beg the question, at least comes to assent to the sentence ‘The number two is green’). But John, we may suppose, also knows a lot of mathematics, and clearly possesses the concept ‘two’. Also, we may suppose, John generally does a perfectly good job of telling which things are green and which are not. It seems odd to say that John does not possess the concept ‘green’. I am not denying that usually if someone seriously utters the sentence ‘The number two is green’ it would be sensible to conclude that they do not understand the meaning of either ‘two’ or ‘green’. (That is, either they do not know the meaning of the English words, or they do not even possess the relevant concepts). Usually, this conclusion will be correct. But it will not always be correct. In particular, we should not conclude that in our example John does not possess the concepts of ‘two’ and ‘green’, and hence we have no reason to reject the claim that John believes that the number two is green.^{32,33}

As to (10), it’s certainly an empirical fact that people often sincerely *report* having had dreams which involve category mistakes. Whether or not we can infer from this that such reports are in fact correct is a complex question which I cannot address in the scope of this paper. I will leave matters by saying that there is at least a putative reason to accept there can be situations where sentences such as (10) are true.

So far, I have defended *M* and *T*. I next wish to argue that under any standard theory of propositional attitude ascriptions *T-entailment* holds, and with some additional assumptions so does *M-entailment*. There are generally three types of views regarding the semantics of propositional attitudes: the *propositional view* according to which propositional attitudes are relations between agents and propositions, the *Fregean view*, according to which propositional attitudes are relations between individuals and meanings, and the *sentential view* according to which propositional attitudes are relations between individuals and sentences or utterances.³⁴

How does each of these views fare with respect to *T-entailment* and *M-entailment*? First, consider the propositional view. Whatever we take propositions to be,

³² One might try to claim that although John possesses ordinary concepts of ‘green’ and ‘two’, when philosophising he uses ‘green’ and ‘two’ with a slightly different meaning from the ordinary. I think this worry can be at least somewhat alleviated by stipulating that an explicit part of John’s theory is the claim that ‘two’ and ‘green’ are used in the theory with exactly the same meaning as when used in everyday contexts.

³³ One could object that it is easier to construct such an argument for the particular category mistake I have chosen (namely, ‘the number two is green’) than for others (e.g. ‘the theory of relativity is eating breakfast’). I believe that with enough ingenuity similar situations can be constructed for almost any category mistake. To briefly point out two further examples of cases where someone might sincerely believe a category mistake: First, some people who suffer from synesthesia perceive certain sounds as always accompanied by a perception of a certain colour. It would probably be quite natural for such persons to believe that ‘This sound is green’ is literally true. Second, Lakoff (1971, p. 332), reports that a common belief among Papagos is that events have minds, so it might well be natural for such persons to believe that ‘My birthday is angry’ is literally true. At any rate, the claim that at least some paradigmatic category mistakes are true will show (given *T-entailment*) that at least some paradigmatic category mistakes are meaningful, and this should suffice to seriously undermine the meaningfulness view.

³⁴ For typical defences of these three views see Barwise and Perry (1990), Frege (1952) and Davidson (2001a), respectively. (I am assuming here that we can interpret Fregean senses as meanings.)

the view entails that if ‘ ϕ ’ denotes a propositional attitude and if ‘S ϕ s that p’ is true, the phrase ‘that p’ successfully denotes a proposition.

Now assume that *T* holds, i.e. that there are possible situations in which the propositional attitude ascriptions are true. Relative to such situations, the category mistakes embedded in the ‘that’-clauses must successfully express a proposition. But meaningless sentences cannot express propositions. It follows that the embedded category mistakes are meaningful. So *T-entailment* holds. What if we reject *T* and adopt only the weaker claim *M*, i.e. we claim that while ascriptions such as (8)–(10) are never true they are nonetheless meaningful? This view is possible under one of the following versions:

(Position 1) The category mistakes embedded in (8)–(10) are meaningful and express a proposition, but the propositional attitude ascriptions containing them are always false.

(Position 2) The embedded category mistakes are meaningful, but nonetheless they do not express a proposition, and consequently, the propositional attitude ascriptions containing them are either truth-valueless or false.

(Position 3) The embedded category mistakes are meaningless, do not express a proposition, but the propositional attitude ascriptions are nonetheless meaningful.

Both Positions 1 and 2 result in the desired conclusion that the relevant category mistakes are meaningful. Position 3 does not result in this conclusion. But Position 3 can be rejected if we add two plausible assumptions: that the meaning of a sentence is composed of the meanings of its constituents, and that sentences (8)–(10) have the embedded category mistakes as constituents. If this is granted, then the embedded category mistakes cannot be meaningless while (8)–(10) are meaningful. So the stronger claim *M-entailment* also holds.

Next, consider the Fregean view. According to this view, if ‘S ϕ s that p’ is true then S stands in the relation ϕ to the meaning of ‘p’. But for this to be the case ‘p’ must have a meaning, i.e. it must be meaningful. So, if it is possible for (8)–(10) to be true, the embedded category mistakes must be meaningful, and so *T-entailment* holds.

