

Nonsense: a user's guide

Manish Oza (moza@uwo.ca)

[Forthcoming in *Inquiry*]

Abstract

Many philosophers suppose that sometimes we think we are saying or thinking something meaningful when in fact we're not saying or thinking anything at all: we are producing nonsense. But what is nonsense? An account of nonsense must, I argue, meet two constraints. The first constraint requires that nonsense can be rationally engaged with, not just mentioned. In particular, we can reason with nonsense and use it within that-clauses. An account which fails to meet this constraint cannot explain why nonsense appears meaningful. The second constraint requires that nonsense does not express thoughts. An account which fails to meet this constraint undercuts the critical force of the concept of nonsense. I offer an account which meets both constraints. The central idea is that to be under the illusion that some nonsense makes sense is to enter a pretence that the nonsense is meaningful.

1 Introduction

Many philosophers suppose that we can think we are saying or thinking something meaningful when in fact we're not saying or thinking anything at all: we are producing nonsense.¹ But what is nonsense?

Early twentieth-century philosophy of language abounded in allegations of nonsense. Russell (1919) suggested that violations of logical type-restrictions result in nonsense. Wittgenstein (*TLP* 6.54) held that much philosophical discourse, including most of his *Tractatus Logico-Philosophicus*, was meaningless. Ayer (1936) argued that verifiability is a semantic constraint, such that metaphysical language is meaningless. Ryle (1949) held that when we make a category mistake, we produce nonsense.

But the idea is still around. Here are two more recent uses of it:

[J]ust as we cannot know a priori or with Cartesian certainty whether any particular thing we think or say is true, so we cannot know a priori or with Cartesian certainty that in seeming to think or talk

¹I would like to thank Dominic Alford-Duguid, John Bunke, Lisa Doerksen, David Dyzenhaus, Mark Fortney, Danny Goldstick, Philip Kremer, Gurpreet Rattan, Melissa Rees, Catherine Rioux, Hamish Russell, Julia Jael Smith, Nick Stang, Dave Suarez, Mason Westfall and Jessica Wright for discussion and criticism. This paper was supported by SSHRC and OGS funding.

about something we *are* thinking or talking about – *anything at all*. We cannot know a priori *that we mean*. (Millikan 1984, 10)

Look, it can be true that people think they have a thought when they don't. They can be deeply attached to a linguistic formulation that upon reflection doesn't say anything. ... So a lot of our thoughts are really not thoughts, they're things masquerading as thoughts. (Appiah 2017)

The idea also plays an important role in recent debates about domain restriction in quantification (Shaw 2015; Mankowitz 2019) and co-predication (Gotham 2016; Liebesman & Magidor 2019).

Rather than giving a definition of nonsense, I'll take the concept to be fixed by its theoretical role. 'Nonsense' is a term of criticism. It picks out a failure to express, or have, a thought: a failure more basic than saying or thinking something significant but false (Routley 1969; Goddard 1970). It wouldn't be a useful term of criticism, however, if failures to make sense were always obvious. Rather, it's implicit in the way philosophers have used the concept that something can appear to make sense even though it does not. This is why many twentieth-century philosophers took it as their job to expose latent nonsense.

The concept of nonsense is closely related to the theory of meaning. The theory of meaning specifies the facts in virtue of which a sentence, as used by a particular speaker or community, has a given content. It thereby entails conditions for saying something meaningful at all; when we violate one of these conditions, we produce nonsense (Cappellen 2013, 25). We might wonder, however, whether meaning is the sort of thing that is suited to being theorized at all. If we experience something as meaningful, how can theory tell us it is meaningless? I'm not going to provide a defence of theorizing about meaning here.² But I would note that the concept of nonsense doesn't require us to exclude experience altogether: we can allow that meaning is ordinarily something that shows up in experience (Zwicky 2019). What we cannot say, if we are to use the concept of nonsense critically, is that the experience of meaningfulness is sufficient for meaning. Theorizing can sometimes show that an apparent thought was really nonsense.

In this paper I'm not going to argue that we should use the concept of nonsense in the way twentieth-century philosophers did; nor will I take a position on which cases count as nonsense.³ Rather, I'm interested in what nonsense would have to be, for the project of those philosophers to be a viable one. I'm interested, in other words, in what would be required of an account of nonsense for the concept to do critical work. In my view, such an account needs to explain two things:

²For discussion of the idea that meaning is what a theory of meaning is a theory of, see Dummett (1975) and McDowell (1976). For criticism, see Wright (1981).

³So I will not be arguing against philosophers who suggest that many cases of apparent nonsense – such as category mistakes – are really just necessary falsehoods (Pap 1960; Goldstick 1974; Camp 2004; Magidor 2013); for responses, see Routley (1969) and Glock (2015, s. 7). However, I will suggest at n. 29 below that my account may provide an alternative explanation for one phenomenon cited by these philosophers.

1. How nonsense can appear meaningful, and
2. How nonsense is in fact not meaningful.

These requirements lead to two constraints on an adequate account of nonsense. The **Engagement Constraint**, discussed in 2 below, is that nonsense can be used in certain ways. The **Austerity Constraint**, discussed in 3, is that nonsense does not express thoughts. Some extant accounts fail, I will argue, because they do not meet one constraint or the other. In 4, I offer an account which meets both constraints and respond to some objections. Finally, in 5 I conclude with a hard question raised by nonsense.

2 The engagement constraint

In this section I'll defend the following constraint on accounts of nonsense:

The Engagement Constraint Nonsense can be engaged with.

'Engagement' is a loose term for various ways in which we seem to use nonsense, not just mention it, bringing to bear our ordinary conceptual capacities. As Annette Baier writes (1967, 520), nonsense 'need be neither useless nor lacking in order and discipline'. We need to account for engagement if we are to explain the appearance of sense that nonsense can present. In particular, there are two main ways of using nonsense that we have to explain: reasoning with nonsense and embedding nonsense within that-clauses.

2.1 Reasoning with nonsense

We draw apparent inferences from nonsense. Different instances of nonsense will appear to license different inferences, depending on the concepts involved. I'll illustrate this point with two examples: a version of the Liar Paradox and a category mistake.⁴

So suppose the following sentence is nonsense:

- (1) 1 is false.⁵

To someone under the illusion that 1 expresses a thought, it will seem that 1 entails

- (2) '1 is false' is false.

Now consider a different instance of nonsense:

⁴Some deny that semantic paradoxes and category mistakes result in nonsense (see previous note). For my purposes, nothing hangs on this, as long as it is conceded that there exists nonsense that looks sufficiently sentence-like to underwrite 'inferences' like the one in the text. Any philosopher who wants to use 'nonsense' as a term of criticism must think that such instances exist.

⁵For accounts on which the Liar expresses no proposition, see Kripke (1975) and Rumfitt (2014).

(3) Goodness is hexagonal.

Someone who thinks that 3 makes sense will think 3 entails

(4) Something is hexagonal.

By contrast, someone who thinks that 1 makes sense would be ‘wrong’ to think that 1 entails 4. We have here a notion of ‘correctness in inference’ for nonsense.

Let us say that someone who reasons from 1 to 2, or from 3 to 4, is *reasoning with* the relevant instances of nonsense. Then our reasoning with 1 differs from our reasoning with 3. Of course, supposing 1 and 3 really are nonsense, these ‘inferences’ are defective. Nonsense doesn’t entail anything.⁶ But an account which fails to capture these sorts of moves is missing a crucial way we engage with nonsense, drawing on the conceptual and linguistic capacities we exercise in ordinary speaking and thinking.

Moreover, reasoning with nonsense can be central to our recognition of it as nonsense. This is characteristic of the critical use of the concept of nonsense by philosophers like Russell and Ryle (Goddard 1970, 12). Consider Russell’s discussion of existence. Russell held that existence was a second-order predicate applicable to first-order predicates. He denied that there was any first-order equivalent (1919, 206):

As regards the actual things there are in the world, there is nothing at all you can say about them that in any way corresponds to this notion of existence. It is a sheer mistake to say that there is anything analogous to existence that you can say about them. You get into confusion through language [...]. I mean, it is perfectly clear that, if there were such a thing as this existence of individuals that we talk of, it would be absolutely impossible for it not to apply, and that is the characteristic of a mistake.

