

Inscrutability and ontological commitment

Berit Brogaard

Published online: 9 August 2008
© Springer Science+Business Media B.V. 2008

Abstract There are two doctrines for which Quine is particularly well known: the doctrine of ontological commitment and the inscrutability thesis—the thesis that reference and quantification are inscrutable. At first glance, the two doctrines are squarely at odds. If there is no fact of the matter as to what our expressions refer to, then it would appear that no determinate commitments can be read off of our best theories. We argue here that the appearance of a clash between the two doctrines is illusory. The reason that there is no real conflict is not simply that in determining our theories' ontological commitments we naturally rely on our home language but also (and more importantly) that ontological commitment is not intimately tied to objectual quantification and a reference-first approach to language. Or so we will argue. We conclude with a new inscrutability argument which rests on the observation that the notion of objectual quantification, when properly cashed out, deflates.

Keywords Ontological commitment · Inscrutability · Quine · Substitutional quantification · Objectual quantification · Rational requirement

1 Inscrutability and ontological commitment

There are two doctrines for which Quine is particularly well known: the doctrine of ontological commitment and the inscrutability thesis—the thesis that reference and

B. Brogaard (✉)
Philosophy Program, RSSH, Australian National University, Coombs Building, Canberra,
ACT 0200, Australia
e-mail: brogaardb@gmail.com; berit.brogaard@anu.edu.au

B. Brogaard
Department of Philosophy, University of Missouri, St. Louis, 599 Lucas Hall (MC 73),
One University Blvd., St. Louis, MO 63121-4499, USA

quantification are inscrutable. Here is a familiar articulation of the doctrine of ontological commitment:

A theory is committed to those and only those entities to which the bound variables of the theory must be capable of referring in order that the affirmations made in the theory be true (1948, pp. 13–14).

The doctrine of ontological commitment is thought to be essential to the project of ontology. In particular, it is thought to constitute the foundation for a proper method of ontology (Quine 1953). The proper method of ontology proceeds as follows: first we determine which theories best account for our empirical observations. Then we determine what objects are needed in order for our theories to come out true.

Inscrutability arguments are a range of arguments which arrive at the conclusion that reference is indeterminate or ‘inscrutable’. They include Quine’s Gavagai arguments, domain permutation arguments and cardinality arguments.¹ I shall assume that the reader is familiar with these arguments. Here is a toy version of the domain permutation argument for a language with the following four expressions: ‘John’, ‘my table’, ‘runs’, ‘is a thing made out of wood’. When we ask speakers of our toy language whether ‘John’ refers to my desk, they will insist that it does not.² After all, they will say, ‘John runs’ is true, whereas ‘my desk runs’ is not. Suppose, however, that ‘John’ picks out my wooden table, that ‘my table’ picks out my athletic friend John, that ‘runs’ picks out things made of wood, and that ‘is a thing made out of wood’ picks out things that run. ‘John runs’ is then true and ‘my desk runs’ false. For, the referent of ‘John’ (viz. my table) is a member of the extension of ‘runs’ (the set of things made of wood), whereas the referent of ‘my desk’ is not. We thus get the right truth-values for the sentences in our language even though ‘John’ refers to my table rather than John and ‘my table’ refers to John rather than my table. Since we cannot rule out that ‘John’ refers to my table and that ‘my table’ refers to John in our toy language simply by looking at the truth-values assigned to the sentences in the language by its speakers, it is indeterminate what the terms in the language refer to.

But as has been shown on countless occasions,³ domain permutations are possible for real languages with an indefinite number of expressions. And, as there is no way to determine whether *our* language is a permuted language, we must conclude that reference is indeterminate or ‘inscrutable’. For any given expression in our language, there is no fact of the matter what it refers to. Further: as existential quantification and reference are intimately connected (‘*n*’ refers to *n* iff ‘there is an *x* such that $x = n$ ’ is true),⁴ there is no fact of the matter what existential quantifiers quantify over.

¹ For the first kind of argument see e.g. Quine (1960, 1968); for the second kind see e.g. Quine (1968); Wallace (1977), Davidson (1979) and Putnam (1980); for the third kind see e.g. Putnam (1980). Cardinality arguments typically rest on the assumption that every consistent 1st-order theory has an infinite number of models which stand to each other in the submodel relation.

² Our speakers also happen to speak a second language with more words.

³ See e.g. the original argument in Putnam (1980).

⁴ Assuming that the quoted sentence and the constant ‘*n*’ belong to a particular language and that every term in the language refers to exactly one expression.

There are various things one can say in response to the inscrutability arguments (see e.g. Lewis 1983, 1984, Williams 2007). For example, one might insist that perceptual contact and speaker intentions can help to fix reference. If I say ‘my table is made out of wood’ when John is not around and I am looking at my table, it is plainly obvious that I am not referring to John. Even if this is correct, however, the problem seems to re-emerge when we refer to things that aren’t present, or when we refer to highly theoretic entities or entities too small to be seen with the naked eye.

Here we will focus on the question of whether one could possibly accept both Quine’s doctrine of ontological commitment and the inscrutability thesis. As several authors have remarked,⁵ the conclusions of the inscrutability arguments may seem to undermine the doctrine of ontological commitment and hence the possibility of Quinean ontology. The argument runs as follows: according to Quine’s theory of ontological commitment, a theory T carries commitment to Fs just in case Fs must be assumed to be among the values of the bound variables of T in order for T to be true. But the inscrutability arguments illustrate that for any theory T and for any predicate F there are acceptable interpretations of T which require that Fs be assumed among the values of the bound variables, *and* there are acceptable interpretations which do not require that Fs be assumed among the values of the bound variables. So, if reference is indeterminate, ontological commitment too would seem indeterminate. As Peter Hylton puts it,

there will be more than one way to interpret the ontological commitments of a theory: a different interpretation of the predicates will go along with a different interpretation of the ontology (2004, p. 128).

But if the very foundation of Quinean ontology is inscrutable, the possibility of Quinean ontology would seem annihilated.

Quine’s own writings foster a reply to this sort of objection to Quinean ontology. We can interpret our best theories relative to a particular framework; viz. our own. As Hylton puts it, ‘the only sense we can make of the idea of “the truth about the world” is in terms of our own theoretical understanding. Hence we take the ontological claims of our own theory seriously’ (2004, p. 133). Matti Eklund (forthcoming) argues that this line of response secures the possibility of ‘immanent ontology’ but not the possibility of ‘transcendent ontology’. That is, we can continue to read off the ontological commitments of our theories but we will never be able to determine what “really” exists independently of our own theoretical perspective. It would thus seem that we *do* have to take a loss if we accept the conclusions of the inscrutability arguments. We must give up hope of doing transcendent ontology (not that this was ever Quine project but still) and settle for immanent ontology. We will argue here, however, that this is no loss at all, as transcendent Quinean ontology is impossible even in the absence of an independent inscrutability argument, the reason being that Quine’s hope of an intimate tie between ontological commitment and objectual quantification cannot be fulfilled. We conclude with a new inscrutability argument which rests on the observation that the notion of objectual quantification, when properly cashed out, deflates.

⁵ See e.g. Hylton (2004); Fogelin (2004); Kemp (2006, p. 118), and Eklund (forthcoming).

2 The notion of ontological commitment

At first glance, Quinean ontology seems simple enough. First we determine which theories best explain our observations. If several competing theories explain our observations equally well, then we rely on other epistemic virtues to single out the best theories (e.g. simplicity, completeness, indispensability, etc.). We then determine the commitments of these theories. The commitments of a theory are the entities we must assume among the values of the bound variables of the theory if the theory is to be true. But, one wonders, what exactly are we to make of the jargon ‘the entities we must assume among the values of the bound variables of the theory if the theory is to be true’? Quine was no friend of modal analysis. In fact, Quine explicitly stated that the notion of ontological commitment belongs to the theory of reference:

As applied to discourse in an explicitly quantificational form of language, *the notion of ontological commitment belongs to the theory of reference*. For to say that a given existential quantification presupposes objects of a given kind is to say simply that the open sentence which follows the quantifier is true of some objects of that kind and none not of that kind (1953, pp. 130–131, italics added).