What if we assume only the weaker claim *M*? We have the following possible positions:

(Position 4) The embedded category mistakes are meaningful, but the propositional attitude ascriptions containing them are false.

(Position 5) The embedded category mistakes are meaningless, but the propositional attitude ascriptions containing them are nonetheless meaningful.

If we adopt Position 4 the desired conclusion follows. Ideally, we would like to rule out Position 5 by appealing to the compositionality of meaning (as we did for Position 3). However, this case is a little trickier because according to the Fregean view we are currently considering the embedded category mistakes do not contribute to the meaning of the propositional attitude ascriptions their regular meaning, but rather a *second-order* meaning. So in order to argue from the claim that

the propositional attitude ascriptions are meaningful to the desired conclusion that the embedded category-mistakes have a *first-order* meaning, we must also adopt the following principle: an expression can only have a second-order meaning if it has a first-order meaning. But it seems that from within the Fregean framework considered this principle should be accepted. Plausibly, the theory should view second-order senses as somehow computable from the first-order senses: otherwise speakers would need to learn two primitive senses for each embedded sentence, and this point is strengthened if one accepts that to accommodate sentences with multiple embeddings each sentence will need a separate *n*th-order sense for any number *n*.³⁵ So, Position 5 is likely to be ruled out by similar considerations that rule out Position 3 above, and *M-entailment* holds.

Finally, consider the sentential view. Admittedly, this view is the best avenue for the proponent of the meaninglessness view: after all, on the sentential view the semantic function of the embedded sentences is to refer to themselves and the claim that the embedded category mistakes successfully denote sentences (or utterances) does not seem in any way at odds with the claim that these sentence are meaningless.

Nevertheless, I think that at least *T-entailment* is plausible even on the sentential view. However the details of the view are filled, it must somehow account for the fact that we can truly ascribe a propositional attitude to an individual even if that individual uses different words to express their own attitude. For example, we can correctly claim that Galileo said that the earth moves, even if Galileo never uttered the English sentence ‘The earth moves’. Consequently, such views of propositional attitudes must assume some kind of relation between the utterances of the ascriber and the utterances or thoughts of the ascribee: a relation such as synonymy or having the same content or, following Davidson, of ‘being such as to make the ascriber and ascribee same-sayers’. But suppose that my utterance of ‘The theory of relativity is eating breakfast’ in (8) merely has the same meaning as or the same content as (rather than being an exact repetition of) John’s original utterance. If my utterance and John’s have the *same* meaning or the same content, it suggests that both utterances have a meaning or content—i.e. they are meaningful.³⁶ As with the argument from synonymy, the proponent of the meaninglessness view can insist that the relation that needs to hold between John’s original utterance and my report is not that of having the same meaning or content, but rather some other relation that does not require our utterances to be meaningful. But as above, this places an explanatory burden on the proponent of the meaninglessness view, and one which is far from trivial to meet.

³⁵ This claim is defended Dummett (1981) and contested in Parsons (1981). For a more recent defence which addresses Parson’s attack see Boisvert and Lubbers (2003).

³⁶ Things are somewhat trickier when we consider Davidson’s version of the view because Davidson wants to avoid reference to intentional entities such as propositions or meanings, and thus does not assume that there is literally a meaning or content that John and my utterances share. Nevertheless, Davidson does insist that the relation same-saying requires synonymy between the utterances of the reporter and the reportee (see e.g. Davidson 2001a, p. 104 or p. 107), and this suggests that he too would accept that the two utterance must be meaningful (cf. Sect. 3 of this paper). He also claims that the utterance of the reporter must ‘serve at least the purpose of conveying the content of what someone said’ (p. 107)—again suggesting that the utterance of the reporter must be contentful.

It should be noted, though, that the argument just presented does depend on the stronger claim *T*, rather than on *M*: if one can never truly report the contents of what John said, then we are under no obligation to claim that his utterance was in some way meaningful. Adopting the sentential view, then, we can plausibly conclude *T-entailment*, though probably not *M-entailment*.

5 The argument from metaphor

Many, if not most, metaphors involve category mistakes. Consider for example ‘The silence was liquid’, ‘A sentence wears its meaning on its sleeve’ and ‘This poem is pregnant’.³⁷ The fact that most metaphors involve category mistakes is not a coincidence. It seems that a big part of what makes metaphors the poetic or figurative devices that they are has to do with connecting objects and properties that normally seem to belong to completely disjoint domains.

Metaphors clearly have some communicative purpose: they are intended to communicate some content, or at least produce some effect in the hearer. What can we infer from this regarding the question of whether category mistakes have *literal* meanings? In this section, I argue as follows. I will begin by defending the claim that for metaphors in general to achieve their *metaphorical* communicative purpose, they must have *literal* meanings. This claim will entail that, since many metaphors involving category mistakes manage to achieve their metaphorical purpose, they must also have literal meanings so category mistakes must be (literally) meaningful.