To show that a certain concept doesn’t exist, Russell supposes that it does exist and considers how it would apply. That is, Russell notes that, if there were a first-order existence predicate, then for any object *o*, it would be true to say that *o* exists using this predicate. And Russell supposes that a genuine predicate cannot be such that it must be true of everything: if you can’t be wrong, then you can’t be right either; so there is no first-order existence predicate. But of course, if Russell is right about this, then there is no such thing as the thought that *o* exists, and therefore no such truth about *o*.

Here is another example. Recall

(1) 1 is false.

In reasoning with 1 I may start by supposing that 1 is false. Then I infer that, since 1 says that it is false, and it *is* false, 1 is true. Of course, if 1 is true, then, since it says that it is false, 1 is false. Now I know I’m in trouble, as 1 seems to be true if and only if it is false, but no sentence can be both true and false. So I

⁶Assuming that the notion of entailment is semantic and not merely syntactic.

conclude that 1 does not express a thought at all: it is nonsense, and is therefore neither true nor false and does not entail anything. Again, reasoning with 1 was crucial to recognizing the problem.

2.2 Nonsense-attributions

This brings us to the second way of engaging with nonsense. Once we recognize that a given thought is illusory, or that a given sentence doesn't express a thought, we can express our recognition of this fact using a that-clause (Sorensen 2002; Cappellen 2013).

The kind of thing I mean is best brought out with an example. Suppose that I've been reading Wittgenstein's *Tractatus* and am gripped by the Picture Theory of Meaning. Yes, I say: a picture really is a fact! (*TLP* 2.141) But as I progress through the book, I begin to suspect that its propositions are meaningless by its own lights. Wittgenstein, I think, was right: 'anyone who understands me eventually recognizes them as nonsensical' (6.54). I express my realization as follows:

- (5) It is nonsense to say that a picture is a fact.

Call such sentences *nonsense-attributions*. They seem to use rather than mention the phrase in the that-clause.⁷ For example, it seems that we could not understand 'It is nonsense to say that goodness is hexagonal' without having the concept of a hexagon. In general, the cognitive value of a nonsense-attribution, the rational difference that accepting it makes (Frege 1892), depends on understanding the concepts used within the that-clause.

Some might argue that nonsense-attributions are metalinguistic – that the nonsense is merely being mentioned. I don't deny that they can be read in a metalinguistic way, but there is also a non-metalinguistic reading.

On the metalinguistic reading, a sentence of the form 'It is nonsense to say that *S*' says the same as one of the form 'The string *S* does not express a thought'. On this reading, the cognitive value of a nonsense-attribution is that a particular sequence of symbols does not express a thought. Note that it's possible to appreciate this without being able to understand any element of the sequence of symbols: for example, I might learn that a particular sequence of letters in shorthand script doesn't say anything, even though I can't read shorthand.

On the non-metalinguistic reading, we need to understand the concepts used within the that-clause to understand the nonsense-attribution. To see that such a reading is available, let's return to the *Tractatus* example:

- (5) It is nonsense to say that a picture is a fact.

⁷I'm aware that nonsense-attributions sound odd. Perhaps this is better: 'To say "A picture is a fact" is nonsense'. I'm happy to grant this, as long as it is clear that this locution can be read in a non-metalinguistic way, where understanding the whole requires understanding the double-quoted words.

Now contrast this with another sentence:

(6) It is nonsense to say that das Bild eine Tatsache ist.⁸

If you don't speak German, you won't be able to engage with the content of the that-clause in 6, while you will be able to engage with its content in 5. For example, you won't be able to identify the concepts used within the that-clause. For this reason, 5 is informative in a way that 6 is not. However, on the metalinguistic reading, both 5 and 6 point out that certain strings do not express thoughts. The reading on which 5 says more than 6 must consider nonsense-attributions as more than metalinguistic.⁹ To be clear, I don't deny that the metalinguistic reading is possible, nor that it may play a role in our engagement with nonsense; however, it is not exhaustive of that engagement.

You might wonder how nonsense-attributions work semantically. For present purposes, I'm saying nothing about this: the point is simply that any viable account of nonsense ought to validate the intuition that nonsense-attributions can be informative in a non-metalinguistic way. This intuition doesn't dictate the semantics we give. Later in the paper, when I give my positive account of nonsense, I will explain what is going on in nonsense-attributions.¹⁰

I've described two ways we seem to use nonsense rather than only mentioning it: first, we reason with nonsense, and second, we make nonsense-attributions. To account for these phenomena, we must satisfy Engagement. I'll now show that a prominent account fails to satisfy the constraint.

2.3 The minimalist account

The *minimalist account* says that, when we produce nonsense, we produce mere words which fail to bear meaning: at the level of language, there is a string of symbols, but at the level of thought, there is nothing. A.W. Moore takes this position (2000, 198-9):

...there is one and only one way in which an utterance can fail to be a representation, namely because the words involved in it have not been assigned suitable meaning: they have not been assigned such meaning as would give the utterance content.

⁸You might object that 8 is ungrammatical (Quine 1961, 135). That might be right. But it would provide further evidence that nonsense-attributions can also be read in a way that is not covertly metalinguistic, as 'It is nonsense to say "das Bild ist eine Tatsache"' is unquestionably grammatical.

⁹Someone who favours the metalinguistic view might point out that they know more about 'A picture is a fact' than about 'Das Bild ist eine Tatsache': they know the meanings of the words in the former string. But on the metalinguistic view, this knowledge plays no part in understanding 5. So 5 is not more informative than 8. At best, 5 can be combined with other knowledge to draw inferences that cannot be drawn from 8. My point is that 5 is in itself more informative than 8 for our Anglophone *Tractatus* reader.

¹⁰Nonsense-attributions are related to another type of sentence in which nonsense is used within a that-clause, e.g. 'Wittgenstein believes that a picture is a fact'. It's disputed whether such a sentence makes sense (Diamond 2000, s. 4; Sorensen 2002), but there's some intuitive pressure to say it does. While I won't have room to discuss this, I believe my account can be extended to deal with such sentences.

Later, he writes that even the most apparently significant nonsense, like the philosophical nonsense we produce ‘when trying to express the inexpressible’, is ‘none other than the the pure and utter nonsense of ‘Phlump jing ux.’”

This position is often thought to follow from Frege’s Context Principle, which says that words have meaning only in the context of a sentence, or that concepts can occur only in the context of a thought (Frege 1884, x), or from a related principle about meaning. Supposing that some sentence does not express a thought, the suggestion is that the elements of the sentence cannot express any concepts either.

I’m going to argue that the minimalist account fails to satisfy Engagement. After this, I’ll show that the principle about meaning does not, on its best interpretation, support this account. The basic problem is that, on the minimalist account, when we produce nonsense we aren’t using our conceptual capacities. But our engagement with nonsense draws on precisely these capacities.

Let’s take the two forms of engagement in turn. First, recall the move from 3 to 4:

- (3) Goodness is hexagonal.
- (4) Something is hexagonal.

If the words in 3 don’t bear any meaning, but the words in 4 do bear meaning, then the string ‘hexagonal’ in 3 doesn’t have the same meaning as it does in 4. But any account of the move from 3 to 4 must invoke this common feature of the two sentences. If ‘hexagonal’ doesn’t have the same meaning in 3 and 4, then it’s unclear why moving from 3 to 4 is better than moving from 3 to, say, ‘Something is square’.¹¹ More generally, reasoning with nonsense relies on the presence of meaningful words and structures in nonsense. So the minimalist account cannot explain reasoning with nonsense.

Second, recall the difference between the nonsense-attributions 5 and 6:

- (5) It is nonsense to say that a picture is a fact.
- (6) It is nonsense to say that das Bild eine Tatsache ist.

If the words within the that-clauses in 5 and 6 do not bear any meaning, then the only difference between 5 and 6 is that they contain different strings of letters. But then we’re forced to adopt the metalinguistic reading of nonsense-attributions: what 5 and 6 tell us is that certain strings lack meaning. And we’re unable to explain why 5 seems to tell an English-speaking reader more than 6 does.

For these reasons, the minimalist account does not satisfy Engagement.

Now, all of this heavy weather about the minimalist account’s defects might seem unnecessary. After all, the proponents of the account don’t argue that we should accept it because it preserves the appearances. Rather, they argue

¹¹The fact that the same string, ‘hexagonal’, occurs in both 3 and 4 is insufficient, given the possibility of homonyms.

that we should accept it because it follows from an important principle about meaning. Here is Moore (2000, 199):

The guiding principle here is that there cannot be as it were positive semantic reasons for an utterance's failing to make sense. It cannot be because of what the parts of the utterance do mean, that the whole thing does not mean anything. The meaning of the parts is their contribution to the meaning of a range of wholes. ... The only thing about a word that can prevent utterances in which it is used from being proper representations is its not having any meaning at all.

