

# Ontological Commitment: Quine and Beyond

## Ontological Commitment in Medieval Logic: Winter School

Philipp Blum, December 13, 2016

To say that the name  $x$  denotes a given object  $a$  is the same as to stipulate that the object  $a$  [...] satisfies a sentential function of a particular type. In colloquial language it would be a function which consists of three parts in the following order: a variable, the word ‘is’ and the given name  $x$ .  
(Tarski 1956: 194)

**My aim for the ‘course’:** To make plausible, by way of a review of the history of Quine’s notorious so-called ‘criterion’ for ontological commitment (“to be is to be the value of a bound variable”) that the role traditionally assigned to it (to ‘reglement’/‘regiment’/‘regulate’ philosophical discussion) is better split into two – aboutness and truthmaking – and to question the importance of ‘logical form’ in ontological enquiry; but then also to argue that we should not throw out the baby with the bathwater, and indeed keep the bathtub, the broadly Quinean paradigm of philosophical investigation, by reformulating the ‘criterion’ in voluntarist terms, by the conscious choice to adopt a certain interpretation of what I (and others) call ‘identity inferences’ of the form ‘ $Fa \wedge Ga \vdash (F \& G)a$ ’.

**Preview of the ‘seminar’:** For the seminar, I propose to have a discussion of the issues mentioned above in the light of recent defenses of (what its proponents call) ‘easy ontology’.

## Directions of fit

**Main claim:** the difference in direction, commonly assumed to mark an important difference between commitment and truthmaking, is less important than has been assumed.

The world and what we say about it are connected by two different relations, which differ in their “directions of fit” – aboutness goes from what we say to the world and truthmaking from the world to what we say: this gives us *oc*, spelt out as *oc’*, which makes it the converse of one – rather orthodox – view of truthmaking, *TME*:

*oc*: A truthbearer  $p$  is *ontologically committed* to an entity  $e$  iff  $e$  has to exist for  $p$  to be true.

*oc’*: A truthbearer  $p$  is *ontologically committed* to an entity  $e$  iff  $p$  entails “ $e$  exists”.

*TME*: An entity  $e$  is a *truthmaker* for a truthbearer  $p$  iff “ $e$  exists” entails  $p$ .

*TME*, however, is implausible: everything will make all necessary truths true; necessary existents will make true only necessary truths.

*oc’* and *TME* differ, moreover, in more than their ‘direction’: while *oc’* spells out an intimate connection between logical form and ontological commitment, which is built into the very notions of the two (and is therefore analytic), *TME* is a substantive (synthetic) metaphysical thesis not about the logical form of sentences but about the grounds of their truth. *oc’*, but not *TME* may plausibly be taken to be stipulative: *oc’* is a *criterion* of ontological commitment in the sense that it offers us a way of making ontological commitments explicit. *TME*, on the other hand, is an *analysis* of an independently existing phenomenon (truth) and true or false depending on whether it correctly characterises it.

Another difference concerns what we may call ‘specificity’. The ontological commitments of a truthbearer are (or at least: should be) endorsed by someone who holds the truthbearer true at least to the degree to which we are committed to logical consequences of what we are saying. Nothing comparable holds of truthmaking: we are at best committed to *there being* truthmakers for what we hold true, not to the obtaining of any particular truthmaking ties. I propose to characterise this difference using the (slightly technical) notion of **aboutness**: what makes our sentences true may be some definite particular of which we have no knowledge or grasp, whereas we can only commit ourselves to things by sentences which are *about* those things.<sup>1</sup>

In this (technical?) sense, “aboutness” is much more demanding than “speaking-of”: if I say “all bankers should go to jail” I may be taken to be speaking of my friend Fred, who unbeknownst to me is a banker, but it seems a stretch to say that the sentence (or my utterance of it) is about him.<sup>2</sup>

1. This notion of aboutness, in particular, is very different from the notion of ‘subject matter’, metaphysically characterised e.g. by Frank Jackson in terms of supervenience and by David Lewis in terms of ... truthmaking (!). Lewis (2001) characterises “ $p$  is (entirely) about  $T$ ” as “for any two possible worlds indistinguishable with respect to  $T$ , “ $p$ ” is either true or false in both of them”.

2. Perhaps ‘speaking-of’ and ‘being-about’ are even two-way independent: if I say “the man over there drinking vodka is a banker”, having in mind Fred who drinks water, perhaps I may be held to be speaking about Fred, but not obviously of him (at least if there is another man over there drinking vodka).

Can such a notion of aboutness be independently and directly characterised? The jury is still out on this. Different theorists have approached it from different directions:

In Quine, aboutness is closely linked to referential opacity or transparency of contexts, notions defined in terms of whether the context in question licenses the inter-substitutivity of coreferential terms *salva veritate*:

...the basis of the principle of substitutivity appears quite solid; whatever can be said about the person Cicero [...] should be equally true when said about the person Tully [...] this being the same person. (Quine 1953d: 17)

Very much in a Fregean vein, Quine also draws the converse implication: if coreferential terms are not substitutable *salva veritate*, then their occurrence within the sentence is not – or “not squarely” – about their usual referent:

If we assert [“Tom believes that Tully wrote the *Ars Magna*”] on the strength of Tom’s confusion of Tully with Lully, and in full appreciation of Tom’s appreciation that Cicero did not write the *Ars Magna*, then we are not giving the term “Tully” purely referential occurrence in our sentence “Tom believes that Tully wrote the *Ars Magna*”; our sentence is not squarely about Tully. If it were, it would have to be true of Cicero, who *is* Tully. (Quine 1957: 18)

Other theorists have connected aboutness to the explanandum of theories of so-called ‘direct reference’, i.e. the supposed semantic difference between definite descriptions (and abbreviations of such) and ‘proper’ names:

Russell’s analysis of statements containing definite descriptions and, by extension, ordinary proper names, shows, he believed, that such statements are not really *about*, do not really *mention*, the denotation of the description or the referent of the name. (Donnellan 1974: 223)

A third possibility would be to give a modal, or at least ‘model-theoretic’ characterisation, perhaps in terms of permutation invariance: a truthbearer is about *e* (perhaps even if, but certainly) only if its truth rules out certain permutations of qualitative properties of *e*. Whether there are such (non-trivial) permutations is one of the issues discussed under the (much abused) label “haecceitism”.

The champion of ontological commitment is undoubtedly Quine: if “to be *assumed as* an entity is [...] to be reckoned as the value of a variable” (Quine 1948: 13), a theory or body of (putative) truths *T* commits us to those entities that are assumed to be in the range of variables to make *T* true.<sup>3</sup> According to Quine, it is the presumed truth of the sentences we use that puts us under ontological obligations and it is our use of first-order quantification in a semi-formalised canonical idiom which makes these obligations explicit.

According to Quine’s criterion, what a sentence commits us to is determined by what its variables range over and hence by how it is formalised. The connection between formalisation and commitment, however, may be in one of two directions:

- (i) Formalisation uncovers ontological commitment: the commitment of a sentence is determined by its logical form.
- (ii) Formalisation is constrained by ontological commitment: sentences have an ‘ontological form’ that a correct formalisation has to confirm.

Quine privileges the first direction, while I think that the second is more important: we cannot simply read of our commitments from the logical forms of the sentences we accept – rather, we are guided in our formalisation efforts themselves by a sense of which commitments are acceptable. But how, one may ask, is it even *possible* that formalisation is constrained by, rather than constrains, ontological commitment? After all we do not, one may think, have an independent grasp of ‘ontological’, as opposed to logical form. But perhaps we do.

There are many inferences we accept as valid but are uncertain about how to formalise. Consider the following inferences:

- (1)  $\frac{\text{Maria is a vixen.}}{\text{Maria is female.}}$
- (2)  $\frac{\text{Sam is slowly buttering his toast.}}{\text{Sam is buttering his toast.}}$
- (3)  $\frac{\text{He ought to } F.}{\text{He is able to } F.}$
- (4)  $\frac{\text{Spaghetti can be cooked 'alla Bolognese'.  
Spaghetti can be cooked 'alla Carbonara'.  
Spaghetti can be cooked 'alla Putanesca'.}{\exists x, y, z(x, y \text{ and } z \text{ are ways of cooking spaghetti and } x \neq y \wedge x \neq z \wedge y \neq z)}$

We recognise these inferences as valid even though (and hence independently of the fact) that we do not want to commit ourselves to some particular formalisation of the sentences they contain. (4) will be considered valid by competent speakers, at least if the conclusion is read as “There are (at least) three different ways of cooking spaghetti”.<sup>4</sup> It is difficult to see why acceptance of the validity of the existential generalisation should commit us to acknowledge an implicit third argument place in “cooks”.

3. Quine was certainly not alone, and perhaps not the first, to focus on variables as the locus of ontological commitment. According to Hintikka (1966: 40) (cited after the reprint), essentially the same criterion was put forward by Kazimierz Ajdukiewicz in his dissertation, published as Ajdukiewicz (1921).

4. It might be replied that it is *only* under this reading that competent speaker accept (4). They would not accept it, the thought goes, if it were read in some other way making the alleged ontological commitment more explicit, e.g. as “Ways of cooking spaghetti exist and there are three of them and they are all different”. But our disinclination (if any) to accept the latter sentence could be just due to the fact that it is bizarre and that competent speakers are reluctant to accept sentences that are barely grammatical. A proponent of the view that ontological commitment can be read off from logical form is committed to some kind of correspondence between quantifier phrases logical notation and more familiar sounding sentences in natural language. It is not clear how such a correspondence could treat the conclusion of (4) differently from, e.g., “ $\exists x, y, z(x, y \text{ and } z \text{ are zebras and } x \neq y \wedge x \neq z \wedge y \neq z)$ ” without presupposing that ways of cooking spaghetti do not (and are not taken to) exist.

