

Meaning and Justification: The Case of Modus Ponens*

JOSHUA SCHECHTER
Brown University

DAVID ENOCH
The Hebrew University

In virtue of what are we justified in employing the rule of inference Modus Ponens? One tempting approach to answering this question is to claim that we are justified in employing Modus Ponens purely in virtue of facts concerning meaning or concept-possession. In this paper, we argue that such meaning-based accounts cannot be accepted as the fundamental account of our justification.

1. Introduction

Consider the rule of inference known as Modus Ponens (MP):

(MP) From the beliefs p and *if p then q* , infer q .

We reason by employing MP, or a rule of inference much like it. We are justified in so reasoning. And we are often justified in believing q upon inferring it from the justified beliefs p and *if p then q* in accord with the rule. In virtue of what are we justified in employing this rule of inference? In virtue of what does the rule MP transmit justification?

It may seem that an answer to these questions is neither possible nor necessary. After all, MP is a plausible candidate for being a basic rule of inference, one of the fundamental rules we employ in reasoning. (In what follows we assume that it is a basic rule of inference.) Justifications must come to an end somewhere, and a basic rule of inference provides a plausible place for such an end.

* This is the penultimate draft of Joshua Schechter and David Enoch. (2006) "Meaning and Justification: The Case of Modus Ponens" *Noûs* 40 (4), 687–715. The definitive version is available at www.blackwell-synergy.com.

But this line of thought is not satisfactory. True, thinkers cannot be required, in order to be justified in holding a belief or employing a rule of inference, to possess a (non-circular) justification for the belief or rule, and a justification for every belief and rule featuring in that justification, and so on *ad infinitum*. Nevertheless, this does not show that an account of the justification for employing MP is unneeded. Even if possessing an account of the justification of MP is not required for *thinkers* to employ the rule justifiably, it is required for *theorists* to have an adequate account of justification. Here is a quick way of seeing that this is so: There are many different possible basic rules of inference that could be employed. Some, such as MP and versions of Inference to the Best Explanation (IBE), we presumably are justified in employing. Others, such as Affirming the Consequent and Inference to the Third Worst Explanation, we presumably would be unjustified in employing. Surely, some account is owed of why the rules in the first class are justified for us to employ and the rules in the second are not. Even if one were to think that the former rules are somehow “default justified” – not in need of positive justification so long as no compelling case can be made against their use – still some principled distinction must be drawn between rules that are default justified and rules that are not. It is highly implausible, after all, that the justificatory status of each rule is a matter of brute fact, and that there is nothing general that can be said concerning why it is that we are justified in employing certain rules and not others. In particular, then, an account is owed of our justification for employing MP.

One tempting approach to providing such an account is to claim that we are justified in employing MP in virtue of facts about meaning or concept-possession.¹ The

¹ We do not sharply distinguish between meanings and concepts. There may be good theoretical reasons to avoid making this identification in general, but they are not, we think, important for our purposes.

idea is that we are justified in inferring q from p and *if p then q* purely in virtue of the nature of the concept *the conditional*. In this paper, we argue that this strategy fails. We do not argue – not here, that is – that meaning-based accounts of justification are extensionally inadequate. Instead, we argue that meaning-based accounts are explanatorily deficient. They require substantial supplementation if they are to do the necessary explanatory work, and even so supplemented they cannot be understood as providing the fundamental explanation of our justification.

Our discussion proceeds as follows: In the next section, we present a few preliminary clarifications. In section 3, we characterize and present meaning-based accounts in general, and Paul Boghossian's recent meaning-based account in particular. We take Boghossian's view as our exemplar of meaning-based accounts because it is the best developed such account of which we are aware, and because there is reason to think that for a meaning-based account to be at all plausible, it must take a shape very similar to Boghossian's. In section 4, we provide our arguments against meaning-based accounts. Section 5 is devoted to arguing that the motivations that have been presented for meaning-based accounts are unpersuasive. We conclude in section 6 by sketching an alternative account of our justification for employing MP, one that ties the justification of a basic belief-forming method to its indispensability for some rationally obligatory project. We suggest that this account – which we explicate and defend more fully elsewhere² – does not have the explanatory shortcomings of pure meaning-based accounts. Indeed, we argue that it is a promising candidate for the fundamental account of our justification, irrespective of whether considerations of meaning are of some relevance to justification.