Cora Diamond makes the same case in a different way. Supposing that the meaning of a term is its contribution to the meanings of sentences, then if a sentence is meaningless, its constituent terms must also be meaningless. So 'a sentence which does not make sense does not contain words which can be said to mean what they do elsewhere' (1981, 21; cf. Diamond 1988; Conant 2002). Now, if the meaning of a term is essential to its identity, then no term in a meaningless sentence is the same as a term in a meaningful sentence (1981, 11):

If I know the rules of the language, I know what a sentence composed in such-and-such a way out of such-and-such Elements says; but I do not know (there is no such thing as knowing) that what I see or hear is this Element, unless the whole of which it is part has a sense to which the meaning of this Element contributes in the way determined by the rules.

In both formulations, the argument depends crucially on the following principle:

CP The meaning of a term is its contribution to the meanings of sentences.¹²

I'm calling this principle 'CP' so as to remain neutral on how, if at all, it is related to Frege's Context Principle. For present purposes, the important questions are not about Frege interpretation; they are whether CP is plausible and whether it supports the minimalist account.

I'm going to argue that there are two ways of reading CP. On the first reading, the meaning of a term is its contribution to the meaning of the particular sentence in which it occurs. Read in this way, CP does support the minimalist account, but it is not plausible. On the second reading, the meaning of a term is its contribution to the meaning of sentences in general in which it occurs. Read in this way, CP is plausible, but it does not support the minimalist account.¹³

¹²To account for contextual contributions to content in Kaplan's sense (1989), we should say 'utterances', but for present purposes we can ignore this. CP is expressed in Moore's statement, quoted above, that 'The meaning of the parts is their contribution to the meaning of a range of wholes', and in Diamond's statement that 'To fix the meaning of ... Elements is to fix their contribution to the sentences of which they are ... parts' (Diamond 1981, 10).

¹³Glock (2015, s. 5) responds to the same argument in a different way.

On the first reading, CP says that the meaning of a term is its contribution to the meaning of the particular sentence in which it occurs. For example, the meaning of ‘hexagonal’ in 4 consists in its contribution to the meaning of 4. On this reading, we cannot ask for the meaning of a term in general – we cannot, for example, ask what contribution ‘hexagonal’ makes to sentences of the form ‘ a is hexagonal’. Rather, we can only ask what contribution a term makes in a particular meaningful sentence.

Read in this way, CP does support the minimalist account. For suppose that 3 has no meaning. Then there is nothing to be said about the contribution that ‘hexagonal’ makes in its occurrence in 3. The fact that ‘hexagonal’ does contribute to the meaning of other sentences, like 4, is irrelevant, given that meaning is particular rather than general. Put differently, if the meaning of a word is its contribution to the meaning of the sentence in which it occurs, then if the sentence in which it occurs has no meaning, then the word has no meaning either. This line of thought seems to underlie both Moore’s and Diamond’s arguments for minimalism. It explains why both Moore and Diamond suggest that the words in a meaningless sentence do not mean what they mean in meaningful sentences (Moore 2000, 199; Diamond 1981, 21).

The problem is that on this reading CP makes it unclear how linguistic understanding is possible. We learn the meaning of a word from a finite set of sentences. If there’s no general contribution that a word makes to each sentence of which it is a part, then it’s unclear how we could ever understand what a word means in a sentence we haven’t encountered before. In Diamond’s terms, we could never be sure that the terms in a new sentence are the same as terms we already know.¹⁴ So we should reject this reading of CP.

On the second reading, the meaning of a term is its contribution to the meaning of sentences in general in which it occurs. As Diamond puts it – stating the view she rejects – ‘it is the *general* possibility a word has of contributing to sense that confers meaning on it’ (1981, 18). Since meanings are general, we can learn the meaning of a word from some finite number of instances and apply this understanding in new cases.

This reading doesn’t trivialize CP. CP tells us that if we don’t understand the contribution a term makes to sentences, then we don’t understand the term. For example, suppose I have a sophisticated theory about some object which I associate with the sign \neg . Nevertheless, if I don’t understand that $\neg p$ is true when p is false, then I don’t understand the meaning of \neg .

Read this way, however, CP doesn’t support the minimalist account. The meaning of ‘hexagonal’, for example, is the general contribution it makes to the meaning of sentences. When the term occurs in 3, its meaning remains the same, because it doesn’t depend on the meaning of the particular sentence in which it

¹⁴Indeed, on Diamond’s view it’s unclear how we can ever understand a new sentence. We can’t know what its terms mean until we know what the sentence means. But how do we know what the sentence means, unless we know what its terms mean? Diamond writes (1981, 21): ‘That such-and-such a word is a working part of a sentence, and that it is *its* content we must grasp to understand the sentence ... cannot be told by observation.’ But she doesn’t say how it can be told.

occurs. The lack of meaning of the whole is no reason to deny that the parts have meaning.

Still less does it follow that, as Moore says, ‘It cannot be because of what the parts of the utterance *do* mean, that the whole thing does not mean anything.’ Suppose that ‘goodness’ refers to a normative property and ‘hexagonal’ ascribes a property applicable only to extended things. Then the reason why ‘goodness is hexagonal’ is meaningless is precisely because of what its parts do mean. Both ‘goodness’ and ‘hexagonal’ have meanings, but the latter is not defined for the kind of thing referred to by the former.¹⁵

I’ve argued that the minimalist account fails to satisfy Engagement. I’ve also shown that CP does not support the minimalist account. But I don’t want to leave the impression that the minimalist account is unmotivated. Rather, its deepest motivation comes from difficulties about nonsense at the level of thought.

Consider a remark Wittgenstein makes in the Preface to the *Tractatus*: ‘to be able to draw a limit to thought, we should have to find both sides of the limit thinkable (i.e. we should have to be able to think what cannot be thought). It will therefore only be in language that the limit can be drawn’. Wittgenstein is pointing to an asymmetry between language and thought: while there’s no difficulty in explaining why a sentence is meaningless without presupposing that it is meaningful, there is some difficulty in explaining why a ‘thought’ is meaningless without presupposing that it is meaningful.

For suppose we say that, when a subject attempts to understand a nonsense sentence, they exercise the same capacities they would exercise in understanding the sentence’s words and structure elsewhere. It’s unclear, then, why a subject attempting to understand nonsense fails to grasp a thought: the subject’s activity looks the same as in the good case. It looks like any substantive explanation of why a given attempt at thinking failed will characterize it in ways that assimilate it to successful thinking. The motivation for the minimalist account is that it precludes any such substantive explanation.¹⁶

I’m not going to respond to this issue here. I’ll respond to it later on by giving an account which allows for such substantive explanations without assimilating nonsense to sense at the level of thought.

3 The austerity constraint

In this section I defend a second constraint on accounts of nonsense:

The Austerity Constraint Nonsense does not express thoughts.

¹⁵I’m not claiming that this kind of explanation (sometimes called a coupling theory (Routley 1966, 180)) explains all cases of nonsense. Where a sentence contains a meaningless word – say Moore’s sentence ‘Scott kept a runcible at Abbotsford’ (Diamond 1981, 7) – it seems plausible that that word’s lack of meaning explains why the sentence is nonsense.

¹⁶Compare Conant (1992), who similarly rejects any substantive explanation of why logically alien thought is impossible.

In a nutshell: if nonsense expresses thoughts, it can play the same role as sense. To the extent that nonsense can play the same role as sense, the concept of nonsense loses its critical force. So nonsense must not express thoughts.

3.1 Nonsense thoughts

To approach the Austerity Constraint, let's consider an account which directly violates it: the *nonsense thoughts* account, on which nonsense expresses genuine but defective thoughts.¹⁷ This account easily satisfies Engagement. Since a nonsense thought is a thought, it can form part of chains of reasoning. A nonsense thought can also fall within a that-clause, as in nonsense-attributions.