But not only do we accept inferences as valid that we do not know how to formalise, we also formalise sentences without endorsing inferences that are licensed by those formalisations. Consider e.g.

$$(5) \quad \frac{\text{Britney lost her virginity}}{\exists x(\text{Britney lost } x)}$$

(5) will undoubtedly be judged unacceptable by (a majority of) competent speakers of English: we may truly say that Britney lost her virginity without thereby committing us to virginites. Does it follow from this that “Britney lost her virginity” does not have the logical form  $R(a, b)$ , even though “her virginity” therein apparently functions as a singular term? Someone defending such a claim could point out that “Britney lost her virginity” is equivalent to (necessarily implies and is implied by or even is synonymous with) a sentence which predicates a monadic property of Britney and that it is this monadic predication which makes apparent the ‘true’ logical form of “Britney lost her virginity”. Such an argument, however, comes with substantial presuppositions:

- (i) necessarily equivalent or, at least, synonymous sentences share their logical form; and
- (ii) logical form is what licenses existential generalisation.

Both these assumptions may be challenged. The first (i) seems especially tenuous. Even if “Sam is a vixen” and “Sam is a female fox” are synonymous, does it follow that “Sam is a vixen” has the logical form  $Fa \wedge Ga$ , which would make “Sam is a vixen” imply “There is at least one female”?<sup>5</sup> Necessary truth-preservation, to be sure, but implication in virtue of *logical form*?

The reluctance of ordinary speakers to read the conclusion of (4) as “Ways of cooking spaghetti exist and there are three of them and they are all different” may also be explained in a different way. Whenever several readings of a natural language sentence are available, Quine counsels us, we should adopt the ontologically least committal. This is expressed by the modal auxiliary in many statements of the criterion:

...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 affirmation made in the theory be true ... (Quine 1948: 13–14, my emphasis);  
 ...entities of a given sort are assumed by a theory if and only if some of them *must* be counted among the values of the variables in order that the statements affirmed in the theory be true ... (Quine 1953b: 103, my emphasis, cf. also 108);  
 To show that some given object is required in a theory, what we have to show is no more nor less than that that object is *required*, for the truth of the theory, to be among the values over which the bound variables range. (Quine 1969a: 94, my emphasis)

The “must” in these formulation, in my view, should not be interpreted as an alethic modality (alethic modalities are, after all, repudiated by Quine),<sup>6</sup> but as a meta-theoretic statement: Whenever a claim may be evaluated as true over different domains, we should interpret it as committed to only those entities common to all these domains.

While some flexibility in setting the ontological price of a statement is clearly needed, quantification over available domains, even when relegated into the meta-theory, brings additional problems: while unions (existential quantifiers, see below) are problematic, intersection is so too. For one thing, it greatly reduces the ontological commitment of many sentences. Quine (1969a: 96) himself has pointed out that some sentence may be true if interpreted with variables ranging over different, but non-overlapping universes, as, e.g., “ $\exists x(x \text{ is a dog})$ ” in universes containing only collies and spaniels respectively. If we take the specific objects required to be those common to all those universes, “ $\exists x(x \text{ is a dog})$ ” will not commit us to anything.

Even after half a century of intense discussion, it is still not clear what exactly the logical form of Quine’s criterion is supposed to be. Various candidates have been proposed, but they all seem to face difficulties. Suppose, for instance that we try to explicate:

$$(6) \quad T \text{ is ontologically committed to (for short: assumes) } a.$$

by

$$(7) \quad T \vdash \exists x(x = a)$$

Does “ $a$ ” in (6) occur referentially or is the context opaque? In (7), the singular term  $a$  occurs in what Quine calls ‘purely referential’ (i.e. quantifiable) position. If (6) is analysed by (7), therefore, we are licensed to infer from it:

$$(8) \quad \exists x(T \text{ assumes } x)$$

While (8) is unproblematic for those who *share* the commitment to  $a$  expressed in (6), it shows that (7) is inadequate as an analysis of the *criterion* of ontological commitment. For to apply (6) as a criterion, we should be able to say that *others* are committed to entities that do not exist, which, in the presence of (8) we cannot.

If “ $a$ ” in (6) does not occur referentially, however, it seems impossible to construe ontological commitment as a relation between a theory and a thing (Scheffler & Chomsky 1958). If we construe it, with Chomsky and Scheffler, as a relation between a theory and a class, however, we get the unacceptable result that “ $\exists x(x = \text{Pegasus})$ ” and “ $\exists x(x = \text{Count Dracula})$ ” have the same ontological commitment, namely  $\emptyset$  (Jubien 1972: 383), and, in addition, that every statement commits its proponent to the existence of classes.

The formalisation of Quine’s criterion as (7) faces another, apparently fatal, objection: Terence Parsons (1967) has shown that any account of ontological commitment as a sentential functor  $\mathcal{C}(\phi)$  from expressions to classes which obeys the following principle:

$$(9) \quad \text{If } \phi \vdash \psi, \text{ then } \mathcal{C}(\psi) \subset \mathcal{C}(\phi).$$

5. Quine (1953b: 114) himself claims that although “ $\exists x(x \text{ is a dog} \wedge x \text{ is white})$ ” and “ $\exists x(x \in \text{dogkind} \wedge x \in \text{class of white things})$ ” are truth-functionally equivalent, they differ in ontological commitment.

6. Based on the presence of the modal auxiliaries, Cartwright (1954: 319) and Chihara (1968: 32) have – falsely, in my view – interpreted Quine’s criterion as intensional.

and either accepts  $\mathcal{C}(\exists x\phi) \subset \{x|\phi\}$  or  $\{x|\phi\} \subset \mathcal{C}(\exists x\phi)$  – i.e. that the ontological commitments of “there is a  $\phi$ ” and the class of  $\phi$  overlap – entails that for any two atomic predicates  $P$  and  $Q$ ,  $\mathcal{C}(\exists x(Px)) = \mathcal{C}(\exists x(Qx))$ . Parsons (1967: 450) concludes that “the only recourse for a meaningful notion of ontological commitment is to move into the domain of a theory of meaning”.

It is not clear how we are to do this, however. If we want to keep ontological commitment as a relation between theories and *things*, “ $T$  assumes  $x$ ” should not be opaque in “ $x$ ”. This rules out a meta-linguistic rendering, as Jubien (1972: 385) has noted:

(10)  $T$  uses “ $a$ ” referentially.

“Being used referentially” is a property not of things, but of their names. When I criticise you for believing in the existence of Pegasus, I am not criticising you for having a false attitude towards (what you take to be) a thing, nor of wrongly using a name – I am chiding you for believing in the existence of something I think does not exist. This was the original puzzle in Quine (1948) and we seem to have made little progress.

Quine’s original advice was to replace talking about Pegasus by talk about the unique Pegasizer, where ‘pegasizing’ is the “*ex hypothesi* unanalyzable, irreducible attribute of being Pegasus” (Quine 1948: 8). Ontological disputes then becomes meaningful: the believer in Pegasus is quantifying over pegasizing things while the disbeliever does not. After elimination of the definite description “the Pegasizer” by Russell’s method, our dispute becomes one solely about the domain of quantification. If such reformulations were generally available, some progress would have been made. But are they?

How are we to understand the new predicate “ $x$  pegasizes”? Suppose the attribute it ascribes is a general feature of things, i.e. could, as far as the semantics of “ $x$  pegasizes” is concerned, be exemplified by more than one thing. We then have the disturbing consequence that “Tom is Pegasus” and “Tom pegasizes” are not synonymous, for the former, but not the latter, implies that there is at most one pegasizing thing. It therefore must be a *semantic feature* of the predicate that it is true of at most one thing.<sup>7</sup> Such predicates, however, are normally semantically complex: paradigmatically, they are composed of the identity predicate and a singular term. If, however, we are to understand Quine’s unanalyzable attribute as

(11)  $x$  pegasizes  $\iff x = \text{Pegasus}$

our understanding of the predicate is derived from our prior grasp of the semantics of the proper name. That “ $x$  pegasizes” is true of at most one object must be stipulated by some meaning postulate for it. The proposed method of Quining away empty terms bites itself in the tail:

...the employment of terms like ‘pegasizes’, with a “guarantee” (whatever that might mean) that they are truly predicable of one and only one thing (or of nothing), does nothing more than introduce “proper names” at the predicate level. (Hochberg 1957: 553)

Hochberg’s point, I think, is well taken: nothing in what Quine says gives us any clue to how we may understand his unanalyzable predicates (especially those that are, necessarily or contingently, unsatisfied).<sup>8</sup> The ‘artificially seeming device’ allows us to say “Pegasus does not exist” without committing ourselves to Pegasus’ existence. But this is just a consequence of the general trade-off between ontology and ideology: if, as Quine (1960a) himself has shown, we can construct languages that do away with variables altogether, we should not conclude from this that sentences couched in such variable-free languages do not incur ontological commitments.

These difficulties, I think, show that Quine’s dismissal of names as paradigmatic bearers of ontological commitment in favour of variables was premature. He was right, however, that *some kind* of semantic ascent is asked for. To say that, contrary to what proponents of  $T$  believe, Pegasus does not exist, is to say that even *their* domain of quantification does not contain any such thing than Pegasus. This is not to say that Pegasus lies *outside* their domain of quantification, but to say that Pegasus is not ‘where’ it would be if it existed (namely in their domain of quantification). We can do this by assuming, for *reductio*, that “Pegasus” refers, devise a predicate true of its putative referent, if of anything, and then use that predicate to say that it is not true of anything – or we can, more simply, just say that “Pegasus”, as used by proponents of  $T$ , does not refer. Are we then back to (10)? Not quite.