What does it mean to say that ‘the notion of ontological commitment belongs to the theory of reference’ as opposed to the theory of meaning proper? The only natural thing it could mean is that the notion of ontological commitment, like that of reference, is extensional or referentially transparent.⁶ If the notion of ontological commitment is referentially transparent, then we can substitute co-extensional expressions in ontological commitment contexts without any change in truth-value. For example, since ‘bachelor’ and ‘unmarried man’ are co-extensional (let’s suppose), “‘John is a bachelor” carries commitment to bachelors’ and “‘John is a bachelor” carries commitment to unmarried men’ have the same truth-value.

However, as Richard Cartwright (1954) has argued, extensional theories of ontological commitment face serious difficulties. As Pegasus and General Oreius do not exist, ‘Pegasus’ and ‘General Oreius’ are empty names.⁷ As they are both empty, they have the same extension (viz., the empty set or a gap). Hence, they are intersubstitutable *salva veritate* in extensional contexts. So, if ontological commitment is an extensional relation, then “‘Pegasus is” carries commitment to Pegasus’ and “‘Pegasus is” carries commitment us to General Oreius’ should have the same truth-value. But intuitively they do not. In Greek mythology, Pegasus is a winged horse, and in the movie *The Lion, The Witch and the Wardrobe*, General Oreius is a centaur, half man and half horse. So, intuitively, “‘Pegasus is” carries commitment

⁶ I am here following the standard practice of taking extensional relations to be referentially transparent. For discussion see e.g. Williamson (2006).

⁷ For Quine, of course, names are not individual constants but names with descriptive content. But even names with descriptive content are intersubstitutable in extensional contexts. For example, since ‘the president’ and ‘the son of George Bush Sr.’ are co-extensional, ‘the son of George Bush Sr.’ can be substituted for ‘the President’ in ‘the President is in Florida’ without any change in truth-value. Thanks to Luca Moretti here.

to Pegasus' and "'Pegasus is'" carries commitment to General Oreius' have different truth-values; the former is true, and the latter false. So, 'Pegasus is' and 'General Oreius is' ought to carry commitment to different things.⁸ 'Pegasus is' ought to carry commitment to Pegasus but not General Oreius, and 'General Oreius is' ought to carry commitment to General Oreius but not Pegasus. Indeed, some of the most famous passages from Quine seem to indicate that he thought that pairs of sentences such as 'Pegasus is' and 'there are centaurs' incur different ontological commitments. Here is one such passage:

We commit ourselves to an ontology containing number when we say there are prime numbers larger than a million; we commit ourselves to an ontology containing centaurs when we say there are centaurs; and we commit ourselves to an ontology containing Pegasus when we say Pegasus is (1948, p. 8).

But this is not what extensional theories of ontological commitment predict. So, extensional theories of ontological commitment must be rejected.

However, Cartwright's argument against an extensional treatment of ontological commitment is not knockdown. For extensionalists could bite the bullet and say that false existential sentences carry the same commitments. The important role for the Quinean notion of ontological commitment, it may be said, is to serve as a criterion for determining what our *very* best theories say exists. And the notion of ontological commitment can play this role even if it yields counterintuitive results for false sentences.

To block this move Terence Parsons (1967) presents a different sort of argument against extensional theories of ontological commitment. Where ' $\vartheta(\phi)$ ' stands for ' $\{x: \phi \text{ OC } x\}$ ', that is, the unique class of those entities to which ϕ is committed (p. 447),⁹ Parsons thinks a proper notion of ontological commitment should satisfy at least the following two constraints.

C1: There are atomic predicates P and Q such that
 $\vartheta(\exists xPx) \neq \vartheta(\exists xQx)$

C2: If ϕ is a logical consequence of γ , then
 $\vartheta(\phi) \subseteq \vartheta(\gamma)$

(C1) says that there are atomic predicates P and Q such that ' $\exists xPx$ ' and ' $\exists xQx$ ' have different commitments, for instance 'rose' and 'tulip'. (C2) says that if ϕ is a logical consequence of γ , then the set of entities to which ϕ is committed will be a subset of the set of entities to which γ is committed. For example, 'there are roses' is a consequence of 'there are red roses'. So, the set of entities to which 'there are roses' is committed is a subset of the set of entities to which 'there are red roses' is committed. Moreover, as the (1953) quote illustrates, Quine thinks there is an

⁸ Of course, our best theories may not include names like 'Pegasus' and 'General Oreius'. But the sentences we utter do. Moreover, the exact same points can be made with respect to empty common nouns such as 'unicorn' and 'centaur'. Thanks to Luca Moretti here.

⁹ Parsons is not claiming that sentences or theories are committed to sets. Rather, he is claiming that if the notion of ontological commitment is extensional, then there is a unique set of all the entities to which the sentence or theory in question is committed.

intimate relation between the set of entities to which ‘ $\exists xFx$ ’ is committed and the class of entities that *satisfy* ‘ Fx ’. Parsons proposes the following three possible relations.

- (A) $\exists(\exists x\phi) = \{x: \phi\}$
- (B) $\exists(\exists x\phi) \subseteq \{x: \phi\}$
- (C) $\{x: \phi\} \subseteq \exists(\exists x\phi)$

(A) says that the set of entities to which ‘ $\exists x\phi$ ’ is committed is identical to the set of entities that satisfy ‘ ϕ ’. (B) says that the set of entities to which ‘ $\exists x\phi$ ’ is committed is a subset of the set of entities that satisfy ‘ ϕ ’, and (C) says that the set of entities that satisfy ‘ ϕ ’ is a subset of the set of entities to which ‘ $\exists x\phi$ ’ is committed.

However, Parsons argues, each of (A), (B) and (C) is inconsistent with the conjunction of (C1) and (C2). Here are simplified versions of the proofs offered by Parsons in refutation of (A), (B) and (C):

To refute (A) and (B): For any sentence ‘ $\exists xFx$ ’ there is a sentence of the form ‘ $\exists x(Fx \ \& \ \sim Fx)$ ’ that entails it. By (A) or (B), the set of entities to which ‘ $\exists x(Fx \ \& \ \sim Fx)$ ’ is committed is a subset of (or is identical to) the set of entities that satisfy ‘ $Fx \ \& \ \sim Fx$ ’. But the set of entities that satisfy ‘ $Fx \ \& \ \sim Fx$ ’ is the empty set. So, by (A) or (B), the set of entities to which ‘ $\exists x(Fx \ \& \ \sim Fx)$ ’ is committed is a subset of (or is identical to) the empty set. As only the empty set is a subset of the empty set, it follows from (A) and (B) that the set of entities to which ‘ $\exists x(Fx \ \& \ \sim Fx)$ ’ is committed is the empty set. But ‘ $\exists xFx$ ’ is a consequence of ‘ $\exists x(Fx \ \& \ \sim Fx)$ ’. So, by (C2), the set of entities to which ‘ $\exists xFx$ ’ is committed is a subset of the set of entities to which ‘ $\exists x(Fx \ \& \ \sim Fx)$ ’ is committed. As the set of entities to which ‘ $\exists x(Fx \ \& \ \sim Fx)$ ’ is committed is the empty set, and only the empty set is a subset of the empty set, the set of entities to which ‘ $\exists xFx$ ’ is committed is the empty set.

To refute (C): ‘ $\exists x(Fx \ \vee \ \sim Fx)$ ’ is a consequence of ‘ $\exists xFx$ ’. By (C2), the set of entities to which ‘ $\exists x(Fx \ \vee \ \sim Fx)$ ’ is committed is a subset of the set of entities to which ‘ $\exists xFx$ ’ is committed. By (C), the set of entities that satisfy ‘ $Fx \ \vee \ \sim Fx$ ’ is a subset of the set of entities to which ‘ $\exists x(Fx \ \vee \ \sim Fx)$ ’ is committed. By transitivity, the set of entities that satisfy ‘ $Fx \ \vee \ \sim Fx$ ’ is a subset of the set of entities to which ‘ $\exists xFx$ ’ is committed. But the set of entities that satisfy ‘ $Fx \ \vee \ \sim Fx$ ’ is the universal set. As only the universal set has the universal set as a subset, the set of entities to which ‘ $\exists xFx$ ’ is committed is the universal set, thus violating (C1).