2. Preliminaries

The central question of this paper is in virtue of what are we – normal adult human thinkers – justified in employing MP. Before assessing the prospects of meaning-based accounts to answer this question, a few preliminary clarifications are in order.

First, it is important to recognize that MP is not the claim that q follows from p and *if p then q* . That claim concerns implication, or what entails what. In contrast, MP is a rule of inference. It licenses patterns of reasoning, the formation of certain beliefs on the basis of antecedently held beliefs. There is undoubtedly some connection between the justified employment of MP and the truth of the statement about implication, but the two are conceptually distinct.³

Second, MP, as stated above, is much simpler than the similar rule or rules of inference that we actually employ. Often mentioned complexities include the fact that sometimes we do not (and ought not) conclude q on the basis of holding the beliefs p and *if p then q* , but rather reassess our commitment to those antecedent beliefs; the fact that we do not (and ought not) believe all of the trivial deductive consequences of what we believe; and the fact that belief is not all-or-nothing but comes in degrees. For simplicity, these complications will be ignored for the remainder of this paper. We believe that our discussion applies, *mutatis mutandis*, to the MP-like rules we actually employ.

MP is plausibly taken to be a basic belief-forming method for ordinary thinkers. That is, it is a belief-forming method that is not employed (by the relevant thinkers) on

² See Enoch and Schechter (forthcoming).

³ For more on the distinction between logic and inference, see Goldman (1986) and Harman (1995).

the basis of other belief-forming methods.⁴ Additional plausible candidates for basic rules of inference include other deductive rules – such as Disjunction Introduction and Conjunction Elimination – as well as ampliative rules of inference – perhaps including forms of Enumerative Induction and Inference to the Best Explanation. Moreover, certain basic belief-forming methods are not rules of inference. These plausibly include methods governing reliance on perceptual faculties, on memory, on introspection, on moral and other normative intuitions, and on intuitions of conceivability and possibility. It will prove important below that our central question generalizes to all basic methods.

The primary notion of justification relevant to our discussion is that of epistemic justification, not pragmatic or moral justification.⁵ This is the sort of justification that is necessary for knowledge, that is typically grounded in evidence, and that typically requires epistemic reasons for belief. (This excludes, for instance, the kind of justification you may have when someone reliably promises you a lot of money to believe a proposition.) Our central question, however, does not directly concern the justification of beliefs, but the justification of a particular rule of inference. Yet, epistemic justification as applied to beliefs and epistemic justification as applied to rules are closely related notions. Roughly, a thinker is justified in employing a rule of inference just in case the beliefs formed by an application of the rule to justified beliefs are *prima facie* justified for the thinker to hold, and are justified, in part, because so formed. Given this connection between the justification of rules of inference and the justification of beliefs,

⁴ Our notion of a basic belief-forming method is closely related to the notion of a basic psychological process in Goldman (1986). Burge (1993) proposes the view that reliance on memory and testimony are justificatorily basic. Pryor (2000) proposes that relying on perception is justificatorily basic. These views fit well with – though they do not entail – the claim that the relevant methods are basic in our sense.

⁵ We use “justification” and its cognates broadly. Certain justified beliefs may be justified in virtue of the relevant thinker possessing an explicit justification. Other justified beliefs may not be justified in this way.

there is no difficulty applying the distinction between epistemic and pragmatic justification to rules of inference as well.