Even if the whole of the ontological weight of a theory, according to Quine (1948: 12), is carried by its quantificational apparatus and proper names can be dispensed with in favour of descriptions, there is a connection between ontological commitment and what we might call “genuine proper names” given by the rule of existential generalisation:

(EG)  $\frac{Fa}{\exists x(Fx)}$  EG

We should resist the suggestion to treat (EG) as “anomalous” (Quine 1953c: 167) and take seriously Quine’s idea that it is “the logical content of the idea that a given occurrence is referential” (Quine 1953d: 146). The validity of (EG) presupposes that “ $a$ ” names something (Quine 1953c: 163) and that it occurs referentially (Quine 1953d: 145); whenever a theory accepts “ $\exists x(x = a)$ ”, it is said to use “ $a$ ” as naming (Quine 1969a: 94).

The ontological commitments of a theory and its range of referring expressions are thus seen to be interchangeable, at least if the latter are understood as including definite descriptions used referentially.<sup>9</sup> As Church (1958: 1009) has remarked, we may therefore turn matters

7. Cf.: “To be able to use definite descriptions instead of proper names one would have to know that the uniqueness clause of the description is fulfilled.” (Hochberg 1957: 551)

8. This holds true even if we – as Fara (2011) urges us to do – take as seriously as we can Quine’s admonition to construe “...is-Pegasus” as a monadic, non-relational predicate. Our lack of understanding of “pegasizes” also prevents us from asserting many other apparently true statements apparently about Pegasus that are not implied by “Pegasus does not exist”. “Pegasus is a fictional entity”, e.g., comes out false under the contextual definition à la (11), but seems true. Marcus (1972: 242) goes as far as calling “Necessarily, Pegasus is Pegasus, and Pegasus is not a fish” an “obvious truth”. They would not be obvious if our only ground to believe the second conjunct were that Pegasus does not exist.

9. The latter proviso is needed to account for sentences like “There are unspecifiable real numbers” (Quine 1969a: 95).

round and say that “a symbol is genuinely a proper name only if it admits the existential inference typified by the inference from  $[Fa]$  to  $[\exists x(Fx)]$ ”:<sup>10</sup>

(12)  $T$  is deductively closed under (EG).

(12), because it only mentions, but does not use “ $a$ ”, has the obvious advantage that it does not commit its proponent to “ $a$ ” having a referent.

## Domains of quantification

**Main claim:** logical form does not give us a grasp of our domain of quantification – rather it is the domain of quantification that constrains admissible logical forms for our sentences (‘ontological form’).

Gottlob Frege, the founder of quantificational logic, says in the section on generality of his *Begriffsschrift* that in “ $Fa$ ”, both “ $a$ ” and “ $F$ ” may be regarded as variable, giving rise to different existential generalisations “ $Fa \vdash \exists x(Fx)$ ” and “ $Fa \vdash \exists \phi(\phi a)$ ” (Frege 1879: §11). Let us consider:

(13)  $\exists \phi(\phi(\text{Socrates}))$

How are we to read (13)? “Socrates is something” and “There is something such that Socrates is it” smuggle in the copula. To see the problem with this, suppose we specialise (13) to the property of being a philosopher and get “Socrates is a philosopher” – may we then say, with Frege, that we replaced  $\phi$  with the referent of “...is a philosopher” in the same way than we specialised “ $\exists x(x \text{ is a philosopher})$ ” to the referent of “Socrates”? Unfortunately, we cannot. For what is it that is said to be identical to  $\phi$  in our shortening “ $\exists \phi(\phi(\text{Socrates}) \wedge \phi = F)$ ” to “ $F(\text{Socrates})$ ”? Dummett (1981: 217) thinks that it is the referent of “what “ $x$  is a philosopher” refers to” and he abbreviates “There is something that “ $x$  is a philosopher” refers to” to “There is such a thing as being a philosopher” which “is quite impossible to deny” (Dummett 1981: 218). The problem with this, as Russinoff (1992: 75) has pointed out, is that “being a philosopher” cannot be substituted by “ $x$  is a philosopher” in the latter sentence (“There is such a thing as  $x$  is a philosopher” is ungrammatical) – so how can these two expressions be co-referential? Even if there are singular terms for properties, they cannot replace  $\phi$  in (13) to make a full sentence – a copula has to be provided. The smuggling in of the copula, however, should make us doubt that the quantification in (13) commits us to an ontology of ways Socrates is.

Some quantifiers then, in particular those Friederike Moltmann (2003) calls “nominalizing quantifiers”, do not range over expressions they could replace, nor over their semantic values.<sup>11</sup> They are therefore neither amenable to an ‘objectualist’ nor a ‘substitutionist’ treatment. They can replace predicative, clausal and adverbial complements (and hence be propredicative, prosentential and proadverbial):

(14)  $\frac{\text{John is wise.}}{\text{John is something admirable.}}$

(15)  $\frac{\text{John believes that Mary married Bill.}}{\text{John believes something funny.}}$

(16)  $\frac{\text{John and Mary behaved awkwardly.}}{\text{John and Mary behaved the same way.}}$

What do these quantifiers range over? No suitable candidates are in sight.

What now of the other connection, from existence to quantification? Can at least this direction of Quine’s criterion be salvaged, i.e. are all existents quantifiable over? Plausibly, this depends on what Crispin Wright (1998: 73) has called the ‘Reference Principle’, namely that “co-referential expressions should be intersubstitutable *salva veritate*, at least in extensional contexts, and intersubstitutable *salva congruitate* in all”. If substitution of coreferential singular terms does not always preserve the well-formedness of the sentences in which they occur (cf the examples given in the discussion session), then the alleged link between existence and quantification is severed: if the singular terms that cannot be replaced *salva congruitate* by coreferential expressions succeed in denoting something (as I think they do), their referents cannot be the values of bound variables replacing their occurrences

Not only paraphrase, however, can help us to limit our ontological commitments; so too can wholesale re-interpretation, not of individual sentences of the theory, but of the alleged subject-matter of the theory as a whole. This is what the formalists do when they “keep[...] classical mathematics as a play of insignificant notations” (Quine 1948: 15), regarding mathematics as a (useful) myth without thereby coming – in their, as in Quine’s, view – under the obligation to retranslate it in any way. But not just formalist mathematicians may find their ontological bills easy to settle. For Quine re-translatability begins at home.

This problem most prominently arises with the indeterminacy of meaning: because empirical evidence only provides information about the assent patterns of *sentences*, there is no fact of the matter in Quine’s behaviouristic semantics which one of the group of sentences with the same stimulus-meaning is the correct translation of a the hitherto un-parsed foreign sentence. The radical interpreter’s decision “what expressions to treat as referring to objects, and, within limits, what sorts of objects to treat them as referring to” (Quine 1957: 3) is not just underdetermined, but *unconstrained* by all possible empirical evidence:

<sup>10</sup>. We do not have to make the further and stronger claim that ontological commitment to entities  $x$  such that  $M$  is *only* given by the assertion of “ $\exists xM$ ” (Church 1958: 1014).

<sup>11</sup>. This is argued for in Moltmann (2003: 452–454). Her main arguments are that nominalizing quantifiers can relate two argument positions which require expressions of different syntactic categories as in “John claimed something Mary never thought about” (ruling out substitutionalism) or which require expressions with different semantic values as in “John became everything Mary had wished for” (ruling out objectualism).

English general and singular terms, identity, quantification, and the whole bag of ontological tricks may be correlated with elements of the native language in any of various mutually incompatible ways, each compatible with all possible linguistic data, and none preferable to another save as favored by a rationalization of the *native* language that is simple and natural to us. (Quine 1957: 4-5)

But even if we impute our referential apparatus on the natives' language, i.e. assume that their ontological categories are roughly the same as ours, facing our 'home language' 'at face value', we still have only an essentially structuralist account of their ontology, leaving open the possibility of unintended, 'Pythagoreanistic' interpretation (Quine 1976b).

Russell clearly recognised the availability of 'Pythagoreanistic' interpretation as a threat to his structuralist theory of scientific knowledge (Russell 1927: 4)<sup>12</sup> and Quine could – and does – adopt Russell's solution to it. Abstract models are ruled out on the ground that they are not appropriately connected with our observations. We 'acquiesce' in our home language, as it were, take it 'at face-value' and determine the ontological commitments of our theories relative to it. Here, I would like to point out some consequences of this move.

If we identify our knowledge of the world with knowledge of only its structure (in so far it is describable by purely general quantified Ramsey-sentences) – if "of the external world we know its structure and nothing more" (Newman 1928: 142)–, then our theories admit of different, incompatible but equally 'good' interpretations *even if the domain of their interpretation is fixed!* The problem arises, as Newman realises, from Russell's claim that we know the structure  $W$  of the world in virtue of a relation  $R$  about which we only know that exists – the problem then is that "any collection of things can be organised so as to have the structure  $W$ , provided there are the right number of them" (Newman 1928: 144):

Thus, on this [Russell's 1927] view, only cardinality questions are open to discovery! Every other claim about the world that can be known at all can be known a priori as a logical consequence of the existence of a set of  $\alpha$ -many objects. (Demopoulos & Friedman 1985: 627)<sup>13</sup>

This is a devastating result, by anyone's standards. Proxy-functions do not only allow for the exchange of the whole domain, but can do away with ontological commitment even within one fixed domain. In a similar vein, Putnam (1977, 1978, 1980, 1981) has argued that the availability of reinterpretations of even our best theories makes the very notion of an "absolutely mind-independent reality" incoherent. This is so, he argues, because it makes no sense to suppose that a theory  $T_1$  "that is "ideal" from the point of view of operational utility, inner beauty and elegance, "plausibility," "simplicity," "conservatism," etc." (Putnam 1977: 485) might be false. From this, he concludes that so-called "metaphysical realism", a position committed to such a possibility, is itself incoherent. A compact version of his argument with respect to an ideal theory  $T_1$  runs as follows:

Pick a model for  $T_1$ ,  $M$ . Relative to the interpretation of 'reference' for  $L$  that yields  $M$ ,  $T_1$  must come out true. How could it not then *be* true? This interpretation must meet all operational (and theoretical) constraints on reference because  $T_1$  is ideal. There can be no further constraints that would rule out  $M$  as the "intended" model. So,  $T_1$  is true in any "intended" model and so must *be* true. The idea that  $T_1$  might be false is unintelligible. (Devitt 1983: 297–298)

Putnam relies on the model-theoretic proof that given a consistent theory  $T_1$  demanding a universe of cardinality  $c$ , and given any set  $s_M$  of cardinality  $c$ , there exists a model of the theory  $M$  the universe of which is  $s_M$ . Because  $M$  and the intended model are in one-to-one correspondence, we can extend the interpretation of  $T_1$  in its intended model to an interpretation of  $T$  in  $M$ . Under this correspondence between the language of the theory and  $s_M$ , he argues, the theory is true of  $s_M$  as a matter of logic – no matter what  $s_M$  is!<sup>14</sup>

We must come up with additional constraints on admissible interpretations. Putnam considers an appeal to a causal theory of reference by the metaphysical realist: but this would be, he answers, just the addition of more theory to  $T_1$  (Putnam 1978: 126) and therefore not help, as long as "reference" is not "glued to one definite relation with metaphysical glue" (Putnam 1980: 477). I think that it is this 'just more theory' reply to realists that may plausibly be rejected:

Constraint  $C$  is to be imposed by accepting  $C$ -theory, according to Putnam. But  $C$ -theory is just more theory, more grist for the mill; and more theory will go the way of all theory. To which I reply:  $C$  is *not* to be imposed just by accepting  $C$ -theory. That is a misunderstanding of what  $C$  is. The constraint is *not* that an intended interpretation must somehow make our account of  $C$  come out true. The constraint is that an intended interpretation must conform to  $C$  itself. (Lewis 1984: 62)

This, I think, is the right answer to Putnam's paradox, but more must be said. I will discuss two apparently different strategies and show that they come to the same thing.

The first answer is to bite the bullet face-on: the reference of "reference" falls indeed within the scope of a theory of reference and belief in it "arise[s] out of a causal interaction between the believers and a reality independent of those beliefs" – but the fact that the determinacy-of-reference question can be raised at any stage does not mean that it is never answered (Devitt 1983: 298):

---

12. For the link between Russell's theory and Newman's critique and Putnam's model-theoretic argument I am indebted to Lewis (1984: 61,fn. 9) and Demopoulos & Friedman (1985).

13. They add: "(The relevant theorem from set theory or second-order logic is the proposition that every set  $A$  determines a full structure, that is, one which contains every relation (in extension) of every arity on  $A$ ; such a structure forms the basis for a (standard) model for the language of second- (or higher-) order logic.)" (Demopoulos & Friedman 1985: 628)

14. As Demopoulos & Friedman (1985: 634) have noted, this is a weakening of Newman's result: while Newman shows that fixing the domain and models up to isomorphism does not fix the intended reference of the  $T_1$ -vocabulary, Putnam shows that it is not fixed by fixing the domain and models up to elementary equivalence (a weaker relation strictly implied by isomorphism).

Putnam's anti-realist argument depends on there being no answer to the question about what determines reference for  $T_1$ . Using a theory of reference there is an answer: reference is determined by causal relations of a certain sort. That answer works for 'causally related' just as it does for 'cat'. (Devitt 1983: 299)

This means that we go back to the answer Russell gave Newman in a letter dated April 24th, 1928:

Is it of course obvious [...] that the only effective assertion about the physical world involved in saying that it is susceptible to such and such a structure is an assertion about its cardinal number. [...] It was quite clear to me, as I read your article, that I had not really intended to say what in fact I did say, that *nothing* is known about the physical world except its structure. I had always assumed spacio-temporal continuity with the world of percepts, that is to say, I had assumed that there might be co-punctuality between percepts and non-percepts... [...] And co-punctuality I regarded as a relation which might exist among percepts and is itself perceptible. (Russell 1968: 176), quoted after Demopoulos & Friedman (1985: 631–632)

Though he never says so in his published works, Russell here seems prepared to assume acquaintance "with a cross category notion such as spacio-temporal contiguity or causality" (cf. Demopoulos & Friedman 1985: 632, who criticise this solution as "quite artificial").<sup>15</sup>

If Russell counts the causal connections in virtue of which we have knowledge of the external world as falling under his notion of "acquaintance", then he introduces for relations, as Newman urged, an analogue of the important/unimportant distinction for domains. Those relations will qualify as important knowledge which is not purely structural. In this sense, the Russell-Newman reaction to reinterpretability comes to essentially the same as the second solution to Putnam's paradox, proposed (on behalf of the realist) by Merrill (1980) and endorsed by Lewis (1984):

...realists typically hold that not only are there objectively existing entities (both observable and unobservable) in the world, but also that *these entities bear to one another certain objective relations*. And according to this latter view *the world* must be represented *not* simply as a *set*, but as a set *together with* a class of relations among the members of that set. To describe *the world* is to describe the entities (or kinds of entities) in it *and their relations to one another*. (Merrill 1980: 72)

The 'ready-made world' dear to metaphysical realists must be understood not just as consisting of mind- and language-independent entities, but as having a *structure* that is independent of language:

...if, as the realist surely must hold, the real world is a *structured* domain, then we are not free to ignore its intrinsic structuring in playing our model-theoretic tricks. (Merrill 1980: 74)

The intrinsic structure of the world, according to Lewis' proposal building on Armstrong's scientific realism about universals, comes from more or less natural properties, the 'real joints' in nature:

Among all the countless things and classes that there are, most are miscellaneous, gerrymandered, ill-demarcated. Only an elite minority are carved at the joints, so that their boundaries are established by objective sameness and difference in nature. Only these elite things and classes are eligible to serve as referents. (Lewis 1984: 65)

Just structure, however, is not enough: the structure of the world must be specific, the elite properties and relations sparse (Merrill 1980: 77). These elite properties and relations will then act as 'reference-magnets', constrain admissible interpretations and make for "an objective inegalitarianism of classifications" (Lewis 1984: 67).

How does this apply to Quine? In our search for the ontological commitments of a theory, the slogan "to be is to be the value of a variable" advises us to look at how the apparatus of quantification in the underlying logic (usually assumed to be first-order) interacts with the predicates of the theory. We are, in short, to look at the extensions of predicates and therefore at the domain of quantification. The addition of the modal auxiliary, in this context, makes our ontological commitments depend on our *choice* of a domain. Logical techniques of exchanging one domain for another and for permutating the domain while preserving the truth-values of sentences then become available – to devastating effect. Quine's reaction, to "acquiesce in the home language", comes to more than just privileging one referential apparatus over another: it amounts, in Ludovician terms, to an inegalitarianism of classifications.

The model-theoretic argument, its ancestors and variants show that Quine's original picture of ontological commitment is wrong in at least two respects: it gets the respective priority of reference and satisfaction wrong and it reverses the priority of reference and ontological commitment. The first point becomes apparent if we consider what it means to postulate 'objective inegalitarianism' as a constraint on admissible interpretations: it means that interpretation must link our predicates and singular terms with the extensions and referents that are most eligible over-all, taking into account possible trade-offs with the eligibility of other extensions or referents of other parts of the language (Lewis 1984: 66). There must be a correlation between things (and classes of things) in the world and the words we use to talk about them over and above the requirement that plugging the singular terms into some sentential function expressing some condition true of all and only the things in the class leaves us with a true sentence. In this sense, then, Putnam's paradox shows the need for *thick* notions of reference and satisfaction: for something to be the referent of "*a*", for example, it is not enough for "*x = a*" to be true of it; for something to satisfy "*Fx*", it is not enough for "*F*" to have an extension including it.

Putnam's paradox also shows that Quine's insistence on the domain of quantification puts the cart before the horse: *In advance* of a criterion of ontological commitment, the very idea of choosing a domain of quantification to interpret our predicates in and have our variables ranging over does not make any sense. We cannot pick the domain over which the variables in a theory range before we know

15. This interpretation is justified because Russell clearly expresses agreement with Newman's conclusion which the latter states as follows: "The conclusion that has been reached is to maintain the view that something besides their existence can be known about the unperceived parts of the world it is necessary to admit direct apprehension of what is meant by the statement that two unperceived events are *causally adjoined*, i.e., happen near each other, temporally and spatially, or overlap or do something of the sort." (Newman 1928: 148)

what its ontological commitments are and we cannot know what the commitments are before we have a domain in which to interpret the theory. Framed this way, I think, it can be easily seen that Lewis’ response to Putnam’s (and Quine’s) trivialisation problems is on the right track: we have to take the idea seriously that domains are not *chosen* by us, but *given* independently of and prior to our interpreting our theories. Putnam’s model-theoretic argument is based on the assumption that “*we* [i.e. our intentions] interpret our languages or nothing does” (Putnam 1980: 482) and can therefore be taken as *reductio* of that assumption. As Lewis (1984) argues, some referents, things and classes of things, are more eligible than others – intended interpretations of our language must maximise eligibility of referents overall (Lewis 1984: 65). Their eligibility does not depend on whether or how we conceive of them; it is itself a natural (and hence, highly eligible) feature of the things que quantify over.

The criterion that an interpretation ascribing commitment to eligible things is, other things being equal, better than one that does not, not only applies to what we take ourselves to be committed to, but also to how we ascribe commitment to others.<sup>16</sup> We should not, as Quine originally thought, *minimise* the ontological commitments of others *tout court*, but instead minimise their divergence from ours.

## Identity inferences

**Main claim:** what really commits us to the existence of entities our statements are about are the inferences we are prepared to draw about their behaviour, and in particular a special kind of ‘identity-inferences’ the validity of which presupposes constancy of reference

We recognise inferences as valid even though (and hence independently of the fact, if it is one) that we do not want to commit ourselves to some particular formalisation of the sentences they contain. We do recognise, however, that their validity depends on what the premisses and the conclusion are about: if I was talking about two different people called “Sam”, for example, the inference (2) from “Sam is slowly buttering his toast” to “Sam is buttering his toast” would be fallacious. We presuppose, in other words, that our uses of names, variables and anaphora are *coordinated*: that identical words and sequences of syntactically dependent words refer – *de jure*, not just *de facto* – to the same thing.

