

Introduction

It is an entrenched and plausible view in philosophy that we can gain knowledge of objective truths by evidence other than sense experience. The clearest candidates of this type of knowledge are our claims about non-spatiotemporal domains, as in pure logic and mathematics, and those expressing analytic truths, independently of whether their intended subject matter is abstract or spatiotemporal. Beyond these paradigm cases, there are some other, more contestable examples as well, including our normative claims or value judgments in ethics, aesthetics and epistemology, and the descriptive claims of metaphysics.

Once we believe in the possibility of *a priori* knowledge acquisition, it becomes natural to ask ourselves: what is happening here, *how* do we learn what is objectively and necessarily the case without relying on the deliverances of sense perception? Clearly, any response to this question will draw heavily on what can be reasonably thought of the nature of those conditions whose obtaining or absence is supposed to determine the truth-value of the relevant claims. A proper explanation of *a priori* knowledge requires an appropriate conception of the meaning of *a priori* claims and the nature of *a priori* truths.

In philosophy of mathematics, the mutual dependence of theories of meaning and truth, on the one hand, and theories of knowledge acquisition, on the other, has long been an established part of common sense. This is in great part due to two seminal papers by Paul Benacerraf, published in 1965 and 1973, which have influenced virtually every writer on the subject since.

The first of these (Benacerraf 1965) became the groundbreaking work of mathematical structuralism. According to a structuralist, the

subject matter of mathematical theories is not a single domain of abstract individuals that, beyond having certain relations to each other, also possess some further properties, which distinguish the system in which they feature from other isomorphic ones. Rather, it is either all systems of individuals exemplifying an abstract structure or the structure itself whose elements are merely positions in the structure lacking any further individuating property, so that questions about what mathematical objects really are cannot be answered beyond what the theory says about the defining relations of these objects to each other.

The second paper (Benacerraf 1973) explicates a dilemma, which can be seen as fuelling many of the debates and inventions in early twentieth-century works on the foundations of mathematics. The dilemma is the following. If we maintain that the truth-value of our mathematical beliefs is determined by the obtaining or absence of those abstract and non-epistemic conditions that these beliefs purport to be about (i.e. whether certain mathematical objects possess certain mathematical properties), then we find ourselves unable to understand how we can, by means of our natural cognitive mechanisms, discover whether or not these conditions obtain. On the other hand, if we suppose that knowledge requires appropriate causal contact between knowing minds and the obtaining truth conditions of true beliefs, then we seem to be forced to conclude that the truth conditions of our established mathematical theories cannot be construed along the standard referentialist lines. Summing up, in philosophy of mathematics, our standard referentialist conception of truth seems to be incompatible with our standard causal theory of knowledge.

The concerns of the two articles are intimately related to each other. On the one hand, in a referentialist semantical framework our ideas of the subject matter of mathematics may be highly significant for our theory of mathematical truth. Some structuralists, for instance, believe that their conception of the subject matter of mathematics also resolves the explanatory puzzle about mathematical knowledge acquisition. On the other hand, an acceptable answer to Benacerraf's dilemma can provide us with reasons for taking a stand on mathematical structuralism as well. Some philosophers, for instance, believe that any acceptable answer to the dilemma must involve the rejection of semantical Platonism, i.e. the realist construal of the abstract subject matter of mathematics, which move may influence their views in the

debate over the structuralist understanding of mathematical objects and properties.

The significance of Benacerraf's observations is not confined to the philosophy of mathematics. Similar questions can be raised in the semantics of any discourse in which we are supposed to acquire knowledge of causally inert subject matters. Logic, for instance, is mostly supposed to concern inferential relations among propositions. Is there anything more to being a proposition than having certain inferential relations to other propositions? Are propositions, together with their inferential relations, real entities existing in an abstract (Platonic) realm, or are they merely projected by human minds? Can we maintain that the truth-value of our logical beliefs is determined by the obtaining or absence of those mind-independent conditions that these beliefs purport to be about? If we maintain this referentialist construal of logical truth, can we properly explain how our reasoning capacities could inform us about the obtaining or absence of such causally inert truth conditions? Isn't it the case that any causal account of how we actually *discover* objective logical truths undermines the adequacy of the standard referentialist construal of the truth conditions of logical beliefs?

The six papers collected in this volume address one or another of these questions in the context of mathematical (Isaacson), logical (Rumfitt, Mišćević) and ethical (Wedgwood) beliefs, and in the case of universal generalizations (Williamson) and beliefs involving ideas of universals in general (Robinson). In the remaining part of this introduction, we will briefly review the main theoretical options that one can adopt in response to Benacerraf's dilemma in the semantics of the above problematic discourses, and then provide a short overview of how the positions discussed in the six papers can be located on this theoretical landscape.

* * *

One way to start a review of the conceivable responses to Benacerraf's dilemma is to identify the crucial presuppositions of the case. Corresponding to the two horns of the dilemma, we can classify these assumptions into two major categories: semantical and epistemological assumptions.

The most fundamental semantical assumption behind Benacerraf's original case is that mathematical claims express genuine propositions that are truth-apt, some being true while others false (1973, 666). We shall call this first tenet *cognitivism* in the semantics of mathematics and the other problematic discourses in general.

The second semantical assumption underlying Benacerraf's dilemma is that metaphysical and epistemological considerations may impose substantive constraints upon a proper theory of meaning and truth. In particular, truth and falsity are substantive properties that play an important explanatory role, among others, in our account of knowledge acquisition about various domains (661, 662, 671). We shall call this second tenet *substantivism about truth* in the semantics of the relevant discourses.

Benacerraf's third semantical assumption is that truth in mathematics is a real, non-epistemic property (664, 665, 668, 674, 675, 676). In other terms, the truth conditions of mathematical claims obtain (or not) independently of anyone's actual knowledge of, or capacity to recognize, this particular circumstance, so no epistemic fact involving the truth-value of a mathematical claim is constitutive of the obtaining or absence of the claim's truth conditions. An ideal thinker can still be claimed to be able to know all mathematical truths, but the conceptual ground of this claim is not an epistemic construal of truth, but instead a realist construal of being an ideal thinker. Generalizing from the mathematical case, we shall call this third tenet *realism about truth* in the semantics of discourses about causally inert subject matters.