One final clarification will prove to be important. In providing an account of what it is in virtue of which we are justified in employing MP, it is not sufficient merely to present necessary and sufficient conditions for justification. Rather, what is owed is an *explanation* of our justification. The account must provide a way of seeing *why it is* that we are epistemically justified in employing MP. It must show the employment of MP in a rationally positive light.⁶

3. Meaning-Based Accounts of Justification

In answering our question – in virtue of what are we justified in employing MP – it is tempting to appeal to considerations of meaning. If someone refused to acknowledge that we are justified in employing MP, this would be good evidence that they failed to understand what was meant by “if ... then”. Moreover, it is difficult to see what could justify inferring q from p and *if p then q* if not facts about the concept *the conditional*. Indeed, our justification for employing MP seems to be *a priori* – epistemically independent from how we find the world to be. It is difficult to envision what *a priori* facts, other than those concerning meanings, could ground the justification for employing the rule.

Of course, MP is not an isolated case. It is natural to think that the rules of Conjunction Elimination are justified in virtue of the nature of the concept *conjunction*

(On a popular way of talking due to Burge (1993), these include beliefs that the thinker is merely *entitled* to hold.)

⁶ The need for an explanation of our justification is endorsed by Boghossian (2000) and Peacocke (2004). See Horwich (2005, p. 169) for an opposing view.

and that the belief that bachelors are unmarried is justified in virtue of the nature of the concept *bachelor*. This suggests that there is a connection between meaning and justification in a wide range of cases.

Meaning-based accounts take such intuitive considerations seriously. They claim that epistemic justification can be had, under some circumstances, simply in virtue of possessing certain concepts – in virtue, that is, of satisfying whatever constraints are necessary for a thinker to have thoughts involving those concepts. This strategy is committed to the following connection, somewhat roughly stated, between meaning and justification:

(Meaning-Justification Link) Merely by possessing certain concepts, a thinker is justified in employing certain belief-forming methods or in holding certain beliefs.⁷

This connection is *prima facie* plausible for the case of MP. It is natural to think that any thinker who possesses the concept *the conditional* is thereby justified in employing MP. By “meaning-based accounts”, then, we shall mean accounts of justification that subscribe to the above Meaning-Justification Link.⁸

Unsurprisingly, meaning-based accounts of justification have perennially been popular. Recently, they have undergone a resurgence.⁹ In contemporary versions,

⁷ The justification in question should be understood to be *prima facie* justification. Presumably, it can be defeated – overridden or perhaps undermined – by other considerations. This qualification should be taken to be implicit in what follows.

⁸ Notice that meaning-based accounts need not reduce epistemic justification to meaning-theoretic properties.

⁹ The logical positivists' analytic theory of the *a priori* is best interpreted as a meaning-based account combined with a conventionalist account of meaning. More recent meaning-based accounts can be found in

meaning-based accounts are typically tied to forms of Conceptual Role Semantics (CRS) – views according to which the meanings of a thinker’s mental representations are determined by their place in the thinker’s cognitive economy, including certain of the inferential relations in which they play a part. In a slogan, the meaning of a concept is its conceptual role.¹⁰

The adoption of CRS suggests the following elaboration of the Meaning-Justification Link:

(CRS Meaning-Justification Link) If a belief helps to constitute the conceptual role for a concept, any thinker possessing the concept is justified in holding the belief. If a belief-forming method helps to constitute the conceptual role for a concept, any thinker possessing the concept is justified in employing the method.¹¹

In other words, if being disposed to hold a certain belief or to employ a certain method is necessary for a thinker to have thoughts involving a certain concept, then any thinker possessing the concept is thereby justified in holding the belief or employing the method.

This claim – henceforth, The Link – is very attractive. It seems to provide a principled distinction between belief-forming methods that are justified for a thinker to

Dummett (1978) and (1991); Peacocke (1992), (1993), and (2004); Boghossian (2000), (2001), and (2003); and the papers collected in Hale and Wright (2001), particularly Wright (1999) and Hale and Wright (2000).