A genuine proper name is a linguistic item that allows us to pick out an object as our subject of predication. This picking out is unconditional, it is not mediated by any descriptive content. The important thing about proper names as opposed to other singular terms is thus that they pick out their referents *reliably*, that they are (comparatively) immune against reference shifts and do not for their reference depend on contingent facts. In metaphysical terms the stability of reference characteristic of proper names comes to this: in their case, the naming relation internal (supervenes on the relata), while it may be external for other singular terms. Proper names track an object and they do so (relatively) invariably. Coordination, in its simplest form, is presupposed whenever we fuse two predications into one:

$$(\mathbf{id}_a) \quad \frac{Fa \quad Ga}{(F \wedge G)a}$$

I will call an inference of the type  $(\mathbf{id}_a)$  an “*inference trading on identities*” or “*identity inference*” for short. The conclusion of  $(\mathbf{id}_a)$  is meant to rule out cases of conjunction introduction where the proper name type is used ambiguously, as in

$$(\mathbf{id}_{DL}) \quad \frac{\text{David Lewis is an American-born philosopher famous for his modal realism.} \quad \text{David Lewis is a Russian-born Canadian lawyer and politician.}}{\text{David Lewis is both a philosopher and a politician}}$$

$(\mathbf{id}_a)$  is valid only if we are able to rule out this kind of equivocation. Identity is crucially involved in the validity of  $(\mathbf{id}_a)$ , as can be seen from the fact that

$$(\mathbf{id}_{\text{something}'}) \quad \frac{\text{Something is red.} \quad \text{Something is round.}}{\text{Something is red and round.}}$$

is not a valid inference. To get validity (under an assignment), we need an additional premise as in

$$(\mathbf{id}_{x,y}) \quad \frac{x \text{ is red.} \quad y \text{ is round.} \quad x = y}{x \text{ is red and round.}}$$

In other words, the assignments of values to  $x$  and  $y$  respectively have to be *coordinated*. How coordination is achieved is a vast, and partly empirical question. Some typical examples may however be useful. A first class of cases involves intentions of co-reference: two utterances are coordinated in this way if one of them essentially involves an intention to use some syntactical-lexical form, some sound pattern or some other physical item *in the same way*, whatever it is, as the salient physical item is used in the other. In typical cases when I am about to use a word, it already exists as a sequence of tokens, ending in one particular token – an intention of co-reference will then determine whether another phonetical or graphical item is a token of the same word or not. In other cases, as in the case of pronouns and anaphora, coreferentiality is encoded in the syntax of the word in question and the structure of some stretch of discourse.

It has long seemed mysterious how our brains achieve such coordination and how it can be formally and conceptually modelled. Many of the different sub-problems discussed under the label of the ‘binding problem’ pertain to this kind of coordination.<sup>17</sup> In the philosophy

<sup>16</sup>. The principle of humanity, a generalisation of Davidson’s ‘principle of charity’, imputes “eligible content, where ineligibility consists in severe unnaturalness of the properties the subject supposedly believes or desires or intends himself to have” (Lewis 1983: 52).

<sup>17</sup>. For relatively accessible surveys of some more recent literature, cf. Treisman (1996), Roskies (1999) and von der Malsburg (1999). The term “binding problem” has been introduced by von der Malsburg (1981) to describe the problem in brain research of modelling connections between cells that are active at the same time and in the same circumstances, held to be a challenge to classical connectionist networks. It has (at least partly) motivated a lot of research in neuroscience, most famously the ‘40 hz’ theory of consciousness of Crick & Koch (1992). Recent neuro-psychological research seems to indicate that the conceptual capacities that underlie our competence with inferences like  $(\mathbf{id}_a)$  are comparatively basic and *sui generis*.



of perception, Clark (2000) has argued that so-called feature integration – perceiving something as being both  $F$  and  $G$ , where  $F$  and  $G$  are sensible properties registered in distinct parallel streams, as in the classical problem of Molyneux or in the ‘Many Properties’ problem of Jackson (1977: 65) – requires a referential apparatus and that perception should therefore not be modelled as simple ‘feature-extracting’, but as a process having propositional structure: features are unified by their locations (in visual space),<sup>18</sup> into so-called ‘proto-objects’ or ‘pre-attentive objects’ (Clark 2004) that are picked out by ‘visual indices’ (Pylyshyn 2001: 127). It is only with such *objects* that we get from mere conjunction to coordination (Clark 2000: sct. 2.5). With this, Quine would agree:

The conjunction [of ‘Lo, a pebble’ and ‘Lo, blue’] is fulfilled so long as the stimulation shows each of the component observation sentences to be fulfilled somewhere in the scene – thus a white pebble here, a blue flower over there. On the other hand the predication [‘The pebble is blue’] focuses the two fulfillments, requiring them to coincide or amply overlap. The blue must encompass the pebble. It may also extend beyond; the construction is not symmetric. (Quine 1990: 4)

According to Clark (2000), coordination is achieved by placing features in regions and re-identifying those, thereby tracking visual objects through visual space.<sup>19</sup> It is not important for our present purposes whether, as Clark (2000, 2004) believes, proto-objects *are* regions or whether, with Cohen (2004) and Matthen (2004), we take them to be (intentional!) *objects* located at regions.<sup>20</sup> Whatever it is an identification of, however, it is not yet identification of particulars, but only a prerequisite for it:

These ultimate facts [of feature-placing] do not contain particulars as constituents but they provide the basis for the conceptual step to particulars. The propositions stating them are not subject-predicate propositions, but they provide the basis for the step to subject-predicate propositions. (Strawson 1959: 218)

While different authors characterise the missing element differently (availability of a sortal is a popular candidate), there is clearly an additional step involved. This suggests that the mechanism captured by ( $\mathbf{id}_a$ ) is more basic than and prior to explicit recognition of particulars within an ontology, and hence more fundamental than ( $\mathbf{EG}$ ), which in turn can be considered a special case of it. Proper names for whatever feature-placing locations there may be are prior to their representation by the quantificational apparatus of the language:

The relevant part of Quine’s programme of paraphrase can most simply be summed up as follows. All *terms* other than the variables of quantification will be found, in canonical notation, to be general terms in predicative position. The *position* of singular terms is reserved for the quantifiers and the variables of quantification; and since quantifiers themselves do not count as *terms*, the only singular terms left are the variables of quantification. But, merely formal distinctions of grammar apart, how was the distinction between *singular terms* and *general terms in predicative position* explained? It was explained in terms of the contrasting roles of predication of the *definite* singular term and the general term in predicative position. This contrast of roles is our fundamental clue to all the theoretical notions employed. So our theoretical grasp of the nature of canonical notation rests upon our theoretical grasp of the identificatory function of singular terms. (Strawson 1961: 49)

Here is another way of making essentially the same point: when Quine says that it is the quantificational apparatus of a language that brings out most clearly what ontological commitments sentences of that language may have, the kind of quantification he has in mind is the *objectual* one. But what is objectual quantification? A familiar explication goes as follows (e.g. Schiffer 2003: 90): To say that the quantification in

(17)  $\exists x Fx$

is objectual is to say that (17) is true iff there exists some entity in the domain of quantification which satisfies the open sentence

(18)  $Fx$

To say that the quantification in (17) is non-objectual is to say that (17) may be true even though there does not exist an entity in the domain of quantification which satisfies (18).<sup>21</sup> This means that to say that we objectually quantify over an  $F$  (and hence, according to the criterion, are ontologically committed to  $F$ s) is to say that there is something in the domain that satisfies “ $Fx$ ”. To say the latter, however, we need a grip on the  $F$  in question. Without such a prior grip on it, we cannot even say that we intend to be ontologically committed to it.

18. More generally, the location is in a quality space, the number of dimensions of which corresponds to the number of independent dimensions of variation that the creature can sense among the stimuli that confront it in that modality.

19. Visual objects are mere intentional objects, treated by the visual system *as if* they were objects, regardless of whether they in fact are (Clark 2004). For a comparison of different object-concepts in developmental psychology and neuroscience, cf. Casati (2004) and references therein. For present purposes, it is important that the ‘proto-objects’ and their localisations may be parochial to just one modality (which would rule out the answer given to Molyneux’s question by Evans (1985b)). As Campbell (2005) points out against Evans (1982), even if deictic and intrinsic locations of objects (‘it is to the left of the ball’, ‘it is on Bill’s left’) are derived from basic egocentric locations, the computational procedure used may be modality-specific. The model is developed mostly for vision and touch and it is not clear whether the binding problem even arises with respect to, say, smells (cf. Clark 2000: 97).

20. I do not think that the arguments advanced by Cohen and Matthen are particularly compelling, however. Rather than talking of places, Clark would perhaps better have formulated his claim in terms of events. The obvious advantage of locations is that they require ontologically only what we have independent reasons to believe in – quality spaces for sense modalities –, while in the case of the ‘objects’ of Cohen (2004) “...there is no reason to expect visual objects to deliver the ontological goods” (Cohen 2004), i.e. no reason to assume them to be real.

21. On this construal, substitutional quantification is just one species of the non-objectual kind. The quantification in (17) is substitutional if (17) is true iff the open sentence (18) has a substitution instance

(19)  $Fa$

which is true and where “ $a$ ” is syntactically a singular term, but where it need not be required that “ $a$ ” refers.