Despite these benefits, nonsense thoughts are a dead end. They make nonsense too much like sense. Consider two important theoretical roles played by thoughts. First, thoughts explain a central kind of understanding: I've understood a sentence if I grasp the thought it expresses. Second, thoughts explain a central kind of communication: I've successfully communicated with you if you grasp the thought I intended to express (Grice 1957). Now, if there are nonsense thoughts, then both roles can be played by nonsense as well as by sense. Suppose that 3 expresses the nonsense thought that goodness is hexagonal. I understand 3 if I grasp this thought; if I grasp a different one, I've failed to understand 3. If I utter 3 to you and you grasp this thought, then I've used 3 to successfully communicate. But if a nonsense thought can do the same work as an ordinary thought, then the concept of nonsense loses its critical force. That a given area of discourse is nonsense becomes an evaluatively neutral fact, not a defect.

In response, the proponent of nonsense thoughts should explain how nonsense thoughts are defective compared to ordinary thoughts, blocking the assimilation of our engagement with nonsense to our engagement with sense. I consider such a response in the next section.

3.2 What's wrong with nonsense?

The natural way to spell out the defectiveness of nonsense, while maintaining that nonsense expresses thoughts, is to hold that nonsense thoughts are not truth-apt.¹⁸ The proposal is that nonsense thoughts exist but lack truth-value – they are neither true nor false – while ordinary thoughts have a truth-value. This is an attractive proposal. It is often tempting to characterize nonsense as neither true nor false. If nonsense expresses thoughts, we have an answer to the Engagement Constraint; if those thoughts are neither true nor false, then – since truth is an important aim of rational discourse (Dummett 1973, ch. 10; Priest 2006, s. 4.5) – we have a reason to avoid speaking nonsense, such that the concept of nonsense retains its critical force. I will argue, however, that when made precise the proposal fails.

¹⁷Sorensen (2002) argues, relatedly, that there are meaningless objects of belief (albeit statements rather than propositions).

¹⁸I'm going to talk about sentences as having truth-values, abstracting from the role of context.

Let's start by observing that the notion of being neither true nor false covers two different cases. First, there are cases like 'The King of France is bald', which Strawson (1950, 330) suggested was neither true nor false. As Dummett (1959) pointed out in response to Strawson, we know how the world would have to be for this statement to be true, and we also know that the world is not that way. So if someone used this statement to make an assertion, we could hold the assertion false: the world fails to be the way it would have to be for the assertion to be true. Despite this, statements like Strawson's are said to be neither true nor false because of their odd compositional behaviour. They have an internal negation ('The King of France is not bald') which, if used to make an assertion, is also false.

Second, there are cases like those we have been discussing in this paper – statements like

(3) Goodness is hexagonal.

Statements like this do display the same odd compositional behaviour as Strawson's example, but this doesn't exhaust the grounds for calling them neither true nor false. In the case of (3), we also have no idea how the world would have to be for the statement to be true. An assertion of (3) doesn't set up a condition that the world can meet or fail to meet, and so we say that the statement is neither true nor false.

In short, then, a statement may be said to be neither true nor false either because of its failure to say anything or because of its compositional behaviour. Dummett draws this distinction more precisely using two levels of content: *assertoric content* and *ingredient sense* (1973, ch. 12).¹⁹ These levels of content are associated with different truth-values. Assertoric content is what is expressed by a sentence in assertion. Dummett maintains that the point of assertion is to exclude possibilities: in asserting p , we indiscriminately exclude all possibilities other than those in which p holds. Either the actual world falls into the excluded class, or it does not; so at the level of assertoric content, every (meaningful) assertion is either true or false.

Ingredient sense, by contrast, tracks the contribution a sentence makes to complex sentences of which it is a part. Different sentences may have the same truth-value at the level of assertoric content, but behave differently when part of complex sentences. For example, consider the sentences a is F and b is F , and suppose that a is an empty name and b names an object that doesn't fall under F . Then at the level of assertoric content, a is F and b is F are both false. But a is not F is also false, while b is not F is true. Negation is sensitive to a distinction among kinds of falsity which doesn't show up at the level of assertoric content. In Dummett's terminology, a is F is false at the level of assertoric content but neither true nor false at the level of ingredient sense.

Now, the proponent of nonsense thoughts must specify at what level nonsense thoughts are neither true nor false. Consider someone who asserts 3. We're supposing that the speaker succeeds in saying something – namely, that goodness

¹⁹For further discussion, see Suszko (1977) and Shaw (2014).

is hexagonal. The speaker succeeds in expressing a certain kind of content – the content characteristic of declarative sentences. Some possibilities are excluded by this content. In fact, all of them are. So the actual world is excluded. So the assertion of 3 is simply false, at the level of assertoric content.

However, 3 will interact differently with negation than an ordinary falsehood. Suppose someone asserts

(7) Goodness is not hexagonal.

There's a reading of 7 on which it is false just as well as 3, and for just the same reason.²⁰ This compositional pattern is what it means for 3 to be neither true nor false at the level of ingredient sense. Thus, if nonsense expresses nonsense thoughts, then the only way in which these thoughts are neither true nor false is at the level of ingredient sense – the way in which Strawson's 'The King of France is bald' is neither true nor false.

The problem is that this is insufficient to capture the defectiveness of nonsense. Assertoric content is bound up with the rational purpose of assertion: it captures the success or failure of an assertion in light of that purpose. Ingredient sense, by contrast, is a technical device serving to capture compositional effects. Truth-values assigned at the level of ingredient sense have no normative weight. To say that an assertion is neither true nor false is not to say that it fails more badly than a false assertion, but only that it is a false assertion of a particular sort, displaying particular compositional behaviour. The assertion is bad only because, at the level of assertoric content, it is false. Nonsense, however, is supposed to be a failure more basic than the failure to say something true.

So a treatment of nonsense as expressing thoughts that are neither true nor false only at the level of ingredient sense obviates the critical force of the concept of nonsense.²¹ To maintain the critical force of the concept, nonsense has to be neither true nor false as a normative status, not merely as displaying distinctive compositional behaviour.

Nonsense, therefore, must lack assertoric content. But to hold that nonsense fails to have assertoric content, we must satisfy Austerity: deny that nonsense expresses thoughts altogether, rather than holding that the thoughts it expresses are neither true nor false.²² This is why the nonsense thoughts account fails.

4 A pretence account of nonsense

I've set out two constraints that an account of nonsense must meet. It's hard to meet both constraints, because they pull in opposite directions. The nonsense

²⁰There is also an 'external' reading of 7 on which it is true. But in ordinary cases the two readings don't come apart (Routley 1966, 181).

²¹For additional arguments against assimilating nonsense to falsity, see Routley (1969) and Goddard (1970). For additional discussion of the idea that nonsense is neither true nor false, see Magidor (2013, chapter 4).

²²This is consistent with using a three-valued logic to represent the interaction of nonsense with sense. Kripke, for example, uses the third value for sentences that 'do not express propositions' (1975, 701). See n. 31 below.

thoughts account satisfies Engagement, but fails to satisfy Austerity. The minimalist account satisfies Austerity, but not Engagement. In short: it's hard to explain our dealings with nonsense without assimilating nonsense to sense.

This tension is not only theoretical. We can see how it shows up in experience in Diamond's description of trying to understand someone who utters nonsense (2000, 157-8):

When you understand someone who utters nonsense, you are not, on the one hand, remaining as it were outside his thought and describing what goes on from the point of view of empirical psychology. But, on the other hand, you are not inside his thought as you are when he makes sense and you understand what he says, because there is no such internal understanding, there is no thought that such-and-such to understand. ... There is, as I said, no inside. But what it is to understand a person who utters nonsense is to go as far as one can with the idea that there is.

For Diamond, this is how we should read the *Tractatus*: by imaginatively entering into the illusion that its sentences express thoughts.

In this section I offer an account of nonsense in the spirit of Diamond's suggestion that engagement with nonsense is an exercise of imagination. First, I explain what semantic pretence is.²³ Next, I introduce some background about understanding and logical form. Finally, I draw on this background to give the account, showing that it satisfies both constraints.

4.1 Pretence and make-believe

Pretence accounts codify games of make-believe, of the sort that children play. A game of make-believe will typically involve some really existing items (*props*), about which something is expressly pretended to be the case (*initial stipulations*), and principles for generating further content in the pretence (*principles of generation*) (Armour-Garb & Woodbridge 2015, 39ff.). Consider, for example, a game about a bank robbery. The props are a child, X, and some newspapers. The initial stipulations are that X is a bank robber and the newspapers are cash. But the game's content is not fixed solely by these stipulations (Crimmins 1998, 5). Real-world facts – in particular, facts about the props and their relations – can be incorporated into the pretence. Thus, when X bolts with an armful of newspapers, in the pretence the bank robber is making a getaway.