How does one defend the claim that for metaphors to achieve their metaphorical purpose they must have literal meanings? I will discuss the most prominent linguistic theories of metaphor, and classify these theories into two categories: those that require metaphors to have literal meanings in order to achieve their metaphorical purpose and those that do not. I will argue that theories that fall under the second category ought to be rejected, and hence that one ought to accept the conclusion that metaphors have literal meanings. To clarify, I will only be concerned with *linguistic* theories of metaphor rather than *cognitive* or psychological theories of metaphor. That is to say, I will be concerned with theories about the nature of metaphorical truth and metaphorical meaning, rather than theories concerning how we process those metaphorical meanings. I maintain that the linguistic and philosophical questions concerning metaphors are by and large independent of the cognitive issues: it may well be right that (at least some) metaphorical meanings are processed directly and not via their literal meanings. This is nevertheless compatible with the claim that metaphors can only have a metaphorical meaning or effect if they have literal meanings. (By analogy, suppose it turns out that we can process the meanings of some familiar sentences directly, without processing the meaning of each of the words in the sentence first. This is still compatible with the claim that the meaning of the sentence is linguistically composed out of the meaning of its parts, and thus that the sentence could not be meaningful unless the words appearing in it

³⁷ A more complicated issue which I shall not address is whether sentences involving explicit type attributions such as ‘John is an ice-cube’ should ultimately count as category mistakes.

were meaningful). It is for this reason that I will not discuss the rich literature on metaphor from psychology and cognitive science.³⁸

Reimer and Camp's recent prominent survey of the field presents the following four classes of theories as representative of the current state of the debate:³⁹ simile theories (and related to them the substitution and expansion theories); interaction theories; Gricean theories; and non-cognitivist theories.

I shall not say much about interaction theory: I find the view too vague to allow for a serious discussion of it⁴⁰ and moreover it seems that any way of making the view more precise has it collapse either into a version of non-cognitivism (if the interaction effects are not taken to produce a specific metaphorical 'content', in the standard sense of the term) or into a version of the substitution view (if the interaction effect is taken to produce a specific content, albeit one that cannot be paraphrased in literal terms). My discussion of these other two kinds of theories should thus apply to interaction theories as well.

Let me turn to simile theories. According to the standard version of the simile theory, a metaphor means that same as its corresponding simile.⁴¹ For example, the metaphorical meaning of 'Juliet is the sun' is the same as the literal meaning of the simile 'Juliet is like that sun'. The simile theory thus understood falls under the category of theories that allow for metaphors to have metaphorical meanings without the metaphorical sentence being meaningful: there is no obvious contradiction in claiming that 'The poem is pregnant' is (literally) meaningless, while the simile 'The poem is like someone who is pregnant' is meaningful.

However, the standard simile view of metaphor ought to be rejected. First, as Davidson notes, it makes metaphors 'too easy' to figure out: It does not explain the general feeling that metaphors cannot be paraphrased in literal terms.⁴² Relatedly, the theory does not explain why we use metaphors as well as similes. But most importantly, even if the theory deals adequately with simple metaphors such as 'Juliet is the sun', it is not clear how it should handle metaphors such as 'John rides his mind at a gallop in search of an idea'.⁴³ For if this metaphor is simply elliptical for 'John is like one who rides his mind at a gallop in search of an idea' the corresponding simile still leaves us with a metaphor.

In response one might opt for a much more sophisticated simile theory—one that does not translate any metaphor of the form 'A is B' into the simile 'A is like B', but into some other, more complex, simile. For example, the above sentence might correspond to a less obvious simile such as 'John is like someone who searches very hard for an idea'. But it is hard to see how such a proposal will get off the ground

³⁸ See Camp (2006) for further defence of the claim that the question of which philosophical linguistic theories of metaphor ought to be adopted is in principle independent of the question of cognitive processing.

³⁹ Reimer and Camp (2006). As Reimer and Camp put it, their survey "is intended to be representative rather than exhaustive".

⁴⁰ At least when it comes to Black's defence of the theory, Reimer and Camp concur with this judgment (*ibid.*, p. 855).

⁴¹ See discussion in Black (1962), Davidson (2001b), and Fogelin (1988).

⁴² Davidson (2001b, p. 254).

⁴³ Taken from Davidson (2001b, p. 253), who quotes Virginia Woolf.

taking into account the meaning of not only words but complete phrases, and ultimately of the metaphorical sentence as a whole (more on this below). So it seems that if the view of metaphor as simile has any plausibility, one will have to resort to a more complex formulation of the theory, a formulation that falls under the first category of theories (namely those that maintain that metaphorical meanings require the metaphorical sentence to be literally meaningful) after all.