Treating a position as referential is not something that a language does by itself – it is a matter of interpreting the language in a certain way. It is highly misleading, therefore, to say that “[w]hat entities there are, from the point of view of a given language, depends on what positions are accessible to variables in that language” (Quine 1951a: 68). The point of view of a language, if there is such a thing as all, is not a stance on what there is, nor on what there is taken to be. At most, it is a stance on what sentences are grammatical and under what constructions grammaticality is preserved. Ontological commitments depend on the point of view not of a language, but of those making claims using this language – making claims liable to figure as premises in ( $\mathbf{id}_a$ )-type inferences.

It is by acceptance of inferences like ( $\mathbf{id}_a$ ), that we commit ourselves to the existence of the referent of “ $a$ ”, hence to the existence of  $a$ .<sup>22</sup> This allows us to bring out what is right in the explication of ontological commitment as referential use (10). A commitment is something one acquires by doing something. If I tell you “You can count on my help”, I commit myself to doing certain things if asked to. It is not required that I actually do them (perhaps you do not need my help) – but I am under an obligation to be prepared to do them. In an entirely analogous way, an ontological commitment to  $a$  means that one is under an obligation to be prepared to draw inferences like (EG) and ( $\mathbf{id}_a$ ) given suitable premises. The commitment we incur is not to  $a$  specifically – it is to *something or other* that makes ( $\mathbf{id}_a$ ) come out sound – the entity we are committed to us is no further characterised than as the substratum, whatever it is, of this inference. In this sense, ontological commitment is unspecific.

Acceptance of an inference is intentional, a matter of the will, an attitude we take towards our statements, not something we express in these statements themselves: ontological commitment then depends on what attitude we adopt to our claims, and may be reduced just by changing that attitude, as it is in the case of the formalists *regarding* mathematics as “a play of insignificant notations”. If we want to understand their talk, however, we have to provide a translation scheme, and sometimes it may prove impossible to find a translation that both preserves their ontological commitment and makes their theory true. In such cases, Quine thinks, semantic ascent is demanded for.<sup>23</sup> But this means that from the ontological commitment of a theory, even a true one, we cannot just read off what there is – we first have to translate the theory into our own language, and then apply the criterion to the translation. But this presupposes that we already have a grip on what there is. Quine himself was quite clear about this:

...how are we to adjudicate among rival ontologies? Certainly the answer is not provided by the semantical formula “To be is to be the value of a variable”; this formula serves rather, conversely, in testing the conformity of a given remark or doctrine to a prior ontological standard. We look to bound variables in connection with ontology not in order to know what there is, but in order to know what a given remark or doctrine, ours or someone else’s, *says* there is; and this much is quite properly a problem involving language. But what there is is another question. (Quine 1948: 15–16)

Quine articulates here the crucial tension of his notion of ontological commitment: at the same time, it is said to depend on logical form and to be removable by paraphrase. The question then becomes: what constrains our assignments of logical forms?

Quine’s criterion and his insistence on the importance of the quantificational apparatus are thus best understood in voluntarist terms: it gives us a way of incurring ontological commitments, *if* we are prepared to do so. Persons, not sentences, commit themselves ontologically – sentences incur commitments only insofar as they are used to do so.<sup>24</sup> This also explains why the criterion is so closely bound up with the idea of the regimentation of ordinary language, i.e. the project of translating our vernacular into “canonical form”, i.e. an idiom (first-order logic) in which our commitments *may* (but do not have to be) unambiguously expressed.<sup>25</sup> And only such regimentation is what makes the criterion applicable:

In a loose way we often can speak of ontological presuppositions at the level of ordinary language, but this makes sense just in so far as we have in mind some likeliest, most obvious way of schematizing the discourse in question along quantificational lines. (Quine 1953b: 107)

If someone resists the invitation to regiment his language in canonical notation, we cannot impute an ontology to that person (Quine 1979: 161–162). Ontological commitment, therefore, has both a dispositional and a normative aspect to it – it is captured by the instances of (EG) and ( $\mathbf{id}_a$ ) we should be prepared to accept, given what we say *and* the standing presumption that we should express ourselves as clearly as possible (which, for Quine, means: in canonical notation).

To say that ontological commitment applies primarily to persons and only derivatively to what they assert helps us solve a problem raised by William Alston (1958): how can it be, Alston (1958) asks, that paraphrase may reduce one’s ontological commitment? Does not a difference in ontological commitment *ipso facto* show that the paraphrase is inadequate? The difficulty can be solved only if we accept, with Quine (1960b: 260), that “paraphrase” means restating in a different way what one wants to say (“accomplishing those same purposes through other channels”), thereby possibly changing what one in fact says. But this means that Quine’s criterion of ontological commitment is a package-deal: we do not just have to accept that to be is to be a value of a variable, but also that we best cast our theories in a language with objectual quantification. But even this is not enough, as we have seen in the case of the ‘spaghetti’-inference

22. It is, of course, difficult to say what the appropriate kind of acceptance of an inference consists in. Presumably, someone accepts the inference iff s/he is unreflectively disposed to treat it as truth-preserving, where this disposition is not derived from the acceptance of any more basic inference pattern. For a more careful discussion, cf. Boghossian (1994).

23. Subsequently, Quine adopted this meta-linguistic stance also toward advocates of substitutional quantification who denied that the existential quantifier has any existential import and that therefore the quantifiers cannot be read as “there is an object such that” and “all objects are such that” respectively (cf. e.g. Marcus (1972: 245) and Martin (1962: 527)).

24. This explains why Quine (1981: 175) can take the criterion to be trivial and even explicate it in clearly intentional terms: “What there are, according to a given theory in standard form, are all and only the objects that the variables of quantification are *meant in that theory* to take as values.” (Quine 1970: 89, my emphasis) To be sure, there are passages (Quine 1953b: e.g., 103), where Quine explicitly says that his notion of ontological commitment applies primarily to sentences and only derivatively to persons. But in these passages, I take him to mean only that unwelcome ontological commitments may be removed by paraphrase.

25. It is just beside the point, therefore, to point out that some uses of “there is” in ordinary language do not carry ontological weight (Azzouni 1998: 4). It is equally false to think, with Chihara (1968: 30), that Quine attempted to *replace* an ordinary locution like “presupposes as existing” with a semi-technical one (as arguable Tarski intended with the concept of truth).

(4): ontological commitment is not only relative to a logic, but to the specific choice of logical form for the sentence by which it is incurred. Ontological commitment and logical form, for Quine, are just two sides of the same coin.<sup>26</sup> The logical form of what you say not only determines which conclusions you should rationally accept, but the conclusions you accept may conversely constrain the ascription of some logical form to what you said. In some, or even most, cases the latter direction is prior, setting the adequacy conditions for a search for logical form that has only just begun. The same happens with ontological commitments: they do not just fall out of a theory of logical form, but are incurred by persons using sentences of the logical form of which they might have no clear conception. Here, I think, is another advantage of (**id<sub>a</sub>**) over (7) and even (**EG**): treating something as recurrent object of predications is not relative and indeed prior to its expression within a logical formalism.

More generally, entities to which we are ontologically committed may be conceived of as what we might call the ‘substrata’ of inferences we should be prepared to draw given what we say.<sup>27</sup> These inferences can be ‘material’, thereby divorcing ontological commitment from logical form. Consider e.g. the following inference:

(20) 
$$\frac{\text{Rupert kills Randolph.}}{\text{Rupert causes the death of Randolph.}}$$

There is no situation in which Rupert may kill Randolph without causing his death. So (20), being necessarily truth-preserving, seems at least *prima facie* valid. Should we therefrom conclude that “killing” means “causing someone’s death in some way” or that the logical form of “Rupert kills Randolph” involves a quantification over an event said to be a death?<sup>28</sup> I think that we are at least not obliged to say so. Let us say instead that the acceptance of (20) shows an ontological commitment to Randolph’s death on the part of someone who asserts the premise. And this, I think, was Quine’s basic point: that in at least some cases, *there are* ways of making explicit our ontological commitment, and that first-order quantification logic may allow us to do so by being used to regiment the way we speak. It is just not the only way, and sometimes not the best way.

I think that there are two main arguments not to settle exclusively, with Quine, on standard first-order predicate logic as our measure-stick for ontological commitment. The first reason is that there are many sentences, with which intuitively we may very well incur ontological commitments, which cannot be formalised in first-order predicate logic.<sup>29</sup> The whole debate about whether second-order logic is committed to sets would not make much sense if ontological commitment were only incurred by first-order quantifiers.

The second reason is that formalisation, even where it is possible, is not only non-unique but often also takes place at some given level of “granularity”. Consider the following example of Chihara (1968: 38):

(21)  $\exists x(x \text{ is a winged horse})$

It seems clear that an assertion of “Pegasus is a winged horse”, via (21), commits us to the existence of a winged horse. But does it commit us to the existence of wings, via (22)?

(22)  $\exists x, y(x \text{ is a horse} \wedge y \text{ is a wing} \wedge x \text{ has } y)$

I think it is too strong to call that consequence “counterintuitive” (Chihara 1968: 38). In some contexts “Pegasus is a winged horse” may very well commit us to wings, in others it does not. It all depends on whether “Pegasus is a winged horse” is taken to be an acceptable possible answer to the question whether there are wings (or better: whether there are entities having wings). In some contexts, the inference from “Pegasus exists” to “Wings exist” is valid, in others it is not. Shall we say that “Every boy loves a girl” “really” contains an occurrence of “if”? In a logic course yes, but not in a language acquisition study.