For present purposes, we're interested in *semantic pretence* – pretence about the meaning of certain sentences or purported thoughts. The props are words, concepts and mental events; the initial stipulations assign meanings to these items which they ordinarily lack. In this way semantic pretence can increase the expressive resources of our language without increasing our stock of words.

²³The classic source for pretence accounts is Walton (1990; 1993). For a thorough recent discussion, see Armour-Garb & Woodbridge (2015). My view differs from that of Armour-Garb & Woodbridge in using pretence to explain our engagement with nonsense, rather than to explain the functioning of ordinary, successful language.

Recent philosophical history offers some reason to think that a semantic pretence account will meet our two constraints. Late in *The Varieties of Reference*, Evans sought to give an account of sentences containing empty names, like ‘Sherlock Holmes does not exist’ (1982, ch. 10). On the one hand, Evans held that empty names are genuinely used, not just mentioned, in such sentences. Assuming that using a name requires knowing what it refers to, this explains why you can’t understand ‘Sherlock Holmes does not exist’ without knowing who Sherlock Holmes is. On the other hand, Evans maintained that, since the sense of a name is a way of presenting its bearer, an empty name has no sense. It follows that sentences containing empty names do not express thoughts. Evans navigated between these two requirements with a semantic pretence account of the use of empty names. Insofar as Evans’ two requirements look a lot like the Engagement Constraint and the Austerity Constraint, we have reason to think that a semantic pretence account will be adequate to nonsense more generally.²⁴

Before considering nonsense, it’ll be useful to begin with an example where pretence operates on an otherwise meaningful sentence. Let’s stay with the bank-robbery pretence. I’ll say that a sentence is ‘make-believedly true’ when an assertion of it would be correct within the pretence (Evans 1982, 354ff.; Crimmins 1998, 15). The basic rules of the pretence, then, are as follows.²⁵

Bank robbery pretence²⁶

Props. The props are the terms ‘The bank robber’, ‘cash’ and sentences containing those terms.

Initial Stipulations. For a sentence P which contains one or more of ‘The bank robber’ and ‘cash’, P is make-believedly true if and only if P [‘The bank robber’/‘X’; ‘cash’/‘newspapers’] is true.

We also have two more general rules (Evans 1982):

Principle of Generation. If P is true, and if there is no set of make-believedly true sentences $Q_1 \dots Q_n$ such that if $Q_1 \dots Q_n$ were true then P would not be true, then P is make-believedly true.

Recursive Principle. If $P_1 \dots P_n$ are make-believedly true and if $P_1 \dots P_n$ were true R would be true and there is no set of make-believe truths $Q_1 \dots Q_n$ such that if $Q_1 \dots Q_n$ were true then R would not be true, then R is make-believedly true.

These rules specify the set of make-believedly true claims. Some claims are expressly made-believe, while others are determined by a function from real-world truths to make-believedly true claims. This makes it possible to recover real-world content from make-believe content, or vice versa (Yablo 2014, s. 10.2).

²⁴The connection between Evans’ pretence theory and nonsense is drawn by Moore (2003a, 188).

²⁵I use P , Q and R as variables running over sentences and P [‘a’/‘b’; ‘c’/‘d’; ...] for the result of replacing ‘a’ with ‘b’, ‘c’ with ‘d’ and so on in P .

²⁶I’m borrowing the helpful format of Armour-Garb & Woodbridge (2015).

Now consider an assertion of ‘The bank robber dropped the cash while fleeing.’ This assertion might lack a truth-value, as phrases like ‘The bank robber’ and ‘the cash’ may lack reference. But given the rules of the pretence, we know what has to be the case for the assertion to be make-believedly true: the assertion is correct, within the pretence, if and only if X dropped the newspapers while fleeing.

4.2 Logical form and partial understanding

A pretence account of nonsense raises special difficulties because the participants pretend not only that something is the case, but that something makes sense. To give such an account, we need some background about how we understand thoughts and sentences.

I’m going to suppose that recognition of logical form is essential to our understanding of thoughts and sentences. By ‘logical form’, I mean the way a thought or sentence is composed of its elements. I’ll speak of a sentence as ‘having’ or ‘being of’ a particular logical form, but this shouldn’t be taken to imply that a given sentence can have only one form (Oliver 2010). For example, ‘If Hegel wrote the *Logic*, then Hegel was brilliant’ is of the form ‘if p , then q ’ and of the form ‘if a is F , then a is G ’.

Why think that recognition of logical form is essential to understanding? Well, it’s commonly thought that our understanding of a thought or sentence consists in grasp either of its truth-conditions (Davidson 1967; Lepore & Ludwig 2002), or of the valid inferences to and from it (Brandom 1994), or both. Logical form contributes both to the truth-conditions of a thought and to its inferential role. So, whatever the right account is, recognizing logical form is essential.

Logical form is *general*. A given form makes a uniform contribution to every thought or sentence which has that form. For example, any thought of the form ‘ a is F ’ will be true if and only if a is F ; any thought of the form ‘ $\neg p$ ’ will be true if and only if p is false. This generality allows us to know things about thoughts and sentences that we don’t understand. Consider the sentence ‘Space is curved’. As long as I recognize the logical form of ‘Space is curved’, then – even if I don’t understand the sentence – I know that ‘Space is curved’ is true if and only if space is curved (Higginbotham 1989).²⁷ I know, too, that ‘Space is curved’ entails that something is curved. We can draw inferences when we recognize a sentence as expressing a thought of a certain form – even if we don’t know which thought.

This kind of understanding can come in stages, which it’s natural to think of as levels of analysis. For example, take

- (8) Space is curved and time is not real.

I may first recognize 8 as a conjunction:

8 is true if and only if ‘Space is curved’ is true and ‘Time is not real’ is true.

²⁷Higginbotham calls this the ‘less demanding notion of meaning’ (1989, 170; see also Higginbotham 1994, 102 and Armour-Garb & Woodbridge 2015, 158).

Next, I realize that the right hand conjunct is negated:

8 is true if and only if ‘Space is curved’ is true and ‘Time is real’ is false.

And so on. Dummett (1974) believed this process essential to the utility of logic: partial analysis reveals some licensed inferences, and further analysis lets us see further structure, licensing new inferences.

Can this process be extended to include non-logical concepts? In principle, it can. For example, suppose that I understand ‘curved’, but don’t know what ‘space’ refers to: with this partial understanding, I might already grasp that ‘Space is curved’ entails that space is not flat. But whether these inferences are valid depends on whether ‘curved’ makes the same contribution to every sentence of which it is a part. It could be that some inferences from ‘curved’ are valid when we’re talking about everyday objects and invalid when we’re talking about space. In other words, it depends on whether non-logical concepts are general in the same sense as logic. So while, in principle, the process could be extended to non-logical concepts, such extensions are *risky* in a way that logic is not.

4.3 Pretentious nonsense: the general idea

We’re now in a position to state a pretence account of nonsense. This is a pretence where something that is not meaningful is make-believable meaningful. In this section I explain the general idea; in the next section, I work through two examples.

When we produce nonsense, we enter a pretence where a certain sentence expresses a thought, or where a certain mental event was a thought. The props are not thoughts. At the level of language, they are nonsensical sentences; at the level of thought, they are mental events: failed attempts at using certain conceptual capacities to think. Neither nonsensical sentences nor failed attempts at thinking stand in logical relations, but we can pretend they do.

The initial stipulation is that a given item is or expresses a thought. Not any old thought, though: a thought of a certain form. Given the stipulation that some bit of nonsense expresses a thought of a certain form, we can exercise our conceptual capacities just as we do with a partially-understood sentence. This allows us to reason with nonsense. When we make a nonsense-attribution, we’re using the expressive resources of the pretence to point out the prop as nonsense.

The pretence account of nonsense satisfies both constraints. Engagement is satisfied because, as I will show, within a pretence we can both reason with nonsense and embed it in nonsense-attributions. Austerity is satisfied because we only pretend that nonsense expresses thoughts. Unlike the bank robbery game and others discussed by some pretence theorists (Armour-Garb & Woodbridge 2015, chs. 2 and 3), pretence does not operate here to allow a sentence to express a content other than the one it appears to express.