The conclusion of Parsons’ argument is that if the notion of ontological commitment is extensional, then for any sentence S, the unique set of the entities to which S is committed is either the empty set or the universal set. So, the set of entities to which ‘there are chairs’ is committed is either the empty set or the universal set, and so on for any other true (or false) sentence. So, all sentences (true or false) have the same commitments. ‘There are chairs’ and ‘there are unicorns’, if they carry commitment to anything, carry commitment to chairs. But accepting this consequence would clearly put an end to Quinean ontology. So, regardless of what Quine actually meant in his (1953), extensional theories of ontological commitment must be rejected. As Parsons points out, the notion of ontological commitment is intensional. So, if it admits of analysis, it will admit of modal analysis.

Though Quine was hostile towards modal analysis, many contemporary defenders of Quinean ontology are not. And, indeed, ontological commitment is most commonly thought of in modal terms. Many authors take a sentence's ontological commitments to be definable in terms of the demands the sentence imposes on the world (see e.g. Hylton 2004; Rayo 2007; Eklund 2006 and forthcoming). The following modal articulation of this criterion seems initially plausible (where T is an interpreted sentence or theory and the modal operator is that of metaphysical necessity):

Modal Criterion (MC)

A first-order fragment T carries commitment to Fs iff necessarily, if T is true, then Fs are among the entities over which the variables of T range.

Unfortunately, modal theories of ontological commitment fare no better than extensional theories. Or so I will now argue. (MC) faces a number of problems, among them the following three.¹⁰

First, there are the paradoxes of strict implication. $\Box q \vdash \Box(p \rightarrow q)$ and $\sim \Diamond p \vdash \Box(p \rightarrow q)$ are theorems in all systems of modal logic. But if necessarily, there are Fs, then for any sentence or theory T, it follows that necessarily, Fs are among the entities over which the variables of T range. So, necessarily, if T is true, then Fs are among the entities over which the variables of T range. If T is inconsistent, then 'T is true' is impossible. So, for any choice of F, it follows that necessarily, if T is true, then Fs are among the entities over which the variables of T range. Hence, any theory or sentence whatsoever carries commitment to any necessary entity, and any impossible theory or sentence carries commitment to anything whatsoever. So, if numbers exist necessarily, any sentence or theory carries commitment to numbers, and because naïve set theory is inconsistent, naïve set theory carries commitment to entities that rule out the existence of sets, hardly a satisfactory result.¹¹

Second, there is the problem of de re necessary properties. Sentences or theories stating that a given individual exists carries commitment to entities whose existence is required for the individual to have the de re necessary properties that it has. For example, in any world in which John exists, there will be zygotes. So, in any possible world in which John exists and it hasn't been determined whether there are zygotes, there will be zygotes. Hence, 'John exists, and it hasn't been determined whether there are zygotes' incurs a commitment to zygotes. But that is odd. A sentence or theory which states that it hasn't been determined whether there are Fs and which doesn't state that there are Fs shouldn't incur a commitment to Fs.¹²

¹⁰ For these and other problems with this criterion see Brogaard (2007b).

¹¹ If numbers exist necessarily, then numbers are among the entities of the domain of any world. Moreover, the variables of a theory range over entities in the domain. So, if numbers exist necessarily, then regardless of which sentence or theory and which world we consider, numbers are among the entities over which the variables of the theory range. By (MC), any theory or sentence is committed to numbers. If naïve set theory is inconsistent, then 'necessarily, if naïve set theory is true, then entities that rule out the existence of sets are among the entities over which the variables of naïve set theory range' is vacuously true. By (MC), naïve set theory incurs a commitment to entities that rule out the existence of sets.

¹² Of course, it might if it analytically entails that there are Fs. Thanks to Amie Thomasson here.

Third, there is the problem of descriptive names. Following Gareth Evans (1979) let ‘Julius’ be a rigid designator of the unique person who actually invented the zip. When ‘Julius’ is so introduced ‘there is a unique person to whom “Julius” (meaning what it actually means) refers’ carries commitment to the existence of Julius. Yet this is not what (MC) predicts. A descriptive name in a counterfactual scenario has the meaning it actually has if the meaning-conferring proposition by which the name is introduced is the same in the two scenarios. In the actual world ‘Julius’ has the same meaning as ‘the unique person who actually invented the zip’. Since the actual inventor of the zip is Whitcomb L. Judson, ‘Julius’ designates Judson in the actual world. But when people in counterfactual scenarios say ‘let “Julius” be a rigid designator of whoever actually invented the zip’, they don’t use the word ‘actually’ to refer to the world that is actual relative to us; rather, they use it to refer to the world that is actual relative to them. Even so ‘Julius’ has the very same propositional meaning in the two worlds; in both worlds it has the same meaning as ‘the actual inventor of the zip’.¹³ So, there are possible worlds in which there is a unique person to whom ‘Julius’ (meaning what it actually means) refers but in which Julius—the person whom ‘Julius’ designates in the actual world—does not exist. (MC) thus wrongly predicts that ‘there is a unique person to whom “Julius” (meaning what it actually means) refers’ does not carry commitment to the existence of Julius, which is not the result we want.

Here is a variation on the Julius problem. Suppose we stipulate in the actual world that ‘red’ is to be a rigid designator of the property that actually causes the seeming of red in normal observers in normal circumstances. Suppose further that in the actual world entities with surface spectral reflectance property *red* causes the seeming of red in normal observers in normal circumstances. ‘Red’ then rigidly designates the reflectance property *red*. ‘There are entities of which “red” (meaning what it actually means) is true’ carries commitment to red things. However, there are worlds in which (1) there are no entities with reflectance property *red*, (2) the reflectance property *blue* causes the seeming of red in normal observers in normal circumstances, and (3) ‘red’ has the same meaning as it does in the actual world, that is, it has been stipulated that ‘red’ has the same meaning as ‘the property that actually causes the seeming of red in normal observers in normal circumstances’. So, there are worlds in which there are entities of which ‘red’ (meaning what it actually means) is true but in which there are no red things. (MC) thus predicts that ‘there are entities of which “red” (meaning what it actually means) is true’ does not carry commitment to red things. Once again this is not the result we want.

Virtually everything that has just been said about (MC) carries over *mutatis mutandis* to the following well-liked criteria of ontological commitment¹⁴:

¹³ The meaning of ‘actually’ is a function from a property or a proposition to an extension at the world of utterance. For insightful discussion of these issues see Rumfitt (forthcoming).

¹⁴ The first criterion is endorsed by e.g. Armstrong (2004, 23f), Cameron (forthcoming), and Williams (forthcoming); the second is proposed by Cartwright (1954); Anderson (1957) and Rayo (2007) as an explication of Quine’s criterion.

Truth-Maker Criterion (TC)

T carries commitment to Fs iff necessarily, if T has a truth-maker, then there are Fs.

Logical Criterion (LC)

T carries commitment to Fs iff ' $T \rightarrow \exists x\phi$ ' is a truth of (free) logic for some predicate ϕ expressing F-hood.

As (TC) is a permutation on (MC), it encounters exactly the same problems as (MC). (LC) runs into versions of these problems as well. First, where T is inconsistent ' $T \rightarrow \exists x\phi$ ' is true for any choice of ' ϕ ', and where Fs exist of logical necessity, ' $T \rightarrow \exists x\phi$ ' is true for any choice of T. So, nominalism about numbers incurs a commitment to numbers, and naïve set theory incurs a commitment to entities that rule out sets. Second, (LC) does not count the analytic entailments of a theory among the theory's ontological commitments. But this has unwelcome consequences. For example, since 'there are bachelors \rightarrow there are unmarried men' is not strictly speaking a truth of free logic, 'there are bachelors' does not carry commitment to unmarried men.¹⁵ Third, 'there is a unique person to whom "Julius" (meaning what it actually means) refers \rightarrow Julius exists' is not a truth of free logic. So, 'there is a unique person to whom "Julius" (meaning what it actually means) refers' does not carry commitment to the existence of Julius. (LC) encounters a further problem. As ' $\exists xFx \rightarrow \exists xFx$ ' is trivial, ' $\exists xFx$ carries commitment to Fs' is trivial. But as the perennial debate over what our theories are committed to shows, the question of what ' $\exists xFx$ ' carries commitment to is hardly trivial.