The fourth semantical assumption, explicitly discussed in Benacerraf's paper, is that the truth conditions of mathematical claims can be specified in terms of the intended subject matter of these claims (i.e. in terms of mathematical objects possessing mathematical properties) (665, 672, 677, 678). The assumption is independent of the previous two, since it does not imply anything substantive concerning the nature of the intended subject matters (664). What it does imply is adherence to the received referentialist construal of mathematical truth in conformity with our notion of truth in the semantics of other segments of natural language. Following Benacerraf's terminology, we shall call this tenet, generally, *referentialism about truth*, emphasizing that the term 'referentialism' has no substantive metaphysical implications here (i.e. that an advocate of this tenet need not commit herself

to any conception concerning the metaphysical status and nature of the relevant subject matters). Putting stress upon the conceptual independence of this tenet from the former two may be significant in the light of two relatively entrenched terminological conventions in present-day philosophy: first, the characterization of construals explaining meaning and truth without any reference to intended subject matters as anti-realist accounts in semantics; and second, the characterization of reference as a substantive relation between representations and represented entities, a notion clearly distinguishable from the deflated concept figuring in the label we suggest here for the received (broadly Tarskian) conception of truth.

The fifth semantical assumption, also explicitly touched upon by Benacerraf, is that the subject matters of mathematical expressions are the kinds of entities they are normally taken to be (673, 675). For instance, numbers and geometrical objects are abstract individuals that are causally inert and have no location in physical space and time. Again, this assumption is clearly independent from the earlier ones. One may maintain that mathematical claims are about abstract (i.e. non-spatiotemporal) states of affairs without subscribing to a substantive, realist interpretation of mathematical objects and properties and also without adopting the referentialist idea that the truth conditions of these claims have to be understood in terms of this abstract subject matter. Again, generalizing from the mathematical case, we shall call this fifth tenet *non-revisionism about subject matter* in the semantics of the relevant discourses.

On the epistemological side, Benacerraf's most fundamental assumption is that at least some of our mathematical beliefs qualify as knowledge (673). We shall call this first epistemological tenet, properly generalized, *anti-skepticism* in the epistemology of the relevant discourses. Since knowledge presupposes the truth of the known proposition, this assumption also implies that the truth conditions of at least some of our mathematical beliefs actually obtain.

The second epistemological assumption behind Benacerraf's case is that the acquisition of knowledge requires an appropriate causal link between the knowing mind and the obtaining truth conditions of the known propositions (671, 672). We shall call this second tenet *a causal theory of knowledge acquisition* in the epistemology of the relevant discourses.

Due to its negative reception in the subsequent literature, the previous assumption is today sometimes replaced by another, which holds that there must be some information-conveying contact between the obtaining truth conditions of known propositions and the knowing minds. The assumption rests on two fundamental convictions: first, that without invoking the existence of such a contact one cannot reasonably explain the reliability of beliefs; and second, that the impossibility of such an explanation undermines any epistemic ground one might have in support of those beliefs (Field 1989, 25–26). We shall call this weaker version of Benacerraf's second epistemological tenet *a contact theory of knowledge acquisition*.

Adopting the five semantical assumptions, we must conclude that the truth conditions of our claims about causally inert subject matters obtain (or not) without exerting any influence upon other existents, including our knowing minds. Adopting the two epistemological assumptions, on the other hand, we must conclude that at least in some cases there is an information-conveying mechanism between the obtaining truth conditions of our claims about causally inert entities and our actual evidence in support of these claims. The two conclusions clearly contradict each other: while the semantical assumptions suggest that there can be no contact between the truth conditions of claims about causally inert subject matters and the human minds, the epistemological assumptions imply that at least in some cases this contact obtains.

Conceivable responses to Benacerraf's dilemma can be classified also into two major categories: those that reject some of the semantical assumptions specified above, and those that abandon some of the epistemological assumptions.

Among the semantical responses, the most radical is the rejection of cognitivism concerning the problematic types of claims. If a claim is not an endorsement of a genuine proposition, thus it cannot be true or false, then it cannot qualify as a piece of genuine knowledge either. Of course, the systematic nature of our linguistic practice may still call for a proper explanation, but this account need not involve reference to the obtaining of any truth conditions. The best example of this *non-cognitivist* treatment of an otherwise problematic discourse is Hare's prescriptivism in metaethics (1952), but the same strategy has been traditionally attributed to metaethical emotivists, such as Ayer (1946) and Stevenson (1944), and more recently, by some authors, to

metaethical expressivists, like Blackburn (1993) and Gibbard (1990). In philosophy of mathematics, Hilbert's instrumentalist account of what he called ideal (infinite) mathematics (1925) is sometimes regarded as an instance of non-cognitivism in the sense specified above. On the current understanding, the crucial tenet of non-cognitivism is that the linguistic practice under scrutiny does not serve the expression of genuine beliefs, because it is not regulated by the detection of the obtaining or absence of some semantically significant conditions that could be regarded as conditions of truth.

A less radical semantical response to Benacerraf's problem is to deny the correctness of the second, substantivist assumption, and adopt a *deflationist* position in the semantics of the relevant discourses. Deflationists maintain that a proper theory of truth and reference is orthogonal to both our conceptions of the metaphysical status and nature of truth or correct declarative use conditions and our theories of how we acquire knowledge of the obtaining or absence of these conditions. In other terms, semantics is autonomous *vis-à-vis* metaphysics and epistemology. The main reason for this is that, on a deflationist understanding, truth is not a substantive property, so there is nothing to say about its nature and its relation to our epistemic capacities. Our notion of truth is fully characterized by the instances of Tarski's Disquotatation Schema or its counterpart for propositions as primary truth-bearers. A deflationist may still wonder how we can acquire knowledge of causally inert subject matters, and maybe even admit that, indeed, there is something theoretically puzzling in this phenomenon. Nevertheless, contrary to Benacerraf's claim, she can maintain that no response to this challenge can undermine the adequacy of referentialism about truth, since playing a substantive explanatory role in theories of knowledge is not a prerequisite for a condition to become constitutive of the truth conditions of a truth-apt representation. Classical versions of deflationism include Ramsey's redundancy theory (1927), Strawson's performative theory (1950), and Quine's disquotational theory (1970), while the most influential recent forms of deflationism are Grover, Camp and Belnap's prosentential (1975) and Horwich's minimal (1990) theories of truth. Beyond these clearly anti-substantivist examples, deflationist conclusions can be derived from Blackburn's (purportedly substantive anti-realist) "quasi-realist" program (1993) in semantics as well: if all distinctive claims of a realist can be endorsed,

on some re-interpretation, by an anti-realist as well, then it may seem quite natural to question the intelligibility of the very contrast between realism and anti-realism, and opt for a deflationist theory of truth.