4.4 Examples of pretentious nonsense

Let's see how this works. I begin with an example of reasoning with nonsense and then give an example of a nonsense-attribution. Recall

- (3) Goodness is hexagonal.

An utterance of 3 doesn't say anything, but we may mistakenly suppose it does. In making this supposition, we enter a pretence defined by the following rules.

'Goodness is hexagonal' pretence

Props. The prop is the sentence 'Goodness is hexagonal'.

Initial Stipulations. 'Goodness is hexagonal' expresses a thought of the forms '*a* is *F*', '*a* is hexagonal' and 'Goodness is *F*'.

Principle of Generation. If *P* is true, and if there is no set of make-believedly true sentences $Q_1 \dots Q_n$ such that if $Q_1 \dots Q_n$ were true then *P* would not be true, then *P* is make-believedly true.

Recursive Principle. If $P_1 \dots P_n$ are make-believedly true and if $P_1 \dots P_n$ were true *R* would be true and there is no set of make-believe truths $Q_1 \dots Q_n$ such that if $Q_1 \dots Q_n$ were true then *R* would not be true, then *R* is make-believedly true.

Before I discuss reasoning with 3, notice that we can simulate a truth-condition for 3 within the pretence. It's make-believedly true that 3 expresses a thought of the form '*a* is *F*'. We know that a meaningful sentence of this form is true if and only if *a* is *F*, so we incorporate this knowledge into the pretence using the Principle of Generation. Then, by the Recursive Principle, it's make-believedly true that

- (9) 'Goodness is hexagonal' is true if and only if goodness is hexagonal.

Of course, 9 is nonsense, since its right-hand side is nonsense. So I don't propose to follow Higginbotham (1989, 156) in claiming that we can know purported truth-conditions like 9. Since 9 is nonsense, it is not true and so cannot be known. Or again, since 9 is nonsense, it does not express a thought, so, assuming the objects of knowledge are thoughts, it cannot be known. Rather, 9 expresses an illusion of understanding; the reasoning sketched explains its etiology.

Next, consider the apparent entailment from 3 to 4. First, note that it's make-believedly true that 3 expresses a thought of the form '*a* is hexagonal'. We know that a meaningful sentence of the form '*a* is hexagonal' entails 4, so we incorporate this knowledge into the pretence by the Principle of Generation. Then it's make-believedly true, by the Recursive Principle, that 3 entails 4.

Let me comment on a few suspicious-looking aspects of this account. First, you might worry that we set out to explain the apparent validity of the inference from 3 to 4, not the apparent truth of the claim that 3 entails 4. In fact, there's no deep problem here. As stated, the Principle of Generation only allows for incorporating *truths* into the pretence. To solve this problem, we'd need to

extend the Principle to allow for incorporating *rules of inference* as well. Just as what's make-believedly true governs what's correct to assert within the pretence, what's make-believedly valid governs what's correct to infer within the pretence. Given such an extension, we could incorporate a rule like 'From a thought of the form '*a* is hexagonal', infer that something is hexagonal', and use it to make-believedly validate the inference from 3 to 4.

Second, you might worry about the Initial Stipulations about the form of 3. For example, the pretence contains the stipulation that 3 is of the form '*a* is hexagonal'. Are we entitled to this, if 3 doesn't express a thought? We are. Earlier, I argued against Diamond's view that words in nonsense sentences don't bear their ordinary meanings; rather, they do bear their meanings, but nonsense sentences don't express thoughts. So there's no reason to deny that 3 involves the ordinary phrase 'is hexagonal': a phrase which invites completion by some appropriate singular term. This obvious fact about 3 shapes the pretence into which we enter when we suppose that 3 is meaningful.

This fits with a plausible idea about nonsense. When we're under the illusion that 3 makes sense, we're not supposing merely that it expresses some thought or other – for example, that it might be a coded expression of a state secret. Rather, we're under the illusion that it expresses a thought which has components in common with 'Goodness is rare' and 'France is hexagonal'. Among other things, then, it must be of the form '*a* is hexagonal'. Why does 3 invite this illusion in particular? Because the illusions associated with a particular item of nonsense are generated and constrained by our ordinary ability to recognize patterns in speech and writing.²⁸

Next, let me address the use of nonsense in nonsense-attributions. The idea will be that, in order to make a nonsense-attribution in a non-metalinguistic way, we have to enter into the pretence that the nonsense makes sense, and use the expressive resources of the pretence to point out 'from the inside' that it does not. This is what Evans called a 'game-to-reality shift' (1982, 369):

[S]omeone who utters such a sentence should be likened to someone who makes a move within a pretence in order to express the fact that it is a pretence. He is not like someone who tries to prevent a theatre audience from being too carried away by jumping up on the stage and saying: 'Look, these men are only actors, and there is no scaffold or buildings here—there are only props.' Rather, he is like someone who jumps up on the stage and says: 'Look, Suzanne and the thief over there are only characters in a play, and this scaffold and these buildings are just props.' The audience must be engaged, or be prepared to engage, in the make-believe, in order to understand what he is saying.

The analogue of a nonsense-attribution in the bank robbery game would be, 'all this cash is really just newspapers'.

²⁸For further discussion of the role of pattern recognition in reasoning, see Besson (2019a).

I argued above that there are two readings of nonsense-attributions: a metalinguistic reading, on which the nonsense-attribution is about a string of symbols, and a non-metalinguistic reading. I want to suggest, now, that the non-metalinguistic reading is available only where there's a pretence in which the nonsense-attribution is make-believable false: a pretence where the content of the that-clause make-believable expresses a thought. Where there's no such pretence, the nonsense-attribution can only be read metalinguistically.²⁹

Where there is an appropriate pretence, some sentence or mental event will be used as a prop. It will be make-believable true that the sentence in question expresses a thought, or that the mental event is a thought. Within the pretence, then, we gain an expressive resource for identifying the sentence or mental event in question: as that which expresses, or is, a particular thought. The nonsense-attribution exploits this expressive resource to identify the sentence or mental event in question and then states that – outside of the pretence – it does not express, or is not, a thought. When read in a non-metalinguistic way, then, the nonsense-attribution will be true if and only if the prop underlying the that-clause does not express, or is not, a thought.³⁰

Let's consider an example. Recall the nonsense-attribution

- (5) It is nonsense to say that a picture is a fact.

On the metalinguistic reading, 5 expresses the fact that the string 'a picture is a fact' is meaningless. But suppose there's a pretence that 'a picture is a fact' expresses a thought. Within that pretence, 5 is make-believable false.

Given that 5 is make-believable false, a non-metalinguistic reading of 5 is available.³¹ On this reading, 5 says that that which make-believable expresses a thought – namely, the sentence 'a picture is a fact' – does not express a thought. Supposing that 'a picture is a fact' really is nonsense, 5 says something true. But it identifies the nonsense using expressive resources available only to one who has entered into the pretence that it makes sense.

I mentioned earlier that nonsense-attributions sound a bit odd. Now we can see why. Within the pretence, 5 is simply false, as 'A picture is a fact' is make-believable meaningful. Outside of the pretence, we aren't engaging with

²⁹Given the role of form in my account, a string which is totally empty of form – say 'xg7*12d' – will, without further set-up, not be the subject of a pretence. Thus, 'it is nonsense that xg7*12d' can only be read metalinguistically. I think this is the correct result. To the extent that a nonsense sentence has some (even partial) syntactic form, there is the possibility of a pretence and a non-metalinguistic reading of the nonsense-attribution.

³⁰While there's no room to develop this, I believe this treatment can be extended to deal with sentences like 'Wittgenstein believes that a picture is a fact'. The truth of this sentence will depend (loosely speaking) on Wittgenstein's participating in a pretence in which the embedded sentence is make-believable true. This would provide an alternative explanation of propositional attitude ascriptions involving nonsense, which have sometimes been used to argue that apparent nonsense is meaningful (Magidor 2013, 59).

³¹This non-uniform treatment of nonsense-attributions entails some non-compositionality. For example, while 'It is nonsense to say that a picture is a fact' involves a game-to-reality shift, 'It is not nonsense to say that a picture is a fact' is more naturally understood as make-believable true. In fact, this is what we should expect. Compare an actor on stage who insists: 'Suzanne and the thief over there are not just characters in a play!'

the nonsense, so 5 can only be read metalinguistically. And surely we don't enter the pretence midway through 5, after 'to say'. So what exactly is going on, semantically? Someone who utters 5 is speaking from within the pretence to point out that, in fact, it is a pretence. They are using expressive resources that are only available within the pretence to communicate something about the pretence itself. This kind of utterance is parasitic on the existence of nonsense; in an ideal language, it would have no place. But in a language where nonsense exists, it has real cognitive value.