It seems fair to conclude that modal and logical criteria of ontological commitment fall short. Of course, it is still open to argue that the notion of ontological commitment can be subjected to counterfactual analysis. One advantage of turning to counterfactual analysis is that some have argued on independent grounds that we need an account of counterfactuals that allows some counterfactuals with necessarily false antecedents to be non-trivial and informative, or more precisely, that allows their truth-values to be affected by the meaning of their consequents.

Here is one consideration in favor of a non-standard analysis of counterpossibles. On the classical Lewisian account of counterfactuals, counterpossibles are vacuously true in that they are always true; an impossibility counterfactually implies anything you like. And relatedly, they are uninformative in the sense that the consequent of a counterpossible makes no contribution to the truth-value, meaning or our understanding of the whole. But the claim that counterpossibles are vacuously true and uninformative runs counter to our intuitions. For example, we should expect that the consequents of 'If Hobbes had (secretly) squared the circle, sick children in the mountains of South America at the time would not have cared' and 'If Hobbes had (secretly) squared the circle, sick children in the mountains of

¹⁵ To avoid this problem Rayo offers the following criterion of ontological commitment: T carries commitment to Gs just in case ' $T \rightarrow \exists x\phi$ ' is a truth of (free) logic for some predicate ϕ expressing F-hood, and part of what it is to be F is to be G.

South America at the time would have cared' make some contribution to their truth-value.¹⁶ If the former is true, we should expect the latter to be false.

A non-trivial account of counterpossibles is suggested (but not defended) by David Lewis (1973, pp. 24–26; 1987, pp. 10–11), and developed by Daniel Nolan (1997); David Vander Laan (2004); Boris Kment (2006) and Berit Brogaard and Joe Salerno (forthcoming). Given such an account, a counterfactual criterion of ontological commitment may be articulated as follows:

Counterfactual Criterion (CC)

A first-order fragment T carries commitment to F s iff, if T were true, then F s would be among the entities over which the variables of T range.

As not every counterpossible is vacuously true, as we envisage the analysis, we do not arrive at the conclusion that inconsistent theories carry commitment to anything whatsoever. But counterfactual criteria encounter a horde of other problems. To mention just one. Counterfactual criteria do not validate what we will call 'Expansion'. Expansion, let's stipulate, says that if Σ carries commitment to χ s, then ' Γ & Σ ' carries commitment to χ s. Here is a counterexample. Let ' Σ ' be 'there are handless humans' and let ' Γ & Σ ' be 'there are handless humans, and all humans are disembodied brains'. The closest worlds in which there are handless humans are worlds in which some humans lost their hands in an accident. At those worlds humans-who-lost-their-hands-in-an-accident are among the entities over which the quantifier in 'there are handless humans' ranges. But at the closest world in which there are handless humans and all humans are disembodied brains, there are no humans-who-lost-their-hands-in-an-accident. So, humans-who-lost-their-hands-in-an-accident are not among the values of the variables of 'there are handless humans, and all humans are disembodied brains'. So, Expansion fails. But Expansion is exceedingly plausible. Here are some unintuitive consequences of rejecting it. If Expansion is invalid, then the commitments of a first-order fragment might change as we discover, say, new laws and add them to the fragment.¹⁷ Likewise, the commitments of a fragment of a theory may differ depending on whether we consider the fragment in isolation from or together with the rest of the theory. But this is exceedingly implausible.

It may perhaps be argued that our counterexample to Expansion commits a contextual fallacy, as the range of worlds with respect to which the premise is evaluated is not kept fixed when the conclusion is evaluated (see Brogaard and Salerno 2008). However, oddities ensue if we impose stricter standards of evaluation on arguments containing counterfactuals. To see this, consider the following instance of Expansion:

'There are blue swans' carries commitment to swans; therefore, 'there are blue swans and square prime numbers' carries commitment to swans.

¹⁶ The example is due to Nolan (1997, p. 544).

¹⁷ It is, of course, intuitive that if you expand a theory, then the commitments of the theory change. However, if Expansion fails, then if we expand a theory, then the commitments of the original (and unchanged) part of the theory change, which is highly unintuitive.

In evaluating the premise, we go to the closest worlds in which there are blue swans. These worlds are relatively near and so are *possible* worlds. But ‘there are blue swans and square prime numbers’ is not true at these worlds. So we must declare the conclusion vacuously true. But this is odd, as we have forsaken a vacuity treatment of counterfactuals with necessarily false antecedents.

Further trouble. Consider the following inference chain.

‘There are blue swans’ carries commitment to swans; hence, ‘there are blue swans and square prime numbers’ carries commitment to swans and everlasting gobstoppers.

Suppose the conclusion is vacuously true relative to a (fixed) sphere of *possible* worlds (i.e., the closest worlds at which there are blue swans). It follows that the inference chain is perfectly in order; for it is vacuously true relative to the sphere of possible worlds that if there are blue swans and square prime numbers, then ζ , for any ζ . But the sequence is not perfectly in order. So, the contextual-fallacy response is not a viable way for friends of (CC) to save face.

It is fair to conclude that there is little hope that the notion of ontological commitment can be successfully subjected to modal, logical or counterfactual analysis. But we are then left with no proper analysis of the notion of ontological commitment. It is tempting now to suggest that the notion should be treated as basic. We are familiar with this move. Defenders of heavy-duty metaphysics (e.g., Schaffer and Sider) like to argue that terms of art such as ‘priority’, ‘grounding’ and ‘fundamental’ are basic terms: they resist analysis. And I have no real quarrel with that. Arguably, we can claim some basic understanding of these terms, given our encounter with them in ordinary language (the truth-values of the disjuncts of a disjunction are prior to the truth-value of the disjunction, etc.).

But can we claim a similar basic understanding of the notion of ontological commitment? If ‘commit’ is intended to have a meaning radically different from that of its ordinary-language equivalent, I think the answer is ‘no’. Heavy-duty metaphysicians simply cannot resort to the view that ‘commitment’ has nothing in common with its ordinary-language equivalent, that it resists analysis but that we nonetheless have a good enough understanding of it. However, I do think we can claim basic understanding of the notion if the notion means what it does in ordinary language. And that just *is* what it means—as Quine is using it anyway. Compare Quine: ‘we commit *ourselves* to an ontology containing centaurs when we say there are centaurs; and we commit *ourselves* to an ontology containing Pegasus when we say Pegasus is’ (1948, p. 8, italic added). Quine here seems to indicate that it is not strictly speaking sentences and theories that have commitments but you and me. We are the primary bearers of ontological commitment, not sentences or theories. We take on ontological commitments to perform certain acts or work towards the realization of certain states of mind, when we say something.

Of course, not any old utterance will commit the utterer to the entities to which the uttered sentence purports to make reference. If I utter the sentence ‘Santa exists’ merely as a demonstration of how one ought to pronounce the words which make up the sentence, my act does not commit me to the existence of Santa. Likewise, if an actress on a stage says ‘just then the witch, to satisfy an itch, was flying on her

broomstick, thumbing for a hitch', she does not thereby commit herself to witches. And someone who says before his job interview 'I have butterflies in my stomach' does not thereby commit himself to stomachs with butterflies in them (see Yablo 1998, 2002, 2005).