The third available semantical reaction to Benacerraf's dilemma is to accept the cognitivist and substantivist assumptions, but deny the adequacy of realism, and adopt an *anti-realist* position about truth in the semantics of discourses about causally inert subject matters. Anti-realists about truth maintain that truth is a substantive epistemic property. In other terms, they hold that the truth conditions of a certain class of claims do not obtain independently of our capacities for recognizing these truths, or more exactly, that some epistemic facts concerning the truth-values of these claims are constitutive of the obtaining or absence of those truth conditions. Anti-realism in semantics and metaphysics has always found its basic motivation in epistemological considerations. No wonder that the doctrine may appear as a solution to Benacerraf's dilemma as well. If the truth-value of our claims about causally inert subject matters are construed in epistemic terms, then the explanation of knowledge acquisition need not invoke an information-conveying link between the knowing mind and something whose existence is fully external to it. Anti-realist replies may differ in their stance to Benacerraf's fourth and fifth semantical assumptions, i.e. whether they maintain or reject referentialism about truth and non-revisionism about subject matter in the semantics of the relevant discourses. It may be worth noting, however, that some influential doctrines from among those which are often classified as anti-realists about truth are arguably realist in the currently adopted sense of the term. Putnam's internal realist epistemisation of truth, for instance, is sometimes presented as a representative of (a referentialist and non-revisionist form of) anti-realism concerning this entity. Putnam, however, has never claimed that epistemic states are constitutive of the obtaining or absence of referential truth conditions.¹ Dummett's verificationist theory of truth cannot be regarded as anti-realist in the current sense either, since the conditions that he takes to be the truth conditions of our beliefs are supposed to obtain also independently of anyone's

¹ What the internal realist Putnam (1981) argues for is that the identity conditions of the intended states of affairs that can be regarded as the referential truth conditions of our beliefs are created by the classificatory work of mind.

actual knowledge of, or capacity to recognize, this particular circumstance.² Three further examples, whose anti-realist status is contestable, are Gibbard's projectivist semantics in metaethics (1990), Blackburn's quasi-realist construal of our claims about moral and modal states of affairs (1993), and Peacocke's conceptualism in the semantics of *a priori* discourses (2005). More plausible examples of anti-realism about truth include the construals of subjective idealists, Carnap's conventionalism about *a priori* (analytic) truth (1934), and maybe Brouwer's intuitionist theory in philosophy of mathematics (1949).

The fourth semantical response to Benacerraf's dilemma is to reject the referentialist assumption, and adopt *non-referentialism about truth* in the semantics of discourses about causally inert subject matters. Non-referentialists maintain that the truth conditions of a certain class of claims cannot be specified in terms of the intended subject matter of the constituents of these claims. For instance, on a non-referentialist construal, the truth conditions of mathematical claims are not mathematical states of affairs, whichever way these would be further understood. Rather, they are conditions that may or may not obtain in a non-mathematical realm. Non-referentialism does not imply anything about the metaphysical status and nature of the relevant subject matters. Nevertheless, it makes realism about truth compatible with anti-realism, fictionalism, eliminativism or quietism about subject matters. Of course, as Benacerraf rightly observed, an advocate of this position must explain what makes her preferred non-referential truth conditions qualify as conditions of *truth*. Once she can deliver this explanation, she can construe the relevant conditions, without reducing the corresponding subject matters, either in anti-realist or in epistemologically unproblematic realist terms. Examples of this non-referentialist strategy may include Dummett's verificationist construal of truth in discourses about epistemically inaccessible domains (1991), Blackburn's (1993) and Gibbard's (1990) projectivist theory in metaethics, and Putnam's (1967) and Hellman's (1989) modal structuralism in philosophy of mathematics. If conative attitudes, possession conditions of concepts and analytic links within our personal system of representation are

² What Dummett's (1991) anti-realist assumes is that the truth conditions of our beliefs are always verifiable (i.e. that we have an effective, though fallible, method to determine whether or not they actually obtain).

construed realistically, then a number of influential accounts that have been developed as an alternative to the epistemologically problematic realist and referentialist construals of truth can be classified as realist in the semantics of discourses about causally inert entities.

The fifth semantical strategy that may be adopted in response to Benacerraf's dilemma is to reject the fifth semantical assumption specified above, and embrace a *revisionist* construal of the referential domain of the problematic discourses under scrutiny. In the case of mathematics, for instance, this would amount to the view that mathematical claims are not about abstract states of affairs whose constituents are causally inert and have no spatiotemporal location, but instead they are either about some aspects of the natural world, or about some concepts in an active intellect, or about some other entities that can influence the human mind. Alternatively, a revisionist can take mathematical claims to be about *in re* or *ante rem* structures, rather than about a single system of individuals whose members, beyond having certain relations to each other, also possess some intrinsic properties that distinguish the system they constitute from isomorphic systems of other individuals. In other terms, she can take these claims to be either about all systems of individuals exemplifying a certain structure or about the structure itself that can be exemplified by those systems. Revisionism in itself does not imply anything about the metaphysical status of the relevant subject matters. A structuralist interpretation of mathematics, for instance, is compatible with a deflationist, an anti-realist and a realist construal of mathematical referents as well. Nonetheless, a major motive behind a revisionist construal of the subject matter of mathematics and other discourses about *prima facie* causally inert subject matters is that this construal allows for the wedding of a substantive realist and referentialist understanding of truth with a causal contact theory of knowledge acquisition in the philosophy of the relevant discourses. Theories falling into this class include Mill's and Kitcher's referentialist naturalism and various forms of structuralism in philosophy of mathematics (Mill 1843; Kitcher 1984). A short note about the structuralist proposals, however, may be in order. Although structuralist construals of Platonic entities are revisionist in character, not all of them support an effective response to Benacerraf's dilemma. In particular, those *ante rem* and *in re* structuralists who maintain that the problematic discourses under scrutiny are about causally inert and real

entities cannot claim that their view is more compatible with a causal theory of knowledge acquisition than the conception that a traditional (non-revisionist) Platonist provides.

In case one does not want to follow any of the five semantical strategies characterized so far, one may try to answer Benacerraf's dilemma by querying at least one of the epistemological assumptions of the case. The most radical epistemological strategy is to deny the possibility of knowledge about causally inert entities. If our received theories of such entities do not qualify as knowledge, then the adoption of Benacerraf's five semantical assumptions concerning the claims and implications of these theories remains compatible with the standard causal theory of knowledge. A limited form of skepticism may result from an error-theorist view of our received beliefs concerning causally inert entities. Such views have been defended by Mackie (1977) in metaethics and Field (1980) in philosophy of mathematics. An error-theorist argues that our received conceptions of a certain domain are false, and thus cannot qualify as knowledge, because the world does not contain those individuals and properties whose existence is required for their truth. Note, however, that an error-theorist need not assume that the existence of the relevant entities is a precondition of *any* truth about the corresponding domains. In absence of this assumption, she may maintain, for instance, that negative existential beliefs about causally inert entities are still true, and as such potentially qualifying as knowledge. A limited skepticism like this cannot resolve Benacerraf's dilemma. In order to save the compatibility of the standard realist and referentialist semantics with the standard causal theory of knowledge, one must deny the existence of any type of knowledge of the relevant causally inert domains.