Other (somewhat related) kinds of theories that arguably fall under the second category are the 'expansion of meaning' and 'substitution' views of metaphor.⁴⁴ According to the 'substitution' view, a word used metaphorically is merely a substitute for another word or phrase that expresses the same meaning literally. Thus for example when 'ice-cube' is used metaphorically (e.g. in the sentence 'John is an ice-cube'), it acts as a substitute for the literal phrase 'person who is cruel and unemotional'. I do not wish to restrict the discussion, though, to cases where there actually exists a literal substitute word that expresses the metaphorical meaning. I will thus understand the substitution view as making the wider claim that a particular word used metaphorically has a second, metaphorical meaning in addition to its literal meaning (whether or not that second meaning can be expressed literally). The 'expansion of meaning' view is a version of the substitution view (understood in the wider sense), according to which the metaphorical meaning is an expansion of the literal meaning. Thus on the expansion view, when used metaphorically the word 'ice-cube' will apply both to cubically shaped frozen pieces of water as well as to people who are cruel and unemotional.

At least at a first pass, both these theories fall under the second category: there is no obvious contradiction in the claim that the literal meaning of 'ice-cube' is not meaningfully applicable to the person John (and hence that 'John is an ice-cube' is meaningless), while the metaphorical meaning of 'ice-cube' is meaningfully applicable to John.

However, thus understood these theories too ought to be rejected. First, it is not clear why, under these views, the metaphorical meaning is not simply a second literal meaning.⁴⁵ If the sentence 'John is an ice-cube' is analysed so that 'John' refers to John, and 'ice-cube' refers to some property that literally applies to John, then at best 'ice-cube' can be seen as ambiguous, and 'John is an ice-cube' as literally true, on one disambiguation. Second, the view faces a similar problem to the one mentioned above concerning more complex metaphors: the metaphor 'John rides his mind at a gallop' does not seem to involve any words being individually used in a special metaphorical sense. Rather, it seems to be the whole complex verb-phrase which receives a metaphorical interpretation here. As above, this suggests that one ought to consider more sophisticated versions of the views, versions that allow the metaphorical meaning to be initially assigned to expressions that are more complex than individual words. But again, this pulls the views towards the first category of theories, namely those on which metaphorical meaning requires the sentence as a whole to be meaningful. Note in particular, that although in the

⁴⁴ See Black (1962), Moran (1997, p. 253) and Davidson (2001b, p. 248) for discussion of these suggestions.

⁴⁵ This point is made by Davidson (2001b, pp. 248–249) and Searle (1979, p. 413).

example we are considering it is the verb-phrase (rather than the whole sentence) that is the minimal unit requiring a metaphorical interpretation, it is precisely the verb-phrase which is responsible for making the sentence into a category mistake in this case, and hence it is the verb-phrase which according to the proponents of the meaninglessness view ought to already be meaningless. If one is willing to assign a literal meaning to the categorically mistaken phrase ‘rides his mind at a gallop’, there should be no further obstacle to taking the whole sentence to be meaningful.

Another problem with the substitution and expansion theories is that the same predicate can have very different metaphorical contributions in different contexts. Consider the sentence ‘Juliet is the sun’ as uttered by Romeo and the sentence ‘Stalin is the sun’ as uttered by a devoted communist. It is not one and the same property that we are ascribing to Juliet and Stalin: in the sense in which Juliet is the sun Stalin is not, and vice versa. Of course, one could claim that ‘sun’ has many metaphorical meanings, and that we are appealing to different meanings in different contexts. But this would have the unappealing consequence that every predicate is massively ambiguous. Perhaps a more adequate way to handle this point would be to argue that expressions have only one metaphorical meaning, but that this meaning is context sensitive, i.e. it can receive different contents in different contexts. The most sophisticated defence of this idea appears in Stern (2000). In a brief, on Stern’s view there is an operator ‘Mthat’ which can take any simple or complex expression ‘ ϕ ’. ‘Mthat(ϕ)’ is a context sensitive expression, such that given a context c , its content is fixed to be those semantic values that are presupposed in c to be ‘m-associated’ with ‘ ϕ ’. Finally, a sentence is interpreted metaphorically whenever it is interpreted as having at least one Mthat-operator in its scope.

I cannot discuss Stern’s view here in the detail that it deserves, but let me make a few brief remarks. First, on the face of it Stern’s view has the odd consequence that metaphorical meanings need not depend even on the literal meanings of the individual words in a metaphorical expression. After all m-association is a relation between agents and *expressions* (rather than agents and semantic-values or meanings), and since Stern claims there can be many different grounds for m-association it is not clear that all these grounds have to involve the meaning of ‘ ϕ ’ rather than the expression itself. But this consequence is odd because it fails to explain why we expect metaphors to be composed of grammatical sentences built out of meaningful words (why can’t one have m-associations to completely meaningless expressions?). Perhaps Stern can insist that as a matter of fact, the grounds for m-association happen to always involve the meaning of ‘ ϕ ’ and that consequently, m-association only holds (and is presupposed to only hold) over meaningful expressions. But this brings us back to the point of complex metaphors such as ‘John rides his mind at a gallop’. Stern’s view is able to accommodate such metaphors because it allows that ‘Mthat’ can operate on complex expressions as well as on words. However, the same line of thought discussed above applies here: if Mthat operates on the complete verb-phrase and m-association requires meaningfulness, then the complete verb-phrase ought to be meaningful. But once the problematic verb-phrase containing the category mistake is meaningful, there is no further reason not to accept that complete sentence is meaningful.