We take on commitments when we make sincere assertive utterances, when we accept a sentence or a theory, when we presuppose something or take it for granted. What are the commitments we take on? No doubt that will depend on the intentional act in question. If we sincerely accept a given sentence, then we may be committed, by the norms of rationality, to believing that certain things exist.¹⁸ If we sincerely assert a given sentence, then we may be committed, by the norms of assertion, to assenting to certain existentially quantified sentences, and if we sincerely take the truth of a given sentence for granted, then we may be committed, by the norms of rationality, to taking certain existential claims for granted. For example, a sincere assertive utterance of the sentence 'there are numbers' requires us, by the norms of assertion, to assent to the sentence 'there are numbers' and a sincere acceptance of the sentence 'there are black holes' requires us, by the norms of rationality, to believe that there are black holes.

A word on the notion of rationality is here in order. The relevant notion cannot be one that prohibits belief in the truth of inconsistent theories or sentences. To see why not consider the following example from Mark Colyvan (2007). 18th century calculus is inconsistent, as it assumes that the value of a particular constant h is zero and non-zero in one and the same calculation. As 18th century calculus was among the best mathematical theories of the 18th century, 18th century mathematicians were committed to believing that there is a number that is zero and non-zero. And, as the example makes vivid, the relevant notion of rationality, of course, cannot be one that requires us to believe all classical consequences of what we believe. For by 'or' introduction, ' h is 0' entails ' h is 0, or there are trolls'. But h is not 0. So, by disjunctive syllogism, it follows that there are trolls. So, it is a theorem of 18th century calculus that there are trolls, and so if rational belief is closed under classical consequence, 18th century mathematicians were committed to believing that there are trolls.

Colyvan thinks that the example shows that we must reject classical logic. But there are further options. One is to argue that rationality requires that we do not accept contradictory theories. If this is so, then accepting 18th century mathematics is not blameless. But if sincere acceptances of 18th century mathematics are not blameless, then, arguably, sincerely accepting 18th century mathematics does not *rationaly* require that we undertake any commitments. In my view, however, this option is not very attractive, as it entails that we are not rationally required to believe the obvious and relevant consequences of our best theories. A second option (which I prefer) is to reject the closure principle that we are required to believe all classical consequences of what we are required to believe. If we take rational belief to be closed under relevant paraconsistent consequence, we can say that 18th century mathematicians were rationally required to believe that h is zero and

¹⁸ If epistemic rationality comes apart from practical rationality, then we can take all occurrences of 'rational' to be shorthand for 'epistemically rational'.

non-zero but deny that they were rationally required to believe that there are trolls. So, we can say that it is a theorem of 18th century mathematics that there are trolls and yet deny that this is a theorem which 18th century mathematicians were rationally required to believe.

So far, so good. But a problem here arises. The present proposal entails that if I were to sincerely accept the sentence ‘there are ghosts’, then I would be rationally required to believe that there are ghosts. But, arguably, I am, as a member of a modern society, rationally required *not* to believe that there are ghosts. So, if I were to accept the sentence ‘there are ghosts’, then I would be rationally required to believe and not to believe that there are ghosts. But, it would seem, the norms of rationality do not require me to do x and also not to do x .

There are two ways one can go about responding to this worry. One could argue that insufficient attention has been paid to the scope ambiguity that resides in conditional requirements. Instead of cashing out my requirement as a requirement to believe that there are ghosts given that I have accepted the sentence ‘there are ghosts’, one could cash out my requirement as the wide-scope conditional requirement that if I accept the sentence ‘there are ghosts’, then I believe that there are ghosts. If my requirement is a wide-scope conditional requirement, then the problem just outlined arises only if we allow the inference from ‘rationality requires that if p , then q ’ to ‘if p , then it is rationally required that q ’. However, this inference (also known as ‘detachment’) is fallacious (see Broome 2007a, b).¹⁹ If my requirement really is a wide-scope conditional requirement, then my sincere acceptance of the sentence ‘there are ghosts’ does not rationally require me to believe that there are ghosts. Rather: rationality requires that if I accept the sentence ‘there are ghosts’, then I believe that there are ghosts. From this, and the assumption that I accept the sentence ‘there are ghosts’, we cannot infer that rationality requires that I believe that there are ghosts or that I am rationally required to believe that there are ghosts. To make this further inference we would need the additional premise that sincerely accepting that there are ghosts is blameless. Here we can stipulate that an acceptance of a sentence or theory is blameless just in case rationality does not require that we do not accept the sentence or theory. Likewise, we will stipulate that an assertive utterance of a sentence is blameless just in case the norms of assertion do not require that we do not assert the sentence. If our ontological requirements are wide-scope conditional requirements, then there are two ways in which an acceptance or assertion can fail to impose commitments on us. An acceptance or assertion fails to incur commitment if it is insincere or prohibited by the norms of assertion or norms of rationality.

It should be noted that the proposal does not entail that there are sincere acceptances or assertions of existentially quantified statements with no commitments attached. Sincerely accepting ‘there are ghosts’ does not entail a rational commitment to believing that there are ghosts. But rationality requires that if we sincerely accept that there are ghosts, then we believe that there are ghosts, *and* it requires that we do not accept the sentence ‘there are ghosts’. So, it follows from the

¹⁹ More precisely, detachment is the inference from ‘rationality requires that if p , then q ’ and ‘ p ’ to ‘rationality requires that q ’.

proposal that Quine was not quite right when he said that we undertake commitments when we make existentially quantified statements. We do not if, for example, the norms of rationality require that we do not accept or assert them. In such cases rationality issues the conditional requirement that if we accept ' p ', then we believe that p , and the unconditional requirement that we do not accept ' p '.

Despite the obvious virtues of this solution to our problem, I am inclined to think that it doesn't get things quite right. One problem with the proposal is that it forces us to deny the intuitiveness of the claim that if I accept the sentence 'there are ghosts', then I am rationally required to believe there are ghosts. Further: it forces us to deny outright that 'there are ghosts' commits us to ghosts. Both consequences seem devastating.

It seems better to set up a coherent meaning for 'rationally requires' that gives the narrow-scope formula.²⁰ On this proposal, if you accept the sentence 'there are ghosts', then you are rationally required to believe that there are ghosts. Of course, this has the undesirable consequence mentioned above that there are conflicts of rationality. But even a wide-scope formula may lead to conflicts of rationality (even if it is less likely). For example, it could be that rationality requires that if p , then q , and also requires that p , and that if p , then not- q . For example, rationality requires that if I accept a sentence ' p ', then I believe that p . So, rationality requires that if I accept the sentence 'there are ghosts', then I believe that there are ghosts. But given that I have conclusive evidence against the existence of ghosts, rationality also requires that I do not believe that there are ghosts.²¹ So, rationality requires that if I accept the sentence 'there are ghosts', then I do not believe that there are ghosts. The right thing to do then is not to accept the sentence 'there are ghosts'. But if an evil demon tells me that he will make me form 500 false beliefs unless I accept the sentence 'there are ghosts', then arguably rationality also requires that I accept the sentence 'there are ghosts'. So, regardless of what I do, I will have complied with and failed to comply with the norms of rationality.

Here is another case.²² Suppose rationality requires that if you believe p , then you believe r , and that if you believe q , then you believe not- r . Suppose further that you have conclusive evidence for believing p and that you have conclusive evidence for believing q . In that case, rationality requires you to believe r and believe not- r . But we know that rationality requires you not both to believe r and to believe not- r . So rationality requires inconsistent things of you. The sort of case that would satisfy these principles is one where you can't help believing p because you can see it with your own eyes, and the same with q , possibly some kind of illusion.²³

Since conflicts of rationality are bound to happen regardless of whether we construe conditional requirements as wide-scope or narrow-scope, I shall, in what follows, assume that if we accept a sentence ' p ', then we are rationally required to believe that p , and that this is so even if rationality requires that we do not accept ' p '.

²⁰ Thanks to John Broome for offering this response.

²¹ As rationality supervenes on the mind, I take having-evidence to be a state of mind rather than a relation to a state of the world outside the mind.

²² This example is due to John Broome.

²³ An example might be the waterfall illusion where a stationary object appears to move and stay still. See Crane (1988).