The second epistemological strategy to follow in response to Benacerraf's case is to insist on the adequacy of the five semantical and the first epistemological assumptions, and query the idea that the acquisition of knowledge requires an appropriate causal link between knowing minds and obtaining truth conditions. Instead of admitting the adequacy of this causal account, the proponents of this position may argue that in the case of discourses about causally inert subject matters the contact between minds and obtaining truth conditions is not causal in character. We shall call this alternative a *non-causal contact theory of knowledge acquisition*. The classic example of this strategy is Gödel's

quasi-perceptionist account of mathematical knowledge (1944), while more recent instances of this category include Bonjour's account of rational insight (1998) and Brown's Gödelian view of mathematical knowledge and the nature of thought experiments (1991).

The third epistemological option one can choose in response to the dilemma is to deny the adequacy of the second epistemological tenet even in its weaker form, and subscribe to a *no-contact theory of knowledge acquisition*. Advocates of this position maintain that, although a proper explanation of knowledge acquisition may require an account of how the epistemic grounds of a knowing mind for adopting a certain class of true beliefs can reliably indicate the obtaining of the truth conditions of these representations, nevertheless, at least in the case of our discourses about causally inert subject matters, this account need not invoke the existence of any contact between these grounds and those conditions. Examples of this category include Wright's (1983) and Hale's (1987) neo-Fregean abstractionist, Balaguer's (1998) full-blooded Platonist, Katz's (1981) and Lewis's (1986) necessity-based, and Shapiro's (1997) and Resnik's (1997) structuralist strategy to account for the possibility of mathematical knowledge. In case one takes his holistic view of science together with his conception of ontological commitment seriously, Quine's empiricist epistemology (1948, 1951) also qualifies as a no-contact theory of mathematical knowledge.

Having reviewed the main theoretical options that can be adopted in response to Benacerraf's dilemma in the semantics of discourses in which we are supposed to acquire knowledge of causally inert subject matters, in the last part of this introduction, we shall provide a brief summary of the six papers appearing in this volume, and locate the positions they discuss on the theoretical landscape specified so far.



The first contribution, Daniel Isaacson's "The Reality of Mathematics and the Case of Set Theory" presents a structuralist account of mathematics that is realist about truth without being committed to the existence of mathematical objects.

Isaacson motivates his position by offering two reasons for rejecting the standard, Platonist semantics of mathematics in terms of particular abstract objects. The first is based on Benacerraf's dilemma:

Platonism is incompatible with the contact theory of knowledge acquisition. The second consideration, going back to Benacerraf's structuralist paper, is that for any allegedly correct interpretation of mathematical discourse in terms of abstract objects, there are other, isomorphic interpretations with equal claim to correctness, since mathematical truth is invariant under isomorphism.

According to Isaacson's alternative semantic proposal, both the subject matter of mathematics and the truth conditions of mathematical statements are constituted by particular mathematical structures, rather than mathematical objects, where the defining mark of particular structures is that all of their exemplars are isomorphic to each other. Particular structures (e.g. the structure of the natural numbers or that of the continuum) are contrasted with general structures (e.g. the structure of groups), which do have non-isomorphic exemplars.

Before expounding his own view on what particular structures are, Isaacson criticizes three rival accounts: the model-theoretic notion of structure, Stewart Shapiro's theory of structures, and modal structuralism. He points out that all of these approaches conceive of particular structures as constituted by objects (although of different kinds), thus contradicting the structuralist insight that the subject matter of mathematics consists exclusively of structures. He also argues that since model theory, and—at least on one interpretation—Shapiro's axiomatic theory of structures are themselves mathematical theories, it is circular, or the beginning of an infinite regress to take them as accounts of the subject matter of mathematical theories in general. As regards the last approach, modal structuralism, Isaacson is deeply skeptical about the role of modality in an *in re* account of particular structures: he thinks that there cannot be a difference between possible and actual existence for objects constituting the subject matter of mathematics.

Turning to his own positive account of structures, Isaacson articulates an *ante rem* realist view, according to which, instead of being constituted by objects, particular mathematical structures are given by coherent and categorical higher-order characterizations. When a characterization satisfying these conditions is established, the characterized particular structure also acquires real existence. In elucidating the notion of coherent characterization, Isaacson relies on Kreisel's conception of informal rigor: it is practicing mathematics through informal rigor which enables us to develop and understand descriptions of

mathematical structures and see their coherence. Claiming that establishing characterizations is an activity of the human mind, the paper arrives at the anti-Platonist thesis that mathematical characterizations are constituted by our concepts, and particular structures do not exist in advance of, or independently from these characterizations.

The semantics based on this conception of particular structures is tested on the foundationally central case of set theory. Isaacson argues that the second-order version of the Zermelo-Fraenkel theory of pure sets, which is quasi-categorical (its non-isomorphic models differ only in their “height”), coherently characterizes the structure of the cumulative set-theoretic hierarchy. Consequently, in his view, if a statement is semantically decided by this theory (in the sense that either the statement itself or its negation is a semantic consequence of the theory), then it has a determinate truth-value. Since Cantor’s continuum hypothesis is such a statement, Isaacson’s position implies that the continuum problem has a determinate answer, which set theorists may reasonably attempt to find.

Isaacson presents a revisionist and referentialist semantics, on which both the truth conditions and the subject matter of mathematical theories are constituted by particular structures. The proposed account is realist about mathematical truth and mathematical entities (although not about mathematical objects). The most important question left open by the paper concerns the metaphysical status of particular mathematical structures. On Isaacson’s account, particular structures are dependent on their conceptual characterizations. On the one hand, this dependence points towards a construal on which structures are constituted by our concepts. In this case, it is difficult to see how they could be different from their characterizations. (In fact, at one point Isaacson writes that structures *are* characterizations.) On the other hand, the fact that particular mathematical structures are considered to be *ante rem* types suggests a non-spatiotemporal interpretation, which would make the dependence on conceptual characterizations much more problematic. The two alternatives differ considerably from an epistemological point of view: Isaacson’s claim that we know mathematical structures through their characterizations and his commitment to the contact theory of knowledge acquisition can be easily combined with the view that structures are conceptual, but are difficult to maintain if structures are considered to be abstract (i.e. non-spatiotemporal) and real. If both the

epistemological advantages of the conceptual option and the intuition that characterizations and structures belong to different ontological categories are to be retained, then a non-referentialist construal of mathematical truth might be more congenial: such an approach could be realist about mathematical truth by claiming that the truth conditions of mathematical statements are constituted by characterizations adopted in space and time, while still holding that the subject matter of mathematics consists of non-spatiotemporal particular structures.