In determining the ontological commitments of a discourse, we cannot presuppose a prior translation into canonical notation. But can we at least go for the values of the variables? I think there are good reasons not to do so – reasons to prefer (**id<sub>a</sub>**) to (**EG**). The first reason is that, as I tried to argue above, Quine’s method for paraphrasing names away in favour of variables relies on a prior understanding of essentially ‘particularised’ attributes, which in turns seems derived from our competence with the proper names they are concocted from – we would better not kick away that ladder, especially if, second, our competence with proper names is underwritten by a more basic capacity of bundling perceptual features into the demonstrata of perceptual indices. Third, and more importantly, Quine’s method for paraphrasing names away would not suffice, even if it were successful. As Quine (1953b: 104) acknowledges and as Quine (1960a) has shown himself, Schönfinkel’s and Curry’s combinatoric logic affords the resources to paraphrase away not just proper names and descriptions, but also the variables themselves. This does not, by itself, show that the criterion is inapplicable, so long as a translation is available:<sup>30</sup>

Once we know the systematic method of translating back and forth between statements which use combinators and statements which use variables [...] there is no difficulty in devising an equivalent criterion of ontological commitment for

26. This seems to apply to inferences more generally: Rational agents do not have deductively closed belief sets and there are good reasons for that. Nevertheless, they are committed to the truth of what follows from what they say according to logical principles they accept (or have no reason not to accept). Confronted with an agent unwilling to draw conclusions we think follow from what he says, we are not *obliged* to blame him for irrationality or logical deviance, however, for we may also change our interpretation of what he said, trying to give another logical form to it such that the disputed conclusion does no longer follow.

27. This has been the use Davidson (1967) made of Quine’s criterion to argue for the existence of events.

28. Such a line would have a couple of problems: If *a* kills *b* entails that there is an event *x* which is *b*’s death, then does not the existence of this very event entail that *b* is dead, i.e. does not exist? But then we are under an obligation to give a logical form to “*a* kills *b*” that does not license existential generalisation (**EG**) and it is not clear what this could be.

29. Standard predicate logic is not expressive enough to handle many quantificational constructions of natural language, like “most”, “many” and “few”. “Most men are right-handed” e.g. cannot be represented adequately in a regimented language if its underlying logic is standard predicate logic cf. also Wiggins (1980: 326) and (Davies 1981: 124–127).

30. Van Fraassen’s injunction therefore seems premature: “Quine’s elimination of names can be carried further; we can also eliminate the variables! [...] Does the use of one language rather than another, when we know that each can be translated perfectly into the other, carry any ontological implications? Of course not. Quine’s programme, to deduce ontology from syntax, was just a mistake.” (van Fraassen 1991: 458, 459)

combinatory discourse. The entities presupposed by statements which use combinators turn out, under such reasoning, to be just the entities that must be reckoned as arguments or values of functions in order that the statements in question be true. (Quine 1953b: 104)

But how do we know that this translation is the right one, i.e. the one preserving ontological commitment? Not all logical equivalences, after all, preserve ontological commitment (cf. Quine 1953b: 114). Because we already know that the commitments of our discourse are the potential referents of our singular terms (“the arguments and values of functions”), had they not already been paraphrased away. Names, especially as used in inferences trading on the identities of their referent, are thus acceptable as clues to someone’s ontology. What counts as a name, however, is not something to be read off from logical form; it should ideally be reflected in the logical form, but often is not.

As things stand, we simply have to acknowledge that we are not (yet, perhaps) in a position to give to all inferences we accept a logical form which makes them valid. This does not mean, however, that no such inferences commit us ontologically nor that they do not commit us via **(EG)**. Ontological commitment does not fall out of a theory of logical form, it guides and constrains it as well.<sup>31</sup> Ontological commitment shares this feature with truthmaker theory: rather than falling out of a theory of logical form which somehow allows us to see through a proposition’s varied clothes, our judgements of ontological commitment constrains what logical forms we find acceptable. A theory of logical form might explain why we find (4), in contrast to (5), acceptable, but it cannot explain why none of them commits us to the things quantified over in their conclusion. The theories of ontological commitment and of logical form, then, deal with different questions: the first explains how inferences we accept depend on how the world is, while the second explains how we can turn such inferences into valid ones by analysing appropriately their premises and conclusions. Rather than shape our semantics after our syntax, we should shape our syntax after our semantics.

---

31. Therefore, I do not agree with Orenstein (2000: 517) that “[t]he essence of a theory of logical form is to explain why some natural language inferences are valid and some not”. Logical form alone cannot do that. What it can do, however, is to analyse inferences we accept as being valid according to some criterion and thereby explain *that* some of them are valid.

## References

- Ajdukiewicz, Kazimierz, 1921. *Beiträge zur Methodologie der deduktiven Wissenschaften*. Lwów: Verlag der Polnischen Philosophischen Gesellschaft in Lemberg.
- Alston, William P., 1958. Ontological Commitments. *Philosophical Studies* 9: 8–16. Reprinted in Benacerraf & Putnam (1964).
- Azzouni, Jody, 1998. On “On What There Is”. *Pacific Philosophical Quarterly* 79(1): 1–18.
- Benacerraf, Paul & Hilary Putnam (editors), 1964. *Philosophy of Mathematics: Selected Readings*. Englewood Cliffs, New Jersey: Prentice-Hall, Inc.
- Boghossian, Paul Artin, 1994. Inferential-role semantics and the analytic/synthetic distinction. *Philosophical Studies* 73: 109–122.
- Campbell, John A., 2005. Information-Processing, Phenomenal Consciousness and Molyneux’s Question. In Bermúdez, José-Luis (editor) *Thought, Reference and Experience: Themes from the Philosophy of Gareth Evans*, pp. 195–219. Oxford: Oxford University Press.
- Cartwright, Richard L., 1954. Ontology and the Theory of Meaning. *Philosophy of Science* 21(4): 316–325. Reprinted in Cartwright (1987: 1–12).
- Cartwright, Richard L., 1987. *Philosophical Essays*. Cambridge, Massachusetts: The MIT Press.
- Casati, Roberto, 2004. Is the Object Concept Formal? *Dialectica* 58(3): 383–394.
- Chihara, Charles S., 1968. Our Ontological Commitment to Universals. *Noûs* 2(1): 25–46.
- Church, Alonzo, 1958. Ontological Commitment. *The Journal of Philosophy* 55(23): 1008–1014.
- Clark, Austen, 2000. *A Theory of Sentience*. Oxford: Oxford University Press.
- Clark, Austen, 2004. Feature-Placing and Proto-Objects. *Philosophical Psychology* 17(4): 443–469.
- Cohen, Jonathan, 2004. Objects, Places, and Perception. *Philosophical Psychology* 17(4): 471–495.
- Crick, Francis & Christoph Koch, 1992. The Problem of Consciousness. *Scientific American* 267(3): 152–60.
- Davidson, Donald, 1967. Causal Relations. *The Journal of Philosophy* 64(21): 691–703. Reprinted in Davidson & Harman (1975: 246–254) and Davidson (1980: 149–162).
- Davidson, Donald, 1980. *Essays on Actions and Events*. Oxford: Clarendon Press. 2nd, enl. edition: Davidson (2001).
- Davidson, Donald, 2001. *Essays on Actions and Events. Philosophical Essays Volume 1*. 2 edition. Oxford: Clarendon Press. Enlarged.
- Davidson, Donald & Gilbert H. Harman (editors), 1975. *The Logic of Grammar*. Encino and Belmont, California: Dickenson Publishing Co.
- Davies, Martin Kinsey, 1981. *Meaning, Quantification, Necessity*. London: Routledge & Kegan Paul.
- Demopoulos, William & Michael Friedman, 1985. The concept of structure in Russell’s *The Analysis of Matter*. *Philosophy of Science* 52(4): 621–639. Reprinted in Savage & Anderson (1989: 183–199).
- Devitt, Michael, 1983. Realism and the Renegade Putnam: A Critical Study of Putnam (1978). *Noûs* 17(2): 291–301.
- Domany, E., J.L. van Hemmen & K. Schulten (editors), 1994. *Modals of Neural Networks*, volume II. Berlin: Springer Verlag.
- Donnellan, Keith S., 1974. Speaking of Nothing. *The Philosophical Review* 83(1): 1–31. Reprinted in Hockney et al. (1975: 93–118) and in Schwartz (1977).
- Dummett, Michael A.E., 1981. *Frege: Philosophy of Language*. 2 edition. Cambridge, Massachusetts: Harvard University Press.
- Evans, Gareth, 1982. *Varieties of Reference*. Oxford: Clarendon Press. Edited by John McDowell.
- Evans, Gareth, 1985a. *Collected Papers*. Oxford: Clarendon Press.
- Evans, Gareth, 1985b. Molyneux’s Question. *Behavioral and Brain Sciences* 8: 79–80. Reprinted in Evans (1985a: 364–399).
- Fara, Delia Graff, 2011. Socratizing. *American Philosophical Quarterly* 48(3): 229–238.
- van Fraassen, Bas C., 1991. *Quantum Mechanics: An Empiricist View*. Oxford: Clarendon Press.
- Frege, Gottlob, 1879. *Begriffsschrift, eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*. Halle a.S.: Louis Nebel.
- Hintikka, Jaakko, 1966. Studies in the Logic of Existence and Modality. *The Monist* 50(1): 55–76. Reprinted, with minor changes, as “Existential Presuppositions and Their Elimination” in Hintikka (1969: 23–44).
- Hintikka, Jaakko, 1969. *Models for Modalities: Selected Essays*. Number 23 in Synthese Library, Dordrecht: D. Reidel Publishing Co.
- Hochberg, Herbert, 1957. On Pegasizing. *Philosophy and Phenomenological Research* 17(4): 551–554. Reprinted in Hochberg (1984: 101–104).
- Hochberg, Herbert, 1984. *Logic, Ontology, and Language: Essays on Truth and Reality*. Analytica: Investigations in Logic, Ontology, and the Philosophy of Language, München: Philosophia Verlag.
- Hockney, D.J., William L. Harper & B. Freed (editors), 1975. *Contemporary Research in Philosophical Logic and Linguistic Semantics: Proceedings of a Conference Held at the University of Western Ontario, London, Canada*. Number 4 in The University of Western Ontario Series in Philosophy of Science, Dordrecht: D. Reidel Publishing Co.
- Jackson, Frank, 1977. *Perception: a Representative Theory*. Cambridge: Cambridge University Press.
- Jacquette, Dale (editor), 2002. *Philosophy of Logic: An Anthology*. Blackwell Philosophy Anthologies, Oxford: Basil Blackwell Publishers.
- Jubien, Michael, 1972. The Intensionality of Ontological Commitment. *Noûs* 6(4): 378–387. Reprinted in Jacquette (2002).
- Lewis, David, 1983. New Work for a Theory of Universals. *Australasian Journal of Philosophy* 61(4): 343–377. Reprinted in Lewis (1999: 8–55).
- Lewis, David, 1984. Putnam’s Paradox. *Australasian Journal of Philosophy* 62: 221–236. Reprinted in Lewis (1999: 56–77).
- Lewis, David, 1999. *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.
- Lewis, David, 2001. Truthmaking and Difference-Making. *Noûs* 35(4): 602–615.
- Loux, Michael J. (editor), 1970. *Universals and Particulars: Readings in Ontology*. Notre Dame, Indiana: University of Notre Dame Press. Second edition: Loux (1976).
- Loux, Michael J. (editor), 1976. *Universals and Particulars: Readings in Ontology*. 2 edition. Notre Dame, Indiana: University of Notre Dame Press. First edition: Loux (1970).
- von der Malsburg, Christoph, 1981. The Correlation Theory of Brain Function. Technical Report 81-2, Max Plank Institute for Biophysical Chemistry, Göttingen. Reprinted in Domany et al. (1994).
- von der Malsburg, Christoph, 1999. The What and Why of Binding: The Modeler’s Perspective. *Neuron* 24: 95–104.
- Marcus, Ruth Barcan, 1972. Quantification and Ontology. *Noûs* 6: 240–250. Reprinted in Marcus (1993: 75–88).
- Marcus, Ruth Barcan, 1993. *Modalities: Philosophical Essays*. Oxford: Oxford University Press.
- Martin, Richard Milton, 1962. Existential Quantification and the ‘Regimentation’ of Ordinary Language. *Mind* 71: 525–529.
- Matthen, Mohan, 2004. Features, places, and things: reflections on Austen Clark’s theory of sentience. *Philosophical Psychology* 17(4): 497–518.
- Merrill, Gary H., 1980. The Model-Theoretic Argument Against Realism. *Philosophy of Science* 47: 69–81.
- Moltmann, Friederike, 2003. Nominalizing Quantifiers. *The Journal of Philosophical Logic* 32: 445–481.
- Newman, M.H.A., 1928. Mr. Russell’s “Causal theory of perception”. *Mind* 37(146): 137–148.
- Orenstein, Alex, 2000. The Logical Form of Categorical Sentences. *Australasian Journal of Philosophy* 78(4): 517–533.
- Parsons, Terence D., 1967. Extensional Theories of Ontological Commitment. *The Journal of Philosophy* 64(14): 446–450.
- Putnam, Hilary, 1977. Realism and Reason. *Proceedings and Addresses of the American Philosophical Association* 50(6): 483–498. Reprinted in Putnam (1978: 123–140).
- Putnam, Hilary, 1978. *Meaning and the Moral Sciences*. London: Routledge & Kegan Paul.
- Putnam, Hilary, 1980. Models and Reality. *The Journal of Symbolic Logic* 45: 464–482. Reprinted in Benacerraf & Putnam (1964) and Putnam (1983: 1–25).
- Putnam, Hilary, 1981. *Reason, Truth and History*. Cambridge: Cambridge University Press.
- Putnam, Hilary, 1983. *Realism and Reason. Philosophical Papers, Volume 3*. Cambridge: Cambridge University Press.
- Pylyshyn, Zenon W., 2001. Visual indexes, preconceptual objects, and situated vision. *Cognition* 80: 127–158.
- Quine, Willard van Orman, 1948. On What There Is. *The Review of Metaphysics* 2(5): 21–38. Reprinted as Quine (1951b), in Quine (1953a: 1–19) and in Loux (1970: 33–43).