3 Consequences for the project of ontology

We have argued that extensional, modal and counterfactual analyses of the notion of ontological commitment are highly problematic. Ontological commitment is not a property of a theory or sentence but something *we* take on when we (sincerely) accept or assert a sentence or theory. To say that T commits us to Fs is just to say that (sincerely) accepting or asserting T commits us to believing that there are Fs or assenting to 'there are Fs'. However, the notion of ontological commitment, so construed, is ontologically thin. Consider a language L in which all quantification is cashed out, meta-linguistically, in substitutional terms.²⁴ ' $\exists xFx$ ' is true iff ' Fx ' is true for some substitution instance substituting 't' for 'x', where 't' is any closed term in the language. So, which existentially quantified sentences are true in L will be determined by which subject-predicate sentences are true in L. Suppose, furthermore, that L is a truth-first language rather than a reference-first language.²⁵ What is true in L is then determined by what competent speakers sincerely assert. If competent speakers in L sincerely assert the sentence 'Mickey Mouse is a fictional mouse that talks', then 'Mickey Mouse is a fictional mouse that talks' is true in L. In these circumstances, if 'there are fictional mice that talk' is true in L and a speaker sincerely utters the sentence 'there are fictional mice that talk', then she commits herself to assenting to 'there are fictional mice that talk'.

What are the consequences of this observation? Well, both defenders of objectual and substitutional quantification have argued that only objectual quantification and a reference-first approach to language implies ontological commitment. According to Jody Azzouni, for example, given substitutional quantification and a truth-first approach to language 'Mickey Mouse is a fictional mouse' can be treated as true and 'Mickey Mouse is a fictional dog' can be treated as false, even if there is no truth-maker that makes the first true and the second false, and even if we have no reason to believe that there are entities the properties of which determine that the first is true and the second false (see Azzouni 2004, p. 57). So, when the quantifier is interpreted substitutionally, ' $\exists xFx$ ' does not incur a commitment to the existence of Fs. Or so goes the classical lore.

Likewise, Thomas Hofweber (2005) has argued that Stephen Schiffer's (1987) something-from-nothing inferences are acceptable given a truth-first and a substitutional-quantifier approach. Thus, from 'Jupiter has four moons', we can apparently legitimately infer 'the number of moons of Jupiter is four' and from that we can legitimately infer 'there is a number that is the number of moons of Jupiter'. We apparently got something from nothing. However, Hofweber argues, this appearance is illusory.²⁶ Read objectually, ' $\exists xFx$ ' follows from 'Ft' only if 't' is a *referring* or *denoting* term. If, however, 'the number of moons of Jupiter is four' is a

²⁴ For a defense of a substitutional interpretation of the quantifiers, see e.g. Wright (2007).

²⁵ This terminology is from Agustin Rayo (pers. comm.). In a truth-first language truth-values are assigned to sentences or theories holistically. What is true in such a language will depend in part on the evidence available to the speakers and what the speakers are willing to assent to given the evidence. In a reference-first language, reference is determined first, and truth-values are then assigned to sentences compositionally.

²⁶ For independent criticism of this argument see Brogaard (2007a). For a response see Hofweber (2007).

pleonastic variant of ‘Jupiter has four moons’,²⁷ then the function of ‘the number of moons of Jupiter’ and ‘four’ is not to refer or denote. Hence, given an objectual interpretation of the quantifiers, the aforementioned inference is invalid.

If the quantifiers are interpreted substitutionally (or inferentially), then the inference is valid, but no ontological commitment arises in going from ‘the number of moons of Jupiter is four’ to ‘there is a number that is the number of moons of Jupiter’. For, on a substitutional interpretation of the quantifiers, the existential (universal) quantifier is a device for asserting the disjunction (conjunction) of its instances. On a substitutional interpretation, ‘there is a number that is the number of moons of Jupiter’ is true in virtue of the truth of one of its instances, namely ‘the number of moons of Jupiter is four’. As ‘the number of moons of Jupiter is four’ carries no commitment to numbers, neither does the consequence ‘there is a number that is the number of moons of Jupiter’. Or so the argument goes.

These observations, however, are not quite right.²⁸ If we sincerely accept the sentence ‘Mickey Mouse is a fictional mouse’ and reject ‘Mickey Mouse is a fictional dog’ (as a speaker of L might do), then we are rationally required to believe that Mickey Mouse is a fictional mouse and disbelieve that Mickey Mouse is a fictional dog. So, we are committed to believing that there are fictional mice. Likewise, if we sincerely accept the sentences ‘the number of moons of Jupiter is four’ and ‘there is a number that is the number of moons of Jupiter’, then we are rationally required to believe that there is a number that is the number of moons of Jupiter. So, we are committed to believing that there are numbers. Whether we get ontological commitment or not depends on which attitude we legitimately take towards the sentence in question and not on the nature of the quantifier. Someone who sincerely accepts the sentence ‘there are three prime numbers between 2 and 9’ is rationally required to believe that there are three prime numbers between 2 and 9, and hence (by minimal closure rules for rational belief) rationally required to believe that there are numbers.

It may be said in response that (pace our earlier considerations) conditional requirements are to be given wide-scope and that an acceptance of an existentially

²⁷ To a first approximation, ‘*p*’ is a pleonastic variant of ‘*q*’ iff ‘*p*’ is ‘the result of movement and extraction that places particular parts of the syntactic material [of ‘*q*’] in special positions’. (Hofweber 2005, p. 267). Pleonastic variants will have the same truth-conditional content in spite of the fact that they differ in surface grammar. Thomasson (2007) argues further that pleonastic variants are analytically equivalent.

²⁸ Defenders of the pleonastic strategy might perhaps say that what they are doing is proposing an additional criterion for being ontologically committed. But, as we have already seen, there is no alternative criterion for being ontologically committed which is not fraught with difficulty. Of course, it might be said that an acceptance or assertion of a sentence or theory carries commitment only if the quantifiers are interpreted objectually. But this would be unmotivated (see below). One might also wonder what determines whether the quantifiers in a sentence or theory are objectual or substitutional. I suppose this will depend in part on the role the quantifiers play in the language. Hofweber (2005), for example, argues that quantifiers sometimes play only an inferential role. For example, when we say that a certain famous detective whose name we don’t remember lives on Baker Street, we do not intend to quantify over real entities in the domain. In such cases the quantifiers are plausibly substitutional. If the speakers have no clear intentions and their linguistic behaviour does not clearly indicate how they interpret the quantifiers, I suppose it is indeterminate whether the quantifiers are objectual or substitutional. Thanks to Amie Thomasson here.

quantified sentence is blameless only if the quantifiers are objectual. If this were so, then accepting a sentence or theory with substitutional quantifiers would not require us to believe the corresponding proposition. We would thus take on commitments only if we were to accept an objectually quantified sentence or theory.

However, this reply is amiss. An acceptance of ' p ' fails to be blameless just in case rationality requires us not to accept ' p '. But rationality imposes this sort of requirement only if we have conclusive evidence against ' p '. For example, rationality requires that we do not accept 'the number of moons of Jupiter is seventeen' because we have conclusive evidence against the truth of this sentence. But if 'the number of moons of Jupiter is four' is truth-conditionally equivalent to 'Jupiter has four moons', then we have evidence for the former just in case we have evidence for the latter. But if we have evidence for 'the number of moons of Jupiter is four', then surely we also have evidence for 'there is a number that is the number of moons of Jupiter' even when the quantifier is given a substitutional interpretation. So, rationality does not rule against accepting it. The objection is amiss.