The second paper, Nenad Mišćević's "Conceptualism and Knowledge of Logic: A Budget of Problems" offers a criticism of conceptualism about logic, a view at the heart of recent attempts to answer Benacerraf's dilemma for logical discourse.

Adherents of conceptualism claim that logical knowledge is apriorist, and its source is our possession and mastery of logical concepts. Although conceptualists share these commitments, and many of them hold that the subject matter of logic is non-spatiotemporal, they still have a variety of theoretical options regarding the metaphysics of concepts, the role logical concepts play in the epistemology and semantics of logic, and the status of logical truth. From a metaphysical point of view, conceptualism is compatible with both realism and anti-realism about concepts, and also with a quietist rejection of both options. As regards the role of logical concepts, one can be conceptualist about justification, holding that the possession of logical concepts is a source of justification for logical knowledge, or about logical truth, characterizing the truth conditions of logical statements in terms of the properties and relations of our logical concepts (of course, the two options are compatible). Conceptualism about logical truth is a non-referentialist position, since the properties and relations of our logical concepts are not what the logical claims built from these concepts are about.

Since Mišćević adopts a broadly realist stance on logical truth, and takes the coincidence of truth conditions with subject matters for granted, the main targets of his criticism are philosophers defending a realist, referentialist, and justification conceptualist view of logic.

The paper examines three purportedly *a priori* sources of justification for logical laws that are prominent in the conceptualist literature. The first source is supplied by the influential thesis that respecting certain basic inference patterns is constitutive of possessing logical

concepts: many conceptualists maintain that this constitutive status can serve as the basis of an *a priori* justification of our belief in the validity of these standard inferences. In addition to this, conceptualist accounts of logical knowledge frequently rely on the obviousness and compellingness of basic inference patterns and principles, often in the form of claiming that finding these logical laws obvious and compelling is also a necessary condition for the possession of logical concepts. The third supposedly *a priori* source of justification is provided by the fact that logic is indispensable for any rational cognitive project.

Focusing mainly on Christopher Peacocke's version of conceptualism, Mišćević claims that the *a priori* justification envisaged by constitution-based accounts moves in an unacceptably narrow circle, and justifying our logical beliefs by a sufficiently wide equilibrium has to take into account the empirical success of our world-directed abilities to imagine certain complex situations while deeming others inconceivable. He also queries, following Timothy Williamson and Tyler Burge, the constitutivity assumption itself by pointing out that persons possessing logical concepts might fall short of the conceptualist requirements, e.g. because of having idiosyncratic views on logic.

Turning to the prospects of *a priori* justification through obviousness and compellingness, Mišćević argues that this characteristic is either a psychological-cognitive or a normatively characterized epistemological property. In the first case the connection between the subjective feeling of certainty and the purported objective warrant calls for an explanation, whereas in the latter an epistemological account of how we can reliably detect this normative property needs to be given. In both cases, Mišćević suggests, the required explanation would have an empirical character, making the amended justification *a posteriori*. In a similar vein, conceptualist accounts relying on indispensability are criticized on the ground that only cognitive projects with a good chance of success warrant their necessary preconditions, and the prospective success of a project is, to a large extent, an empirical matter.

Although his paper is mainly concerned with criticizing justification conceptualism about logic, Mišćević indicates that in his view the justification of logic is ultimately *a posteriori*, and is inherently related to the evolutionary success of our reasoning practices, with a special emphasis on the success of the "built in" ability to reason about concrete situations. From the perspective of Benacerraf's problem, the

main difficulty with this proposal concerns the ontological status of the truth conditions of general logical principles. On the one hand, if they are non-spatiotemporal, then how can we have *a posteriori* information about their obtaining? A well-known response to this problem, which Mišćević might be attracted to, is accepting a radical empiricist form of confirmation holism. This move, however, can hardly satisfy Benacerraf's curiosity, since an advocate of this doctrine would still owe us an account of how our empirical evidence could reliably indicate the obtaining of the purported abstract truth conditions of our logical beliefs. On the other hand, finding viable spatiotemporal truth conditions for principles of pure logic can be a daunting task in a referentialist framework, since our referential intentions appear to exclude the spatiotemporal construal of logical objects and properties.

As has been mentioned, Mišćević's arguments against conceptualism are targeted on approaches that are realist and referentialist about logical truth. This leaves open the question how effectively an aposteriorist could argue against accounts combining a concept-based, apriorist epistemology for logic with a non-referentialist, conceptualist construal of logical truth.

Ian Rumfitt's paper, "What is Logic?" develops a new, revisionist account of logic's subject matter, and applies it to the problem of explaining the epistemological role of deductive reasoning.

According to the standard view, logic investigates the unique relation of consequence that holds between the premises and the conclusion of a sound argument, and the *relata* of this relation are truthapt objects, which are, or can be, expressed by declarative utterances. Challenging this widely accepted picture, Rumfitt argues that we do not have a pre-theoretical grasp of a single relation of logical consequence, since the standards of soundness vary from context to context. Consequently, he claims, it is a mistake to describe logic as a science investigating a unique relationship. He also suggests that the existence of objects that are expressed by utterances should not be assumed early in an inquiry into the nature of logic, since our ordinary standards for assessing whether two utterances "say the same thing" are inadequate to sustain judgements of strict identity. In keeping with these reservations, Rumfitt adopts the traditional terminology and calls the *relata* of consequence relations propositions, but proposes to construe a

proposition as an ordered pair consisting of a declarative sentence type and a possible context of utterance.

Having criticized the standard view on logic's subject matter, Rumfitt turns to expounding his alternative account, according to which there exist a large number of different consequence relations, and logic is concerned with finding general truths valid for all of them. Consequence relations are set apart from other relations among the same type of *relata* by their modal characterization: for each space of possibilities there is a corresponding consequence relation that holds between a set of premises and a conclusion if and only if the conclusion in its actual sense is true in all possibilities of the space in which the premises in their actual sense are true. This modal conception of consequence relations differs not only from Russell's extensionalist construal of consequence as a relation that holds when either one of the premises is actually false or the conclusion is actually true, but also from Tarski's theory of consequence, which is formulated in terms of the actual truth-values of the premises and the conclusion under possible reinterpretations.