It is also worth noting that Stern's view faces various difficult problems. For a start, the view does not address the question of why metaphorical meaning isn't just a case of a second literal meaning. Moreover, although Stern's move from many metaphorical meanings to one single highly context sensitive meaning addresses the objection concerning massive ambiguity, it comes at the price of creating new problems. To point out two (both raised in Camp 2005): because Stern allows embedding any expression in a sentence under an *M*that-operator, Stern's view avoids massive lexical ambiguity at the price of introducing massive structural ambiguity.⁴⁶ Also, the move to context sensitive terms causes complications in propositional attitude ascriptions. If 'p' is a context sensitive term, we expect its content to be fixed by the context of utterance, even if it is uttered as part of an embedding of a propositional attitude ascription such as 'John said that p'. This suggests that on Stern's view there is no reading according to which the ascription 'Romeo said that Juliet is the sun' ascribes to Romeo whatever property *he* (rather than the reporter) presupposes to be *m*-associated with 'the sun'. But this seems wrong.⁴⁷ The upshot of this brief discussion is first, that Stern's view faces some serious challenges, but more importantly that the treatment of complex metaphors on his theory seems to ultimately require at least some metaphorical sentences—including ones which are themselves category mistakes—to be meaningful as a whole.

Next, we have Gricean theories of metaphor, which take metaphorical meanings as a species of the general phenomenon of conversational implicatures.⁴⁸ Different Gricean theories differ in the details of which implicatures generate the relevant metaphorical meanings, but that details need not concern us. What is important is that Gricean theory in general assumes that conversational implicatures are generated via literal contents and hence that a sentence cannot generate an implicature without being literally meaningful. The upshot is that Gricean theories fall under the first category, namely theories that imply that if a sentence has a metaphorical meaning it must be literally meaningful.

Finally, we have non-cognitivist theories of metaphor, most prominently defended by Davidson.⁴⁹ According to non-cognitivists there is no such thing as metaphorical meaning. Of course metaphors can produce (and be intended to produce) various effects in the hearer, but those effects are secondary to the literal meaning of the metaphor.⁵⁰ The upshot is again that the view falls under the first

⁴⁶ Camp (2005, p. 717).

⁴⁷ See Camp (2005, Sect. 3). There is some scope for addressing this problem via more complex theories on how context interacts with propositional attitude ascriptions. I shall not explore this possibility here.

⁴⁸ See e.g., Martinich (1984) and Searle (1979).

⁴⁹ Davidson (2001b).

⁵⁰ It is important to note that Davidson explicitly endorses the claim that metaphors have literal meanings, and that it is via these meanings that one achieves the metaphorical effect. For example, he says that the metaphorical effect "is something brought off by the imaginative employment of words and sentences and depends entirely on the ordinary meanings of those words and hence on the ordinary meanings of the sentences they comprise" (ibid, p. 247).

category: it entails that all metaphors, and in particular metaphors involving category mistakes, have a literal meaning.⁵¹

To summarise: I have discussed what are taken to be the most prominent linguistic theories of metaphor. I argued that many of these theories require a metaphorical sentence to be literally meaningful in order to achieve its metaphorical purpose. I have also presented several theories that do not require metaphors to be literally meaningful (for example versions of the simile and substitution views), but I argued that such approaches should be rejected. I thus conclude that metaphorical sentences, and in particular metaphorical sentences involving category mistakes, must be literally meaningful.

While I have discussed what are taken to be the prominent theories of metaphor, I was obviously not able to discuss *all* theories of metaphor. But it is worth noting that the considerations which were raised against theories that do not take metaphors to have literal meanings point us to some general challenges that would face any view which allows for metaphorical meanings without literal meanings: metaphorical meanings would have to take into account not only the literal meanings of the words in the metaphor, but also complex interactions between these words—meanings, and the way in which they combine into increasingly complex phrases. Yet on the proposed views these complex interactions are supposed to fall short of assigning a meaning to the metaphorical sentence as a whole. It is doubtful that any theory can succeed in walking the fine line between these two constraints, and thus these challenges leave us, I think, with good reasons to be sceptical about the prospects of theories of metaphor that allow for metaphors literally meaningless.