- Quine, Willard van Orman, 1951a. On Carnap's View on Ontology. *Philosophical Studies* 2: 65–72. Reprinted in Quine (1966b: 203–211).
- Quine, Willard van Orman, 1951b. On What There Is. *Proceedings of the Aristotelian Society, Supplementary Volume* 25: 217–234. Reprinted in Quine (1953a: 1–19).
- Quine, Willard van Orman, 1953a. *From a Logical Point of View: 9 Logico-Philosophical Essays*. Cambridge, Massachusetts: Harvard University Press.
- Quine, Willard van Orman, 1953b. Logic and the Reification of Universals. In Quine (1953a), pp. 102–129.
- Quine, Willard van Orman, 1953c. Meaning and Existential Inference. In Quine (1953a), pp. 160–168.
- Quine, Willard van Orman, 1953d. Reference and Modality. In Quine (1953a), pp. 139–159.
- Quine, Willard van Orman, 1957. Speaking of Objects. *Proceedings and Addresses of the American Philosophical Association* 30. Reprinted in Quine (1969b: 1–25).
- Quine, Willard van Orman, 1960a. Variables Explained Away. *Proceedings of the American Philosophical Society* 104: 343–347. Reprinted in Quine (1966a: 227–235).
- Quine, Willard van Orman, 1960b. *Word and Object*. Cambridge, Massachusetts: The MIT Press. New edition: Quine (2013).
- Quine, Willard van Orman, 1966a. *Selected Logic Papers*. New York: Random House. Second edition: Quine (1995).
- Quine, Willard van Orman, 1966b. *The Ways of Paradox and other essays*. Cambridge, Massachusetts: Harvard University Press. Second, enlarged, edition: Quine (1976a).
- Quine, Willard van Orman, 1969a. Existence and Quantification. In Quine (1969b), pp. 91–113. Repr. 1969.
- Quine, Willard van Orman, 1969b. *Ontological relativity and other essays*. New York: Columbia University Press. Repr. 1969.
- Quine, Willard van Orman, 1970. *Philosophy of Logic*. Cambridge: Cambridge University Press.
- Quine, Willard van Orman, 1976a. *The Ways of Paradox and other essays*. 2 edition. Cambridge, Massachusetts: Harvard University Press. First edition: Quine (1966b).
- Quine, Willard van Orman, 1976b. Whither Physical Objects? In Cohen, Robert S., Paul K. Feyerabend & Marx W. Wartofsky (editors) *Essays in Memory of Imre Lakatos*, pp. 497–504. Number 39 in Boston Studies in the Philosophy of Science, Dordrecht: D. Reidel Publishing Co.
- Quine, Willard van Orman, 1979. Facts of the Matter. In Shahan, Robert W. & Chris Swoyer (editors) *Essays on the Philosophy of W. V. Quine*, pp. 155–169. Norman, Oklahoma: University of Oklahoma Press.
- Quine, Willard van Orman, 1981. Responses. In *Theories and Things*, pp. 173–186. Cambridge, Massachusetts: Harvard University Press.
- Quine, Willard van Orman, 1990. *Pursuit of Truth*. Cambridge, Massachusetts: Harvard University Press.
- Quine, Willard van Orman, 1995. *Selected Logic Papers*. 2 edition. Cambridge, Massachusetts: Harvard University Press. First edition: Quine (1966a).
- Quine, Willard van Orman, 2013. *Word and Object*. Cambridge, Massachusetts: The MIT Press. First edition: Quine (1960b).
- Roskies, Adina L., 1999. The Binding Problem. *Neuron* 24: 7–9.
- Russell, Bertrand Arthur William, 1927. *The Analysis of Matter*. London: Routledge & Kegan Paul.
- Russell, Bertrand Arthur William, 1968. *The Autobiography of Bertrand Russell. Vol. 2: 1914–1944*. London: George Allen & Unwin.
- Russinoff, Ilene Susan, 1992. Frege and Dummett on the Problem with the Concept “Horse”. *Noûs* 26(1): 63–78.
- Savage, C. Wade & Curtis Anthony Anderson (editors), 1989. *Minnesota Studies in the Philosophy of Science, Volume XII: Rereading Russell: Essays on Bertrand Russell's Metaphysics and Epistemology*. Minneapolis, Minnesota: University of Minnesota Press.
- Scheffler, Israel & Noam Chomsky, 1958. What is Said to Be. *Proceedings of the Aristotelian Society* 59: 71–82.
- Schiffer, Stephen, 2003. *The Things We Mean*. Oxford: Clarendon Press.
- Schwartz, Stephen P., 1977. *Naming, Necessity, and Natural Kinds*. Ithaca, New York: Cornell University Press.
- Strawson, Peter Frederick, 1959. *Individuals: An Essay in Descriptive Metaphysics*. London: Methuen & Co.
- Strawson, Peter Frederick, 1961. Social Morality and the Individual Ideal. *Philosophy* 36(136): 1–17.
- Tarski, Alfred, 1956. *Logic, Semantics, Metamathematics*. Indianapolis, Indiana: Hackett Publishing Co. Trans. J.H. Woodger, 2nd ed. 1983 edited and introduced by John Corcoran published by Oxford University Press.
- Treisman, Anne, 1996. The Binding Problem. *Current Opinion in Neurobiology* 6(2): 171–178.
- Wiggins, David, 1980. ‘Most’ and ‘All’: some comments on a familiar programme, and on the logical form of quantified sentences. In Platts, Mark (editor) *Reference, Truth, and Reality: Essays on the Philosophy of Language*, pp. 318–346. London: Routledge & Kegan Paul.
- Wright, Crispin, 1998. Why Frege Did Not Deserve His “Granum Salis”: A Note on the Paradox of “The Concept Horse” and the Ascription of *Bedeutungen* to Predicates. *Grazer Philosophische Studien* 55: 239–263. “New Essays on the Philosophy of Michael Dummett”, ed. Johannes Brandl and Peter Sullivan.