Rumfitt motivates his account by showing how it explains the epistemological role of deduction: if a thinker is reliable in deducing a conclusion from some premises only when the contextually relevant consequence relation holds, then she can apply her deductive capacity to premises she already knows, and splice together different pieces of knowledge in a conclusion whose truth is guaranteed by the consequence relation in question. In addition to being true, her belief in the conclusion will have been produced by a reliable belief forming procedure, satisfying one of the most important necessary conditions of knowledge. Rumfitt emphasizes that besides deducing conclusions from known premises, we can also attain new knowledge by reasoning deductively from premises that are *supposed* to be true. The conditions of soundness in these cases can be explicated only by invoking modal notions. Relying on this analysis of deductive capabilities and his construal of logical laws as general truths about consequence relations, Rumfitt accounts for the usefulness of logic by observing that learning a logical law enables us to extend our deductive capacity with respect to *all* consequence relations and gain new knowledge by deduction in every context we encounter.

Despite his rejection of the view that logic is concerned with a single consequence relation, Rumfitt holds that it is possible to identify

a broadly logical relation of consequence, which he characterizes, following Ian McFetridge, as being applicable and truth-preserving in the presence of any supposition. He also accepts, although with important qualifications, McFetridge's thesis that the correlative, broadly logical possibility is the weakest non-epistemic possibility: any proposition, which is possible in a non-epistemic sense (physically, metaphysically etc.), is possible in this weakest sense as well.

Rumfitt's referentialist account of logical truth characterizes the conditions of an argument's soundness and the truth conditions of logical laws in modal terms. At least some of these modal conditions are assumed to be real, as he emphasizes that broadly logical modality is not epistemic. Since non-actual possibilities are presumably causally isolated from the actual world, a variant of Benacerraf's dilemma might be applicable to Rumfitt's proposal: if the truth conditions of logical claims are causally isolated, then it is difficult to see how we could have a non-skeptical epistemology of logic that is compatible with the contact theorist account of knowledge acquisition.

Timothy Williamson's "Absolute Identity and Absolute Generality" argues that we can sharpen our understanding of certain issues concerning absolutely general quantification by drawing on analogies with corresponding issues concerning identity.

The paper starts with an interpretative question: what makes it the case that a speaker means identity by her predicate 'identical'? Williamson observes that if the speaker is committed to the reflexivity of 'identical,' and accepts the indiscernibility principle "if something is identical with something, then whatever applies to the former also applies to the latter," then—assuming that both of these commitments are correct, and her other words can be interpreted homophonically—it is provable that identity in her sense is coextensive with identity in our sense. Although the proof is elementary, it depends on the assumption that the quantifier 'whatever' in the speaker's indiscernibility principle ranges over a domain that includes identity in our sense. Williamson therefore turns to examining the problem in a more general setting, where the homophonic interpretation of words different from 'identical' is not taken for granted.

Formalizing the speaker's commitments in first-order logic, he points out that if the indiscernibility principle is construed as a schema

covering only formulas of the speaker's language, then there will be models that verify all commitments of the speaker, and still do not interpret 'identical' by identity. Williamson rejects excluding these non-standard models simply on the ground that they are not models for first-order logic with identity, and also criticizes Quine's methodology of interpretation, which requires us to interpret languages in a way that the strongest expressible indiscernibility relation is identity. Nonetheless, he argues that it would be a mistake to adopt Peter Geach's relativistic account of identity, according to which a predicate can play the role of identity only relative to a given language.

In Williamson's view, the key to formulating a satisfactory account of our grasp of absolute (i.e. not language-relative), identity lies in the fact that our commitment to the indiscernibility principle is open-ended: our predicate 'identical' refers to absolute identity at least partly because instead of treating the indiscernibility principle as limited to our current language, we have a general disposition to accept its instances in extensions of our language as well. In support of this account, Williamson recalls that reflexivity and the open-ended indiscernibility principle together uniquely characterize identity: if two predicates belonging to different first-order languages both satisfy these principles, then they are coextensive over the intersection of the two domains of quantification.

Turning to problems concerning absolute generality, Williamson establishes an analogous unique characterization result about universal quantification by proving that if two quantifiers belonging to different languages are both governed by open-ended analogues of the standard $\forall$ -introduction and $\forall$ -elimination rules, then they are logically equivalent. He claims that since the proof is purely syntactic, and does not depend on a semantic analysis of the universal quantifier, the result can be used to support the thesis that we have an idea of absolute, unrestricted generality in virtue of our commitment to the open-ended $\forall$ -introduction and $\forall$ -elimination rules.

Williamson formulates and answers a series of possible objections to this account of our grasp of absolute generality. Perhaps the most important among them is that the unique characterization argument presupposes the logical validity of the open-ended rules on the unrestricted reading, whereas they can be only materially valid, since the universal quantifier can be interpreted as ranging over any domain, not

exclusively over everything. In response to this objection, Williamson points out that we have to distinguish between the variable and the constant unrestricted accounts of the universal quantifier. According to the former view, for any things, the universal quantifier can be legitimately interpreted as ranging over just those things, while on the latter account the quantifier is mandatorily interpreted as unrestricted. Williamson's argument presupposes the correctness of the constant unrestricted account, as the open-ended rules are not logically valid on the variable construal.

Williamson does not argue against the variable account, but he notes that the analogy between absolute generality and absolute identity can help us to see that one argument against the constant unrestricted account is problematic. The fact that every interpretation that is legitimate on the constant account is also legitimate on the variable account but not *vice versa* cannot be used to support the claim that the variable account is preferable because of its greater generality, as the analogous argument for preferring first-order logic without identity to first-order logic with identity clearly fails.

The paper concludes with the observation that two frequently voiced objections to the absolute conception of generality are incompatible, as relativists cannot claim that generality absolutism is both inarticulate and inconsistent without undermining their own position. The incompatibility of the two charges is shown by an indirect argument: if absolutism was both inarticulate and inconsistent, then there would be a legitimate relativist interpretation of the universal quantifier on which everything accepted by the absolutist is true, and a valid proof of an explicit contradiction could be given from premises accepted by the absolutist. This proof would also demonstrate the inconsistency of generality relativism, as it would show that the relativist is committed to the existence of a legitimate interpretation on which an explicit contradiction should be true.