6 Arguments in favour of the meaninglessness view?

So far I have presented four arguments against the meaninglessness view. But it may be instructive at this point to return to the positive motivations in favour of the view. Proponents of the meaninglessness view rarely discuss these motivations explicitly. Rather they tend to take it for granted that in the absence of any compelling counter-arguments the default or natural view to take should be the meaninglessness view.⁵² However, as far as I can see there are four main ideas that can be seen to motivate the meaninglessness view. The most serious one is that category mistakes are highly anomalous and the meaninglessness view seems to provide a

⁵¹ Camp presents an argument that resembles my argument from metaphors (Camp 2004, pp. 223–226). However, she chooses to attack Davidson's view as being in general an incorrect account of metaphors (ibid, pp. 225–226). Whether or not her attack is successful, this move strikes me as dialectically unnecessary in this context because Davidson's view anyhow takes metaphors to be literally meaningful. One might try to develop an alternative non-cognitivist view according to which the metaphorical effects are achieved directly via the literal meanings of the individual words in the metaphorical sentence, rather than through the literal meaning of the sentence as a whole. This will not do since the grammatical structure of the sentence clearly matters to the meaning of the metaphor: 'This man is a stone' is not the same metaphor as 'This stone is a man'. Moreover, is it not merely superficial aspects of the grammatical structure that play a role here, as is apparent from translating metaphors to languages with quite different syntactic structures than English.

⁵² See for example Diamond (2001, p. 96) or Pap (1960, p. 41).

simple and compelling explanation for their anomaly. I think this is indeed a strong motivation in favour of the meaninglessness view and I take it that the best way to defuse this motivation is to provide an alternative explanation for the anomaly of category mistakes—one which is consistent with the claim that category mistakes are meaningful. The task of providing such an explanation, however, is beyond the scope of this paper and I reserve it for other work.⁵³ In the remainder of this section I will briefly discuss other motivations for the view and explain why I do not find them very compelling.

6.1 The imagination motivation

One motivation which sometimes mentioned for taking category mistakes to be meaningless is that one cannot even imagine ‘what it would take for ‘Two is green’ to be true’. Underlying this complaint is perhaps the thought that meaning is truth-conditions and that a sentence cannot have truth-conditions if one cannot imagine ‘what it would take for it to be true’.

But the best sense I can make behind this somewhat obscure complaint is that it is merely a convoluted way of saying that one cannot imagine a situation in which ‘Two is green’ is true.⁵⁴ And if so, the complaint is a rather puzzling one: since it necessarily false that two is green it is hardly surprising that we cannot imagine a situation in which ‘Two is green’ is true. Conceivability may not be an infallible guide to possibility, but it seems odd to conclude from the fact that in this case conceivability is a *successful* guide to possibility, that the sentence in question is meaningless.

Nonetheless, the objector might insist that for all other examples of meaningful but necessarily false sentences one *can* imagine situation in which the sentence in question is true. For example, even though ‘There is a counterexample to Fermat’s last theorem’ is a necessarily false sentence, one can perhaps imagine a situation in which it is true (for example, I imagine the newspaper headlines announcing ‘Famous mathematician finds mistake in proof for Fermat’s last theorem and produces a counterexample!’). Similarly, even though the sentence ‘The number 5 is even’ is necessarily false, I can perhaps imagine a situation in which it is true (for example, I imagine someone dividing five by two and being left with no remainder). But, the argument goes, the same is not true of category mistakes: as hard as one tries one cannot imagine a situation in which ‘Two is green’ is true.

I am not sure that if the suggested disanalogy in our imagination powers were true that would have proved much regarding the meaningfulness of category mistakes. But more straightforwardly, I don’t think there is any such disanalogy. If ‘imagining a situation in which *s* is true’ is interpreted in a sense which is permissive enough to include imagining dividing five by two and being left with no remainder, then it seems to me obvious that we can also imagine that ‘Two is green’

⁵³ I develop and defend a pragmatic account of category mistakes, one that is consistent with the claim that they are meaningful and truth-valued in Magidor (2007), and in a forthcoming monograph.

⁵⁴ I assume that situations in which ‘Two is green’ simply has a different meaning from the one it actually has are discounted.

is true (imagine that you check what colour the number two has, and it turns out to be green). Moreover, I take it that cases such as that of the philosopher of mathematics described in Sect. 4 provide us with a way of imagining that ‘Two is green’ is true which does not require very wild imagination. The imagination motivation is thus unconvincing.

6.2 The motivation from alternative theories of meaning

The imagination motivation was most naturally linked to the idea of meaning as truth-conditions. A second motivation for the meaninglessness view stems from a preference for some alternative theory of meaning coupled with the thought that one’s favoured theory of meaning is incompatible with the meaningfulness view.

Consider for example verificationism about meaning. According to verificationism, the meaning of a sentence is its verification conditions, and a sentence is meaningful if and only if it is either verifiable or falsifiable. According to at least some version of verificationism (roughly, the version endorsed by the logical positivists—let us call this ‘traditional verificationism’) only two kinds of verification methods count as legitimate: judging that a sentence is true or false on the basis of direct sense experience and judging that a sentence is true or false on the basis of its being an analytic or a logical truth or falsehood.

One might argue that traditional verificationism entails that category mistakes are meaningless. No sense experience, the argument goes, shows that ‘Two is green’ is either true or false. But neither is ‘Two is green’ analytically true or analytically false. So ‘Two is green’ fails to have legitimate verification conditions and hence by the lights of traditional verificationism it is meaningless.⁵⁵

Or consider conceptual role semantics. According to conceptual role semantics “the meaning of a representation is the role of that representation in the cognitive life of the agent, e.g. in perception, thought and decision-making”.⁵⁶ One might argue that conceptual role semantics entails that category mistakes are meaningless. Sentences such as ‘Two is green’, the argument goes, play no role in the cognitive life of any agent and hence by the lights of conceptual role semantics they are meaningless.