4.5 Objections and replies

This concludes my account of nonsense. Of course, semantic pretence accounts face no shortage of objections (Richard 2000; 2013; Azzouni 2018; Kroon 2018; Woods 2018). In this section I discuss two of the most serious.

First, as Jody Azzouni has noted, inference 'is invariably language-wide in scope. Sentences, nearly enough, from any area of discourse, may be employed to deduce results.' This raises the problem of how 'to understand inferences when they involve both pretence and non-pretence sentences from which non-pretence conclusions are drawn' (2018, 700-701). For present purposes, the worry concerns reasoning that moves from partly nonsensical premises to non-nonsensical conclusions. Suppose we go from 10 and 11 to 12:

- (10) A picture is a fact.
- (11) If a picture is a fact, then some facts are hard to understand.
- (12) Some facts are hard to understand.

What's happening here? 10 is wholly nonsense. 11 is (let us suppose) also nonsense because of its antecedent.³² But 12 is fine. The problem is that the reasoning from 10 and 11 to 12 seems to be either irrelevant or invalid. For if it takes place within a pretence, it can be, at best, make-believedly valid, establishing only the make-believe truth of 12, and is thus irrelevant to the genuine truth of 12. And if it takes place outside of a pretence, it is invalid, vitiated by the meaninglessness of 10 and 11.

There's no real problem here. We should cheerfully accept the first horn of the dilemma: the reasoning from 10 and 11 to 12 is valid only within the pretence that 10 expresses a thought. So it is only make-believedly valid, and it shows only the make-believe truth of 12. None of this undermines the fact (if it is one) that 12 expresses a truth. It only means that 10 and 11 are not a route to knowing that truth.³³

³²If nonsense is contagious in this way, we might use Bochvar's logic B3 (also known as the Weak Kleene scheme) to represent the interaction of nonsense with sense. Indeed, Bochvar (1937) interprets the third truth-value as 'nonsense' or 'meaningless'. However, the question of the correct 'logic of nonsense' is subtle and beyond the scope of this paper. See Halldén (1949), Åqvist (1962) and Goddard & Routley (1973, ss. 5.3 and 5.4).

³³This leaves intact the cases where nonsense is used in coming to know that it is nonsense – e.g. the Liar Paradox and Russell on existence. In such cases, the fact that a given item is nonsense is not known as the conclusion of an argument with nonsensical premises.

Second, you might worry that participating in a pretence requires the intention to participate in a pretence (Azzouni 2018, 693-694). But those who are under the illusion that some item of nonsense makes sense have no such intention. So they cannot be participating in a pretence.

I don't think that participating in a pretence requires the intention to participate in a pretence. Rather, it requires intentions to proceed in accordance with certain suppositions – suppositions which may not be true, but need not be known to be false. In the present cases, these are suppositions to the effect that a certain sentence is meaningful or that a certain mental event is a thought. In some cases (as in 3 above), anyone engaged in such reasoning would know that the sentences are nonsense. In other cases (as in the *Tractatus*), the usual speaker doesn't believe that the sentences in question are nonsense. Nor do they positively believe that these rather strange sentences make sense. Rather, they take for granted that the sentences make sense. Where the sentence is nonsense, proceeding on such a supposition suffices for entering a pretence.³⁴

It might seem puzzling how an ordinary speaker could intend to follow the rather technical rules that govern what is correct to say within a pretence. After all, ordinary speakers need not possess concepts like *logical form* or *make-believedly true*. But this problem isn't specific to nonsense. It's an instance of the more general problem of how speakers are related to the rules that govern what is correct to say within a language, the rules we codify in a theory of meaning. These rules are formulated using concepts like *truth-condition*, which speakers need not possess. The solution, in my view, is to distinguish between the practical ability a speaker has and the theoretical representation of that ability (Evans 1981; Fricker 1982; Peacocke 1986; Wright 1986b; Davies 1989). The theoretical representation should specify the features to which a speaker is sensitive, but it need not do so in terms which a speaker would recognize. In the case of nonsense, the rules for a pretence specify the features to which speakers are sensitive when they participate in the pretence, but they need not do so in terms which speakers would recognize.

5 Conclusion. Is there a transparent level of sense?

I've argued that an account of nonsense must meet two constraints: the Engagement Constraint and the Austerity Constraint. While many existing accounts fail to meet one constraint or the other, I've proposed an account which meets both. Of course, I haven't shown that there are no other ways of accounting for

³⁴As this suggests, my notion of a pretence is fairly lightweight. A pretence can exist in virtue of a single individual's mental states; pretences may, but need not, be shared. This allows for pretences involving little-understood scientific or technological concepts: these pretences are not widely shared, but they can exist in virtue of the mental states of individuals who possess the relevant concepts. It would be possible to go further and hold that pretences are abstract objects such that, for any description of a pretence, the pretence exists, but my account does not require going this far.

nonsense. So, while I claim that pretence is what explains our engagement with nonsense, this claim is subject to other explanations that might be devised.

I'll end by setting out a hard question which should be of interest to anyone who finds the idea of nonsense compelling.

Philosophers have often thought that illusions of sense can be long-lasting and difficult to overcome: while it may only take a second's reflection to see through 'Goodness is hexagonal', it may take a lifetime to get past deep-rooted logico-metaphysical confusions (Moore 2003a, 185). This raises the prospect of *content skepticism*: skeptical doubt about whether we are really expressing thoughts. Can we ever be certain that we are having a thought, or may our attempts always turn out to be nonsense?³⁵

We could reject this skepticism if there were a level of thoughts where illusions of sense are impossible: a level where there's no distinction between sense and the appearance of it. Call this a *transparent level of sense*. Such a level of sense would be analogous to a level of perceptual content at which it's impossible to be mistaken: for example, while you can be mistaken about whether you're listening to a trumpet, or looking at a copy of *Naming and Necessity*, perhaps you can't be mistaken about whether you're hearing a brassy tone, or seeing an orangey-red hue. Of course, it's disputed whether there is such a level of perceptual content. And it's just as unclear whether there is a transparent level of sense (cf. Millikan 1984, 92).

Many early analytic philosophers thought there was. Frege and the early Wittgenstein, for example, thought that a logically perfect language would display its structure with total lucidity, such that nonsense could never appear to be sense. A statement like 'the Good is more identical than the Beautiful', expressed in such a language, would be an incoherent jumble of signs (*TLP* 4.003). Later philosophers, like Austin and Ryle, may have implicitly supposed that the domain of transparent sense was the domain of ordinary language. When we go to the store and ask where the toothbrushes are, or give somebody directions to the library, there is just no prospect that our utterances will turn out to have been nonsense.

But I'm not sure. Nobody ever managed to find the logically perfect language, and the line between ordinary and non-ordinary language has proven hard to draw. It seems clear, though, that the question of a transparent level of sense is closely bound up with the problem of how meaning can belong both to theory and experience. This is a problem for future research.

³⁵We should distinguish this from a global skepticism, about whether all of our sentences might turn out to be nonsense. It might seem possible to hold that we can know whether we are having a thought, that there is a knowable condition for sense, while allowing that all our attempts might fail to meet it. This position is, however, fraught. The condition for having a thought would itself be the content of a thought; but then it isn't clear how we could know the condition. These problems of self-undermining are endemic to the philosophical terrain (Kripke 1982, 69-71; Heal 1989, ch. 6).