But if our ontological commitments do not depend on the nature of the quantifiers in our language, then the claim that only objectual quantification implies ontological commitment is false. If we (sincerely) accept or assert a sentence that employs existential quantifiers, we take on commitments, regardless of whether the quantifier is objectual or substitutional. If, however, the quantifier is substitutional, then there is no guarantee that the entities "quantified over" are "really" in our domain of quantification. So, accepting a true (and indispensable) theory and taking on commitments is no foolproof way to determine what there really is. Quine's method is not a method of transcendent ontology.²⁹

This returns us to the problem presented by inscrutability arguments. The problem was that if we accept the conclusions of the inscrutability arguments

²⁹ A word on the work by Amie Thomasson is here in order. Thomasson argues that Quine's criterion is not necessary for ontological commitment (see e.g. her 2007). This is so, she argues, because what a theory is committed to cannot always be read off of the overt quantificational claims of the theory. If a given language allows the pleonastic transformation from 'there are nine planets' to 'the number of planets is nine', then both of these sentences will commit us to numbers. The reason both commit us to numbers is that (i) the second sentence is an analytic consequence of the first, and (ii) the application conditions for the word 'number', as used by ordinary folks, are satisfied, which, she argues, suffices for the word to have an extension. I am sympathetic to this proposal. But the thesis offered in this paper differs from Thomasson's in at least the following two respects: whether or not the second sentence commits us to numbers will depend on whether or not one sincerely accepts it. It is possible to refuse to accept the second sentence and in that way avoid commitment to numbers. Further, whether the first sentence commits us to numbers will depend on whether sincerely accepting it requires us to rationally believe that there are numbers. And whether this is so will depend on the norms governing rational behavior. In the philosophy room it may be rationally permissible to accept the first sentence and reject the second. Hence, in the philosophy room, it may be perfectly rational to believe that there are nine planets and yet deny that the number of planets is nine. In these circumstances, then, it may be that the second sentence commits us to numbers whereas the first does not. Of course, it is also possible to argue that if there really is an analytic entailment from 'there are nine planets' to 'the number of planets is nine' (as Thomasson would argue), it's not possible to accept the first and reject the second. Certainly, the standard reply will be to say that in the philosophy room it may be possible (even if not while using plain English). So the question that remains is what these different "higher" standards are for accepting, e.g., that there are numbers (where this is uttered in the philosophy room). One might say, with Thomasson, that most of these appeals to higher standards don't make sense (what more is it supposed to take to "really" think that there are numbers?) Thanks to Amie Thomasson for discussion of these issues.

(e.g., gavagai and permutation arguments), then transcendent ontology is impossible. For, given that there are many equally acceptable interpretations of our theories, our theories do not have determinate commitments apart from particular interpretations. So, the best we can do is to read off the ontological commitments of our theories when given a homely interpretation. But our home language could be a permuted language. So, even if our best theories are true, this approach will not necessarily help us determine what really exists. It will help us determine what really exists only on the further assumption that our home language is not a permuted language.

As we have now seen, however, we need not accept gavagai or permutation arguments to arrive at this sort of conclusion. All we need to do to arrive at this sort of result is ponder the notion of ontological commitment employed in the Quinean method. As we have seen, the notion of ontological commitment—or at least the only one that is not fraught with difficulty—is neutral on the question of how to interpret the quantifiers in our theories. So, by spelling out the commitments of our theories we have made no progress towards determining what really exists, not even on the assumption that the theories are true. At best, we have made progress towards determining what really exists only on the further assumption that the quantifiers in our theories are objectual (and that the function of all apparently referential terms is to refer).

4 A new inscrutability argument

We have argued that the Quinean method of ontology will help us determine what there is only on the further assumption that the quantifiers in our theories are objectual. We will now argue that there is no non-trivial way to draw an ontological distinction between substitutional and objectual quantification. So, even on the assumption that our theories are true and the quantifiers are objectual, we have no way of determining what really exists.³⁰

The standard way to draw the semantic distinction between the two kinds of quantification is as follows:

Objectual quantification: ‘ $\exists xFx$ ’ is true iff there is at least one object in the range of the variables that satisfies ‘ Fx ’.

Substitutional quantification: ‘ $\exists xFx$ ’ is true iff ‘ Fx ’ is true for some substitution instance substituting ‘ t ’ for ‘ x ’, where ‘ t ’ is any closed term in the language.

This distinction between objectual and substitutional quantification is unproblematic. However, the semantic condition for objectual quantification employs an existential quantifier, and for the objectual quantifier to be ontologically thick—for it to quantify over a domain of *real* or *ontologically thick* entities—this meta-linguistic quantifier must itself be ontologically thick. But what does it mean to say that the meta-linguistic quantifier is ontologically thick? It might be suggested that the meta-linguistic quantifier is ontologically thick just in case it is objectual. This,

³⁰ This argument owes much to Jody Azzouni’s (2004) argument for ontological independence as a mark of real existence, though his conclusion differs from mine.

however, won't do. Given the semantic condition, the meta-linguistic quantifier is objectual just in case there is at least one object in the range of the variables that satisfies 'is an object in the range of the variables that satisfies "Fx" '. But the right-hand side of this instance of the semantic condition contains a new existential quantifier. As we have no clue as to how to interpret this meta-meta-linguistic quantifier non-trivially, this suggestion fails to give us any insights into what it means to say that the meta-linguistic quantifier is ontologically thick.

Perhaps we can simply say that the quantifier in the semantic condition quantifies over a domain of *real* or *ontologically thick* objects. We might articulate this proposal as follows:

Objectual quantification: ' $\exists xFx$ ' is true iff there is at least one *real* object that satisfies 'Fx'.

However, this condition is not an improvement on the original. Like the original, the new condition leaves us with the question of how to interpret the quantifier. Saying that the meta-linguistic quantifier is objectual won't do, for reasons already articulated. It might perhaps be thought that we can just add further material to the right-hand side of the semantic condition. For example, we might take the right-hand side of the semantic condition to say: 'there is at least one object in an actual domain of real objects that satisfies "Fx"'. But this is to no avail. For, regardless of how much material we stuff into the semantic condition, we still have no clue as to how to interpret the meta-linguistic quantifier. Jody Azzouni makes a related point in term of ontological commitment: 'one can't read ontological commitments from semantic conditions unless one has already smuggled into those semantic conditions the ontology one would like to read off' (2004, p. 55). Though I disagree with Azzouni that one cannot read off ontological commitments from semantic conditions, I agree that the claim that the objectual quantifier quantifies over a domain of real objects is at best a postulate and isn't guaranteed by the semantic condition for the quantifier.

What are the consequences of this observation for the project of ontology? We argued above that the Quinean method of ontology will help us determine what there is only on the further assumption that the quantifiers in our theories are objectual. As it turns out, more is needed. Assuming that our best theories are true and that the quantifiers in our theories are objectual does not entitle us to draw any conclusions about what really exists. We are entitled to draw these kinds of conclusions only on the further assumption that the objectual quantifier quantifies over a domain of real objects. But there seems to be no way to demonstrate the truth of this further assumption. Of course, even if we cannot demonstrate the truth of this assumption, there may turn out to be hard facts about whether or not the objectual quantifier quantifies over a domain of real objects. But if these facts are not semantic, and this, as we have seen, they cannot be, then it is not clear what sorts of facts they could be. In virtue of what would it be true that the objectual quantifier quantifies over a domain of real objects if not the semantics governing the language containing the quantifiers?

It seems fair to conclude that if the semantic distinction between substitutional and objectual quantification does not suffice for drawing an ontological distinction between objectual and substitutional quantification, then there is no significant ontological distinction to be drawn. But then it will make no ontological difference

how we interpret the quantifiers in our language. Whether an existentially quantified statement, ‘ $\exists xFx$ ’, is true will depend on whether it has true substitution instances or whether it is true to say that there is an object that satisfies ‘ Fx ’, and whether the latter holds will depend at least in part on the conventions and norms endorsed by speakers of the language. But if the quantifiers in our language are not ontologically thick, then the Quinean method of ontology does not help us determine what there really is. Even if we could determine that our best theories are true, and that the quantifiers in our theories are objectual, we are left with no insights into the real nature of things.

5 Conclusion

On the face of it, Quine’s method of ontology would seem to be inconsistent with the inscrutability thesis (the conclusions of e.g. gavagai argument and permutation arguments). Given Quine’s method of ontology, we first identify the theories which best corroborate our observations and which are otherwise epistemically superior. We then determine the theories’ commitments by looking at what the theories quantify over. The conclusion of the inscrutability arguments is that for any referring expression it is indeterminate what it refers to. It follows that there is no determinate domain over which our theories quantify. For, on the assumption that every term refers to exactly one object, ‘ n ’ refers to n iff ‘there is an x such that $x = n$ ’ is true. So, if it is not determinate whether ‘ n ’ refers to n , then it is not determinate whether ‘there is an x such that $x = n$ ’ is true. But if our theories do not determinately quantify over anything, then the project of ontology is pointless.