It is hard to address such worries properly without entering into a lengthy discussion of the subtleties of various theories of meaning. But I would like to offer two brief responses. First, it remains to be shown that any suggested theories of meanings are in fact inconsistent with the meaningfulness view. The two arguments I have sketched above are certainly questionable. It is not clear that general verificationism about meaning entails the meaningless view, for one might argue that we can easily verify that ‘Two is green’ is false.⁵⁷ Nor is it even clear that traditional verificationism entails the meaninglessness view. If one concedes that sense experience of non-black non-ravens can verify the claim that ravens are black, then one

⁵⁵ See Reimer and Camp (2006, p. 847) for the suggestion that verificationism was a motivation for the meaninglessness view.

⁵⁶ Block (1998).

⁵⁷ Compare this with the example of the continuum hypothesis, for which there is much stronger case to be made for the claim that it is neither verifiable nor falsifiable.

might also argue that sense experience of green non-numbers can verify the claim that numbers, and in particular the number two, are not green. Alternatively, one might try to argue that ‘Two is green’ is analytically false in much the same way that ‘ $2 + 2 = 5$ ’ or ‘Something is green all over and blue all over’ can be argued to be analytically false. Similarly, it is not clear that conceptual role semantics really entails the meaninglessness view. One might argue, for example, that the uses of category mistakes in metaphors entail that they *do* play a role in the cognitive lives of agents. It is perhaps on the basis of her belief that (metaphorically) a poem is pregnant, that a publisher decides to publish it.

My second response to the ‘theories of meaning’ argument is this. If it nonetheless turns out that a certain theory of meaning X is inconsistent with the meaningfulness view, then all the worst for that theory of meaning. The arguments I have provided in this paper in favour of the meaningfulness view did not rely on a particular theory of meaning. Rather they relied on recognising certain linguistic phenomena such as metaphorical uses of category mistakes or embeddings of category mistakes in propositional attitude ascriptions, as well as on some general principles of language such as the principle of compositionality—principles which I assume that any reasonable theory of meaning should accommodate. It is insufficient for proponents of X to object to my arguments on the grounds that my conclusion is inconsistent with their favoured theory X. One must either show that my arguments are unsound, or else abandon X in the light of my conclusion.

6.3 The nonsense motivation

Perhaps a final somewhat flatfooted motivation for the meaninglessness view is simply the intuition that sentences such as ‘The number two is green’ or ‘The theory of relativity is eating breakfast’ should be classified as ‘nonsense’.

But even if this is so, I think one ought to be careful not to confuse the philosophical claim that something literally has no sense, with the natural English use of the phrase ‘nonsense’. My friend can say to me ‘My work is worthless!’, and I can reply ‘That’s nonsense, and you know it!’. Of course, I do not mean to claim that the sentence she uttered was literally meaningless. What I mean to say is that what he said was obviously false, or a ridiculous claim to make. In this sense, it seems perfectly appropriate to categorise category mistakes as nonsense: at least in most contexts, they are rather ridiculous sentences to utter. But this does not entail that they literally have no sense or are meaningless. On the contrary, as I hope to have established in this paper, category mistakes are perfectly meaningful sentences, and the meaninglessness view cannot be the correct explanation for their anomaly.

Acknowledgments For helpful comments and discussion of this paper and various distant ancestors of it, many thanks to audiences in Durham, Jerusalem, Lisbon, MIT, NYU, Oxford, and Rutgers; to the anonymous referees; to Dorothy Edgington, John Hawthorne, John Hyman, Stephen Kearns, Hemdat Lerman, Adrian Moore, Sarah Moss, Anders Nes, Carl Posy, Richard Price, Katrina Przyjemski, Aynat Rubinstein, Judith Jarvis Thomson, Sebastian Watzl, and above all to Timothy Williamson.