References

- Åqvist, L. (1962). Reflections on the logic of nonsense. *Theoria* 28(2), 138–157.
- Appiah, K. A. and D. Kodzi (2017). Interview with Kwame Anthony Appiah – Pt. I. *Oxford Review of Books*.
- Armour-Garb, B. and J. A. Woodbridge (2015). *Pretense and Pathology: Philosophical Fictionalism and its Applications*. Cambridge University Press.
- Ayer, A. J. (1936). *Language, Truth, and Logic*. London: V. Gollancz.
- Baier, A. (1967). Nonsense. In P. Edwards (Ed.), *The Encyclopedia of Philosophy, volumes 5 and 6*. Collier-MacMillan.
- Balcerak Jackson, B. (2009). Understanding and semantic structure: Reply to Timothy Williamson. *Proceedings of the Aristotelian Society* 109(1pt3), 337–343.
- Bochvar, D. A. and M. Bergmann (1937/1981). On a three-valued logical calculus and its application to the analysis of the paradoxes of the classical extended functional calculus. *History and Philosophy of Logic* 2(1-2), 87–112.
- Camp, E. (2004). The generality constraint and categorial restrictions. *Philosophical Quarterly* 54(215), 209–231.
- Cappelen, H. (2013). Nonsense and illusions of thought. *Philosophical Perspectives* 27(1), 22–50.
- Conant, J. (1992). The search for logically alien thought: Descartes, Kant, Frege, and the Tractatus. *Philosophical Topics* 20(1), 115–180.
- Conant, J. (2002). The method of the Tractatus. In E. H. Reck (Ed.), *From Frege to Wittgenstein: Perspectives on Early Analytic Philosophy*. Oxford University Press.
- Crimmins, M. (1998). Hesperus and Phosphorus: Sense, pretense, and reference. *Philosophical Review* 107(1), 1–47.
- Davidson, D. (1967). Truth and meaning. *Synthese* 17(1), 304–323.
- Davies, M. (1981). *Meaning, Quantification, Necessity: Themes in Philosophical Logic*. Routledge & Kegan Paul.
- Davies, M. (1989). Tacit knowledge and subdoxastic states. In A. George (Ed.), *Reflections on Chomsky*. Blackwell.
- Descartes, R. (1996). *Meditations on First Philosophy: With Selections From the Objections and Replies*. Cambridge University Press.
- Diamond, C. (1981). What nonsense might be. *Philosophy* 56(215), 5–22.

- Diamond, C. (1988). Throwing away the ladder. *Philosophy* 63(243), 5–27.
- Diamond, C. (2000). Ethics, imagination and the method of Wittgenstein's Tractatus. In A. Crary and R. J. Read (Eds.), *The New Wittgenstein*, pp. 149–173. Routledge.
- Dummett, M. A. E. (1959). Truth. *Proceedings of the Aristotelian Society* 59(1), 141–62.
- Dummett, M. A. E. (1973). *Frege: Philosophy of Language*. London: Duckworth.
- Dummett, M. A. E. (1975). What is a theory of meaning? (I). In S. Guttenplan (Ed.), *Mind and Language*. Oxford University Press.
- Evans, G. (1976). Semantic structure and logical form. In G. Evans and J. McDowell (Eds.), *Truth and Meaning*, pp. 49–75. Oxford University Press.
- Evans, G. (1981). Semantic theory and tacit knowledge. In S. Holtzman and C. M. Leich (Eds.), *Wittgenstein: To Follow a Rule*. Routledge.
- Evans, G. (1982). *The Varieties of Reference*. Oxford University Press.
- Frege, G. (1892/1960). On sense and reference. In P. Geach and M. Black (Eds.), *Translations from the philosophical writings of Gottlob Frege*. Basil Blackwell.
- Frege, G. and J. L. Austin (1884/1953). *Foundations of Arithmetic*. Blackwell.
- Fricker, E. (1982). Semantic structure and speakers' understanding. *Proceedings of the Aristotelian Society* 83(1), 49–66.
- Glock, H. (2015). Nonsense made intelligible. *Erkenntnis* 80(1), 111–136.
- Goddard, L. (1970). Nonsignificance. *Australasian Journal of Philosophy* 48(1), 10–16.
- Goddard, L. and R. Routley (1973). *The Logic of Significance and Context*. New York: Wiley.
- Goldstick, D. (1974). Against 'categories'. *Philosophical Studies* 26(5-6), 337–356.
- Gotham, M. (2016). Composing criteria of individuation in copredication. *Journal of Semantics*.
- Grice, H. P. (1957). Meaning. *Philosophical Review* 66(3), 377–388.
- Halldén, S. (1949). *The Logic of Nonsense*. Uppsala Universitets årsskrift.
- Heal, J. (1989). *Fact and Meaning: Quine and Wittgenstein on Philosophy of Language*. Blackwell.
- Higginbotham, J. (1989). Knowledge of reference. In A. George (Ed.), *Reflections on Chomsky*, pp. 153–74. Blackwell.

- Higginbotham, J. and G. Segal (1994). Priorities in the philosophy of thought. *Aristotelian Society Supplementary Volume 68*(1), 85–130.
- Kaplan, D. (1989). Demonstratives: An essay on the semantics, logic, metaphysics and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, and H. Wettstein (Eds.), *Themes From Kaplan*, pp. 481–563. Oxford University Press.
- Kripke, S. (1975). Outline of a theory of truth. *Journal of Philosophy* 72(19), 690–716.
- Kripke, S. (1982). *Wittgenstein on Rules and Private Language*. Harvard University Press.
- Kroon, F. (2018). Pushing the boundaries of pretence. *Analysis* 78(4), 703–712.
- Liebman, D. and O. Magidor (2019). Copredication, counting, and criteria of individuation: A response to gotham. *Journal of Semantics* 36(3), 549–561.
- Magidor, O. (2013). *Category Mistakes*. Oxford University Press.
- Mankowitz, P. (2019). Triggering domain restriction. *Mind and Language* 34(5), 563–584.
- McDowell, J. (1976). Truth-conditions, bivalence, and verification. In G. Evans and J. McDowell (Eds.), *Truth and Meaning*. Clarendon Press.
- Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories*. MIT Press.
- Moore, A. W. (2000). *Points of View*. Oxford University Press.
- Moore, A. W. (2003a). Ineffability and nonsense. *Aristotelian Society Supplementary Volume 77*(1), 169–193.
- Moore, A. W. (2003b). *Noble in Reason, Infinite in Faculty: Themes and Variations in Kants Moral and Religious Philosophy*. Routledge.
- Oliver, A. (2010). The matter of form : Logic’s beginnings. In T. J. Smiley, J. Lear, and A. Oliver (Eds.), *The Force of Argument: Essays in Honor of Timothy Smiley*, pp. 165–185. London: Routledge.
- Pap, A. (1960). Types and meaninglessness. *Mind* 69(273), 41–54.
- Peacocke, C. (1986). Explanation in computational psychology: Language, perception and level 1.5. *Mind and Language* 1(2), 101–123.
- Quine, W. V. O. (1961). Notes on the theory of reference. In *From a Logical Point of View*. Harper Torchbooks.

- Richard, M. (2000). Semantic pretense. In T. Hofweber and A. Everett (Eds.), *Empty Names, Fiction, and the Puzzles of Non-Existence*, pp. 205–32. CSLI Publications.
- Richard, M. (2013). *Context and the Attitudes: Meaning in Context, Volume 1*. Oxford University Press.
- Routley, R. (1966). On a significance theory. *Australasian Journal of Philosophy* 44(2), 172–209.
- Routley, R. (1969). The need for nonsense. *Australasian Journal of Philosophy* 47(3), 367–384.
- Russell, B. (1919). The Philosophy of Logical Atomism, lectures 5-6. *The Monist* 29(2), 190–222.
- Shaw, J. R. (2014). What is a truth-value gap? *Linguistics and Philosophy* 37(6), 503–534.
- Shaw, J. R. (2015). Anomaly and quantification. *Noûs* 49(1), 147–176.
- Sorensen, R. A. (2002). Meaningless beliefs and Mates’s problem. *American Philosophical Quarterly* 39(2), 169–182.
- Strawson, P. F. (1950). On referring. *Mind* 59(235), 320–344.
- Walton, K. L. (1990). *Mimesis as Make-Believe: On the Foundations of the Representational Arts*. Harvard University Press.
- Walton, K. L. (1993). Metaphor and prop oriented make-believe. *European Journal of Philosophy* 1(1), 39–57.
- Wittgenstein, L. (1921/1961). *Tractatus Logico-Philosophicus* (Trans. Pears and McGuinness). Routledge.
- Wittgenstein, L. (1953/2009). *Philosophical Investigations, 4th Edition* (Trans. Hacker and Schulte). Wiley-Blackwell.
- Wittgenstein, L. (1969). *On Certainty* (Ed. Anscombe and von Wright). Harper Torchbooks.
- Woods, J. (2018). Pretendism in name only. *Analysis* 78(4), 713–718.
- Yablo, S. (2014). *Aboutness*. Princeton University Press.
- Zwicky, J. (2019). *The Experience of Meaning*. McGill-Queen’s University Press.