Although Williamson's paper makes no reference to this fact, there is an important link between absolutely general, unrestricted quantification and the structuralist approach to the semantics of set-theory. As we have already seen in connection with Isaacson's contribution, categorical theories play a key role in structuralist accounts, since they are considered as characterizing a unique structure, which determines the truth-value of every sentence of the theory's language. It has also been

mentioned that second-order Zermelo-Fraenkel set theory is only quasi-categorical. Nonetheless, an important result due to Vann McGee (1997) shows that the categoricity of set theory can be saved by switching from the standard variable construal of the universal quantifier to the constant unrestricted account: if only those interpretations are considered legitimate whose domain of quantification includes absolutely everything, then second-order Zermelo-Fraenkel set theory with urelements can be shown to be categorical in the sense that all of its models have isomorphic pure set hierarchies. In view of McGee's result, Williamson's vindication of the constant unrestricted construal of universal quantification can be seen as paving the way for a structuralist account of set theory that is capable of conferring a determinate truth-value on every set-theoretic statement.

Ralph Wedgwood's "The Refutation of Expressivism" argues against expressivist accounts in the semantics of normative discourse, and in favor of the rival, truth-conditional approach. Starting with a characterization of expressivism as the view that the fundamental explanation of the meaning of normative statements is to be given in terms of the mental states they express, without making reference to their truth conditions, the paper focuses on two variants of the position: emotivism and Allan Gibbard's semantic theory.

According to an emotivist, instead of representing properties, normative terms express certain feelings towards the evaluated objects or actions. Consequently, normative statements lack truth conditions and cannot be characterized as true or false. In response to this theory, Wedgwood recalls Peter Geach's critical observation that emotivism cannot account for the meaning of normative expressions in a variety of contexts, for instance when they occur within the scope of truth-functional sentential operators, or in propositional attitude attributions. Wedgwood argues that an adequate semantics has to explain how normative terms can occur in these contexts in the same sense as in others, must provide a uniform interpretation of the operators within whose scope they are embedded, and has to account for the fact that our customary inference patterns are valid in the case of normative statements as well.

Allan Gibbard's theory combines a quasi-realist stance on normative discourse (involving a thin interpretation of truth, proposi-

tions, properties etc.) with a psychologistic semantics on which normative statements express mental attitudes that can be characterized in terms of “hyperplans”: complete, consistent plans about what to do and think in every conceivable situation. Wedgwood admits that Gibbard’s theory satisfies the constraints imposed by Geach’s objection. Nevertheless, he claims that an adequate construal of our normative claims must also meet some further conditions. Building on Crispin Wright’s observations, he argues that our normative discourse is a disciplined enterprise, in which speakers aim to comply with certain standards of justification or warrantedness. Consequently, an acceptable semantics for normative claims has to provide an account of these standards as well. Wedgwood claims that any account satisfying this condition will assume that there is a property of normative statements, which is the point or purpose of normative discourse in the sense that speakers try to conform to the standards of justification and warrantedness because they aim at only accepting statements having this property. In addition, the property in question must have certain features that are possessed only by truth: among others, it is preserved in valid arguments, and satisfies an analogue of the T-schema. Wedgwood concludes that truth and truth conditions have an explanatory role in any adequate semantics of normative discourse, and consequently expressivism in general and in its specific Gibbardian form is untenable.

Wedgwood emphasizes that his discussion of expressivism relies on a substantive view about normative truth. In his book *The Nature of Normativity* (Wedgwood 2007), he develops a realist and referentialist truth-conditional semantics of normative discourse, and combines the revisionist thesis that non-physical normative facts can be causally efficacious with the view that in certain cases our knowledge of normative truths is based on *a priori* intuitions.

The motivation Wedgwood offers for classifying Gibbard’s account as expressivist shows that his notion of truth-conditional semantics is restricted to theories that explain meaning in terms of referential truth conditions. By contrast, his central argument against expressivism seems to establish only that an acceptable semantics of normative discourse has to refer to truth conditions *simpliciter*—he does not argue for the stronger claim that only referentialist construals can be adequate. As a consequence, Wedgwood’s argument does not rule out

the possibility of formulating an adequate truth-conditional, but nonetheless non-referentialist account that specifies the truth conditions of normative statements in the psychologistic terms characteristic of expressivist explanations of meaning.

The last paper of the volume, Howard Robinson's "Benacerraf's Problem, Abstract Objects and Intellect" argues that a variant of Benacerraf's dilemma can be raised with respect to any discourse, as all thought involves abstract objects in the form of universals. Robinson sketches a solution to the problem, according to which abstract objects exist only as concepts, and the ability of concept apprehension is a primitive and defining property of intellect.

The paper begins with a reconstruction of Benacerraf's original dilemma. On Robinson's interpretation, by pointing out that Platonism about mathematical truth and a naturalist, causal theory of knowledge acquisition entail that there is no mathematical knowledge, Benacerraf sets out a problem for philosophers who wish to combine realism about mathematics with naturalism. Robinson argues that it is possible to formulate a more general, discourse-independent variant of this challenge to naturalism: the ability to grasp abstract objects in the form of universals is a necessary precondition for thinking about any domain, but there is no viable, naturalistically acceptable account of this capacity. In support of this negative claim, Robinson criticizes the causal theory of conceptual representation, which he considers the standard contemporary naturalist solution to the problem of concept apprehension. In addition to pointing out that the causal theory ignores the connection between thinking and consciousness, the main problem he adduces is that an analysis of apprehension in terms of causal relations cannot account for the normativity of conceptual content.

Having argued against the feasibility of a reductive causal analysis, Robinson finishes his case against naturalistic theories of concept apprehension by dismissing the non-reductivist position that takes apprehension as primitive, but still maintains that grasping a universal necessarily involves causal contact with its instances. He claims that if we had a causally based primitive grasp of certain empirical features of the world, then *a priori* reflection on these features could acquaint us with a virtually unlimited number of further universals, but allowing

that the mind has an irreducible grasp of a vast range of (potentially uninstantiated) universals is incompatible with naturalism.

After his criticism of naturalistic accounts of concept apprehension, Robinson turns to examining the ontological status of universals. Surveying the three traditionally distinguished theoretical options, Platonism, *in re* realism, and conceptualism, he argues that neither of the realist approaches is viable, because Platonism cannot explain the predicative nature of the supposedly separate and self-subsistent universals, while *in re* realism is incompatible with realism about possibilities involving uninstantiated universals.

In line with his arguments against reductive theories of apprehension and realism, Robinson's own account of universals is a conceptualist view, according to which the ability to grasp concepts is a primitive and defining mark of intellect. Universals exist as concepts, but not as concepts of humans, because this would imply that universals unthought by humans do not exist, but as concepts of an objective intellect. Robinson indicates that this construal, which he terms a neo-Platonic account of universals, can be generalized to other classes of abstract entities, in particular to mathematical objects. On the general neo-Platonic view he advocates, the existence of abstract objects is independent from and presupposed by the existence of individual objects, and abstract entities exist as modes of understanding exercised by an objective intellect.