References

- Barwise, J., & Perry, J. (1990). Semantic innocence and uncompromising situations. In A. P. Martinich (Ed.), *The philosophy of language* (2nd ed.), Oxford: Oxford University Press.
- Beall, J. C., & van Fraassen, B. C. (2003). *Possibilities and paradox*. New York: Oxford University Press.
- Black, M. (1962). Metaphor. In M. Black (Ed.), *Models and metaphors*. Ithaca, NY: Cornell University Press.
- Block, N. (1998). Semantics, conceptual role. In E. Craig (Ed.), *Routledge encyclopedia of philosophy*. London: Routledge. Retrieved December 01, 2006, from <http://www.rep.routledge.com/article/W037>.
- Boisvert, D., & Lubbers, C. (2003). Frege's commitment to an infinite hierarchy of senses. *Philosophical Papers*, 32, 31–64.
- Camp, E. (2004). The generality constraint and categorical restrictions. *Philosophical Quarterly*, 54, 210–231.
- Camp, E. (2005). Critical study of Joseph Stern, Metaphors in context. *Nous*, 39, 715–731.
- Camp, E. (2006). Metaphor in the mind: The cognition of metaphor. *Philosophy Compass*, 1(2), 154–170.
- Chomsky, N. (1957). *Syntactic structures*. The Hague: Mouton.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Davidson, D. (2001). On saying that. In D. Davidson (Ed.), *Inquiries into truth and interpretation*. Oxford: Clarendon Press.
- Davidson, D. (2001). What metaphors mean. In D. Davidson (Ed.), *Inquiries into truth and interpretation*. Oxford: Clarendon Press.
- Diamond, C. (2001). What nonsense might be. In C. Diamond (Ed.), *The realistic spirit*. Cambridge, MA: MIT press.
- Drange, T. (1966). *Type crossings*. The Hague: Mouton.
- Dummett, M. (1981). *Frege: Philosophy of language* (2nd ed.), London: Duckworth.
- Evans, G. (1982). *The varieties of reference*. Oxford: Clarendon Press.
- Fine, K. (2003). The non-identity of a material thing and its matter. *Mind*, 112, 195–234.
- Fogelin, R. (1988). *Figuratively speaking*. New Heaven: Yale University Press.
- Frege, G. (1952). On sense and reference. In M. Black & P. Geach (Eds.), *Translations from the philosophical writings of Gottlob Frege*. Oxford: Blackwell.
- Haack, R. J. (1971). No need for nonsense. *Australasian Journal of Philosophy*, 49, 71–77.
- Heim, I., & Kratzer, A. (1998). *Semantics in generative grammar*. London: Blackwell Publishers.
- Hodges, W. (2001). *Logic*. Harmondsworth: Penguin.
- Lakoff, G. (1971). Presupposition and relative well-formedness. In A. Jakobovitz & D. Steinberg (Eds.), *Semantics: An interdisciplinary reader*. Cambridge: Cambridge University Press.
- Lambert, K. (1968). On the no type theory of significance. *Australasian Journal of Philosophy*, 46, 79–86.
- Maddy, P. (1980). Perception and Mathematical intuition. *Philosophical Review*, 89, 163–196.
- Magidor, O. (2007). *Category Mistakes*. D.Phil dissertation, University of Oxford.
- Magidor, O. (2009). The last dogma of type confusions. *Proceedings of the Aristotelian Society*, 109, 1–29.
- Martinich, A. P. (1984). A theory of metaphor. *Journal of Literary Semantics*, 13, 35–56.
- McCawley, J. (1970). Where do noun phrases come from. In R. Jacobs & P. Rosenbaum, (Eds.), *Readings in English transformational grammar*. Boston: GINN, pp. 166–183.
- McCawley, J. (1971). Interpretative semantics meets Frankenstein. *Foundations of Language*, 7, 285–296.
- McDaniel, K. (unpublished manuscript). Categories of being.
- Montague, R. (1974). The proper treatment of quantification in ordinary English. In R. Thomason, (Ed.), *Formal philosophy: Selected papers by Richard Montague*. New Heaven, CT: Yale University Press.
- Moran, R. (1997). Metaphor. In B. Hale & C. Wright, (Eds), *A companion to the philosophy of language*. Oxford: Blackwell Publishers.
- Pap, A. (1960). Types and meaninglessness. *Mind*, 69, 41–54.
- Parsons, T. (1981). Frege's hierarchies of indirect senses and the paradox of analysis. *Midwest Studies of Philosophy*, 6, 35–57.
- Reimer M., & Camp, E. (2006). Metaphor. In E. Lepore & B. Smith (Eds.), *The Oxford handbook of philosophy of language*. Oxford: Oxford University Press.
- Routley, R. (1966). On significance theory. *Australasian Journal of Philosophy*, 44, 172–209.
- Routley, R. (1969). The need for nonsense. *Australasian Journal of Philosophy*, 47, 367–384.

- Russell, B. (1905). On denoting. *Mind*, 14, 530–538.
- Russell, B. (1908). Mathematical logic as based on the theory of types. *American Journal of Mathematics*, 30, 222–262.
- Ryle, G. (1938). Categories. *Proceedings of the Aristotelian Society*, 38, 189–206.
- Searle, J. (1979). Metaphor. In A. Ortony (Ed.), *Metaphor and Thought*, Cambridge: Cambridge University Press.
- Sorensen, R. (2001). *Vagueness and contradiction*. Oxford: Oxford University Press.
- Stanley J. (2002). Nominal restrictions. In G. Peters & G. Preyer (Eds.), *Logical form and language*. Cambridge, MA: OUP.
- Stern, J. (2000). *Metaphor in context*. Cambridge, MA: MIT Press.
- Strawson, P. F. (1950). On referring. *Mind*, 59, 320–344.
- Strawson, P. F. (1952). *Introduction to logical theory*. London: Methuen.
- Williamson, T. (1994). *Vagueness*. London: Routledge.