A standard reply to this sort of argument is to say that even if we accept the conclusions of the inscrutability arguments, we can do ‘immanent ontology’ (Eklund forthcoming). We can read off the commitment of our theories by assuming a homely interpretation—we can rely on our home language. However, the prospects for ‘transcendent ontology’ are dim.

In this paper we have argued that we do not need an independent inscrutability argument to arrive at the conclusion that transcendent ontology is impossible. Quine’s method does not give us any insights into what really exists because the only viable notion of ontological commitment is neutral with respect to the interpretation of the quantifiers. Even if it has been determined that ‘there are Fs’ commits us to Fs, and that ‘there are Fs’ is true, we cannot conclude that there really are Fs.

Acknowledgements The paper has benefited from correspondence with J.C. Bjerring, David Chalmers, Mark Colyvan, Patrick Greenough, Al Hajek, Michael Lynch, Joe Salerno, Jonathan Schaffer, Susanna Schellenberg, Wolfgang Schwartz, Manidipa Sen, and Weng Hong Tang and the audience discussion in Sydney. Special thanks to John Broome, Luca Moretti and Amie Thomasson for written comments on the paper and to Luca Moretti for organizing a great conference.

References

- Anderson, A. (1957). Richard Cartwright, R. *Ontology and the theory of meaning*, *Philosophy of Science* (vol. 21, pp. 316–325). *Journal of Symbolic Logic*, 22, 393–394.

- Armstrong, D. (2004). *Truth and truthmakers*. Cambridge: Cambridge University Press.
- Azzouni, J. (2004). *Deflating existential consequence. A case for nominalism*. New York: Oxford University Press.
- Brogaard, B. (2007a). Number words and ontological commitment. *The Philosophical Quarterly*, 57, 1–20.
- Brogaard, B. (2007b). Making sense of ontological commitment. *Ontological Commitment Conference*, Sydney, Nov. 30–Dec. 1, 2007.
- Brogaard, B., & Salerno, J. (2008). Counterfactuals and context. *Analysis*, 68, 39–46.
- Brogaard, B., & Salerno, J. (Forthcoming). Remarks on counterpossibles. In V. Hendricks, et al. (Eds.). *Between logic and intuition: David Lewis and the future of formal methods in philosophy*. Synthese Library.
- Broome, J. (2007a). Wide or narrow scope? *Mind*, 116, 359–370.
- Broome, J. (2007b). Requirements. In T. Rønnow-Rasmussen, B. Petersson, J. Josefsson, & D. Egonsson (Eds.), *Homage à Wlodek: Philosophical papers dedicated to Wlodek Rabinowicz*. www.fil.lu.se/hommageawlodek.
- Cameron, R. (Manuscript). *Truthmakers and ontological commitment: Or, how to deal with complex objects and mathematical ontology without getting into trouble*.
- Cartwright, R. (1954). Ontology and the theory of meaning. *Philosophy of Science*, 21, 316–325.
- Colyvan, M. (2007). The ontological commitments of inconsistent theories. In *Ontological Commitment Conference*, Sydney, Nov. 30–Dec 1, 2007.
- Crane, T. (1988). The waterfall illusion. *Analysis*, 48, 142–147.
- Davidson, D. (1979). The inscrutability of reference. *Southwestern Journal of Philosophy*, 10, 7–20.
- Eklund, M. (2006). Neo-Fregean ontology. In J. Hawthorne (Ed.). *Philosophical perspectives*. (Vol. 20, pp. 95–121).
- Eklund, M. (Forthcoming). The ontological significance of inscrutability. *Philosophical Topics*.
- Fogelin, R. (2004). Aspects of Quine's naturalized epistemology. In R. Gibson (Ed.), *The Cambridge companion to Quine* (pp. 19–46). Cambridge: Cambridge University Press.
- Hofweber, T. (2005). A puzzle about ontology. *Nous*, 39, 256–283.
- Hofweber, T. (2007). Innocent statements and their metaphysically loaded counterparts. *The Philosophers' Imprint*, 7, 1–33. <http://www.philosophersimprint.org/007001/>.
- Hylton, P. (2004). Quine on reference and ontology. In R. Gibson (Ed.), *The Cambridge companion to Quine* (pp. 115–150). Cambridge: Cambridge University Press.
- Kemp, G. (2006). *Quine: A guide for the perplexed*. New York/London: Continuum.
- Kment, B. (2006). Counterfactuals and the analysis of necessity In J. Hawthorne (Ed.). *Philosophical perspectives* (Vol. 20, pp. 237–302).
- Lewis, D. (1973). *Counterfactuals*. Oxford: Blackwell.
- Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, 61, 343–377.
- Lewis, D. (1984). Putnam's paradox. *Australasian Journal of Philosophy*, 62, 221–236.
- Lewis, D. (1987). Counterfactuals and comparative possibility. *Journal of Philosophical Logic*, 2(1973), 418–446. Reprinted in *Philosophical Papers*, Vol. II, June 1987, pp. 3–32.
- Nolan, D. (1997). Impossible worlds: A modest approach. *Notre Dame Journal for Formal Logic*, 38, 325–527.
- Parsons, T. (1967). Extensional theories of ontological commitment. *The Journal of Philosophy*, 64, 446–450.
- Putnam, H. (1980). Models and reality. *Journal of Symbolic Logic*, 45, 464–482.
- Quine, W. V. (1948). On what there is (Reprinted from *From a logical point of view*, pp. 1–19, by W. V. Quine, Ed., 1953, Cambridge, MA: Harvard University Press).
- Quine, W. V. (1953). Notes on the theory of reference. In W. V. Quine (Ed.), *From a logical point of view* (pp. 130–138). Cambridge: Harvard University Press.
- Quine, W. V. (1960). *Word and object*. Cambridge: MIT Press.
- Quine, W. V. (1968). Ontological relativity. *Journal of Philosophy*, 65, 185–212.
- Rayo, A. (2007). Ontological commitment. *Philosophy Compass*, 2/3, 428–444.
- Rumfitt, I. (Forthcoming). *Logical necessity*.
- Schiffer, S. (1987). *Remnants of meaning*. Cambridge: MIT Press.
- Thomasson, A. (2007). Deflating existence. In *Ontological Commitment Conference*, Sydney, Nov. 30–Dec. 1, 2007.
- Vander Laan, D. (2004). Counterpossibles and similarity. In F. Jackson & G. Priest (Eds.), *Lewisian themes* (pp. 258–275). Oxford: Clarendon Press.

- Wallace, J. (1977). Only in the context of a sentence do words have any meaning. *Midwest Studies in Philosophy*, 2, 144–164.
- Williams, J. R. G. (2007). Eligibility and inscrutability. *Philosophical Review*, 116, 361–399.
- Williams, J. R. G. (Manuscript). *Fundamental and derivative truths*.
- Williamson, T. (2006). Indicative versus subjunctive conditionals, congruential versus non-hyperintentional contexts. In E. Sosa & E. Villanueva (Eds.), *Philosophical issues, Vol. 16: Philosophy of language* (pp. 310–333). Oxford: Blackwell.
- Wright, C. (2007). On quantifying into predicate position: Steps towards a new(tralist) perspective. In M. Potter (Ed.), *Mathematical knowledge*. Oxford: Oxford University Press.
- Yablo, S. (1998). Does ontology rest on a mistake? *Aristotelian Society, Supplementary Volume*, 72, 229–261.
- Yablo, S. (2002). Abstract objects: A case study. *Philosophical Issues*, 12, 220–240.
- Yablo, S. (2005). The myth of the seven. In M. E. Kalderon (Ed.), *Fictionalism in metaphysics* (pp. 88–115). Oxford: Oxford University Press.