Relying on a revisionist conception of universals and abstract objects in general, Robinson provides a theory of *abstracta* that is realist from the human point of view, as no epistemic facts about humans are considered to be constitutive of the existence of abstract objects, but anti-realist from a global perspective, since abstract objects are supposed to exist only inasmuch as they have a role in rendering the world intelligible for an objective intellect. As a reply to Benacerraf's dilemma, the most important challenge Robinson's theory has to face is explaining how humans attain knowledge about abstract entities. In view of the strong intuitive appeal of a contact theory of knowledge acquisition, it seems questionable that by construing abstract objects as concepts of an objective intellect, neo-Platonists are significantly better placed to meet this challenge than proponents of traditional Platonism.

References

- Ayer, A.J. 1946. *Language, Truth and Logic*. 2nd rev. and reset ed. London: Victor Gollancz.
- Balaguer, M. 1998. *Platonism and Anti-Platonism in Mathematics*. New York: Oxford University Press.
- Benacerraf, P. 1965. "What Numbers Could not Be." *Philosophical Review* 74: 47–73. Repr. in Benacerraf and Putnam 1983, 272–294.
- . 1973. "Mathematical Truth." *Journal of Philosophy* 70: 661–679. Repr. in Benacerraf and Putnam 1983, 403–420.
- Benacerraf, P. and H. Putnam, eds. 1983. *Philosophy of Mathematics*. 2nd ed. Cambridge: Cambridge University Press.
- Blackburn, S. 1993. *Essays in Quasi-Realism*. New York: Oxford University Press.
- BonJour, L. 1998. *In Defence of Pure Reason: A Rationalist Account of A Priori Justification*. Cambridge: Cambridge University Press.
- Brown, J.R. 1991. *The Laboratory of the Mind. Thought Experiments in the Natural Sciences*. London: Routledge.
- Brouwer, L.E.J. 1949. "Consciousness, Philosophy, and Mathematics." In *Proceedings of the Tenth International Congress of Philosophy*, vol. 3, 1235–1249. Amsterdam: North Holland Publishing Company. Excerpts repr. in Benacerraf and Putnam 1983, 90–96.
- Carnap, R. 1934. *Logische Syntax der Sprache*. Vienna: Springer. Translated as *The Logical Syntax of Language*. London: Kegan Paul Trench, Trubner & Co., 1937.
- Dummett, M. 1991. *The Logical Basis of Metaphysics*. Cambridge, Mass.: Harvard University Press.
- Field, H. 1980. *Science without Numbers. A Defence of Nominalism*. Princeton, N.J.: Princeton University Press
- . 1989. *Realism, Mathematics, and Modality*. Oxford: Basil Blackwell.
- Gibbard, A. 1990. *Wise Choices, Apt Feelings*. Cambridge, Mass.: Harvard University Press.
- Gödel, K. 1944. "Russell's Mathematical Logic." In *The Philosophy of Bertrand Russell*, edited by P.A. Schilpp, 125–153. Evanston: Northwestern University Press. Repr. in Benacerraf and Putnam 1983, 447–469.
- Grover, D., J. Camp, and N. Belnap, 1975. "A Prosentential Theory of Truth." *Philosophical Studies* 27: 73–125.
- Hale, B. 1987. *Abstract Objects*. Oxford: Basil Blackwell.
- Hare, R.M. 1952. *The Language of Morals*. Oxford: Clarendon Press.
- Hellman, G. 1989. *Mathematics without Numbers*. Oxford: Clarendon Press.
- Hilbert, D. 1925. "Über das Unendliche." *Mathematische Annalen* 95: 161–190. Translated as "On the Infinite," in van Heijenoort 1967, 367–392. Partially translated as "On the Infinite," in Benacerraf and Putnam 1983, 183–201.
- Horwich, P. 1990. *Truth*. Oxford: Basil Blackwell.

- Katz, J. 1981. *Language and Other Abstract Objects*. Totowa, N.J.: Rowman and Littlefield.
- Kitcher, P. 1984. *The Nature of Mathematical Knowledge*. New York: Oxford University Press.
- Lewis, D. 1986. *On the Plurality of Worlds*. Oxford: Basil Blackwell.
- Mackie, J.L. 1977. *Ethics: Inventing Right and Wrong*. New York: Viking Press.
- McGee, V. 1997. "How We Learn Mathematical Language." *Philosophical Review* 106: 35–68.
- Mill, J.S. 1843. *A System of Logic*. London: Longmans, Green, and Company.
- Peacocke, C. 2005. "The *A Priori*." In *The Oxford Handbook of Contemporary Philosophy*, edited by F. Jackson and M. Smith, 739–763. New York: Oxford University Press.
- Putnam, H. 1967. "Mathematics without Foundations." *Journal of Philosophy* 64: 5–22. Repr. in Benacerraf and Putnam 1983, 295–311.
- . 1981. *Reason, Truth and History*. Cambridge: Cambridge University Press.
- Quine, W.V. 1948. "On What There Is." *Review of Metaphysics* 2: 21–38. Repr. in Quine 1980, 1–19.
- . 1951. "Two Dogmas of Empiricism." *Philosophical Review* 60: 20–43. Repr. in Quine 1980, 20–46.
- . 1970. *Philosophy of Logic*. Englewood Cliffs, N.J.: Prentice-Hall.
- . 1980. *From a Logical Point of View*. 2nd ed. Cambridge, Mass.: Harvard University Press.
- Ramsey, F. 1927. "Facts and Propositions." *Proceedings of the Aristotelian Society, Supplementary Vol. 7*: 153–170.
- Resnik, M. 1997. *Mathematics as a Science of Patterns*. New York: Oxford University Press.
- Shapiro, S. 1997. *Philosophy of Mathematics: Structure and Ontology*. New York: Oxford University Press.
- Stevenson, C.L. 1944. *Ethics and Language*. New Haven: Yale University Press.
- Strawson, P. 1950. "Truth." *Proceedings of the Aristotelian Society, Supplementary Vol. 24*: 129–156.
- van Heijenoort, J., ed. 1967. *From Frege to Gödel: A Source Book in Mathematical Logic, 1897–1931*. Cambridge, Mass.: Harvard University Press.
- Wedgwood, R. 2007. *The Nature of Normativity*. Oxford: Clarendon Press.
- Wright, C. 1983. *Frege's Conception of Numbers as Objects*. Aberdeen: Aberdeen University Press.