¹⁰ Conceptual Role Semantics is not a single view, but a heterogeneous family of views. See Block (2000) for an overview. For the purposes of our discussion, distinctions among versions of CRS can safely be ignored.

¹¹ This principle is related to what Peacocke (1993) calls “the metasemantic account”, to what Boghossian in (2000) calls “Principle (L)” and in (2003) calls the “Meaning-Entitlement Connection”, and to what Hale and Wright (2000) call “the traditional connection”.

employ and those that are not. Moreover, assuming the adequacy of CRS, the methods we are intuitively justified in employing are precisely the methods that plausibly constitute acceptable conceptual roles for a wide range of cases.¹² For instance, on most versions of CRS as applied to logical constants, a rule resembling MP is constitutive of the concept *the conditional*. The rules of Conjunction Elimination are constitutive of the concept *conjunction*. And the claim *all bachelors are unmarried* is plausibly constitutive of the concept *bachelor*. In contrast, the rule of Affirming the Consequent and the belief *all bachelors are handsome* plausibly do not help constitute acceptable conceptual roles.

But there is an immediate problem. As Boghossian (2003) notes, there are many apparent counterexamples to The Link – concepts with constitutive methods that thinkers are not even *prima facie* justified in employing. Here are the four examples Boghossian discusses:¹³

(1) Tonk

“Tonk” is stipulated to have the following associated conceptual role:

From p , infer p tonk q ; from p tonk q infer q .

(2) Boche

The meaning of “boche” is plausibly constituted by inference rules such

as:

From x is *German* infer x is *boche*; from x is *boche* infer x is *cruel*.

¹² For discussion of general difficulties for CRS, see Fodor and Lepore (1991), Block (2000), and Boghossian (2000). Williamson (2003) presents several difficulties for CRS aimed specifically at its place in meaning-based accounts.

(3) Aqua

“Aqua” is stipulated to have the following conceptual role:

From *x is water* infer *x is aqua*; from *x is aqua* infer *x is H₂O*.

(4) Flurg

“Flurg” is stipulated to have the following conceptual role:

From *x is an elliptical equation* infer *x is flurg*; from *x is flurg* infer *x can be correlated with a modular form*.

It is intuitively clear that even if the rules stated constitute conceptual roles for a thinker, this does not, by itself, justify the thinker in inferring an arbitrary *q* from an arbitrary *p* (via *p tonk q*), in inferring *Hans is cruel* from *Hans is German* (via *Hans is boche*), in inferring *sample s is H₂O* from *sample s is water* (via *sample s is aqua*), or in inferring *equation e can be correlated with a modular form* from *equation e is an elliptical equation* (via *equation e is flurg*). These rules do not gain any justificatory respectability merely from their apparent roles in constituting meanings. Of course, a thinker may have other ways of justifying these inferences, and in the latter two cases such justifications are now available.¹⁴ But simply introducing the terms “aqua” and “flurg” with their associated rules of inferences is insufficient for being justified in employing these rules.

¹³ The rules for “tonk” were first introduced in Prior (1960). The example of “boche” was first used by Dummett (1973, p. 454).

¹⁴ The Taniyama-Shimura Conjecture states that all elliptical curves can be correlated with modular forms. An important subcase was proved by Wiles on his way to proving Fermat’s Last Theorem. The conjecture was proved in full generality 1999 by Breuil, Conrad, Diamond, and Taylor, building on earlier work of Wiles and Taylor.

Had it been, science and mathematics would have been much easier than they actually are.

Clearly, then, if we are to rule out these examples, some qualification is in order. As Boghossian indicates, two main strategies can be pursued. First, certain apparent concepts could be ruled out as spurious.¹⁵ Second, The Link could be restricted so that a thinker is justified in employing only certain of the belief-forming methods that help constitute genuine concepts. (Of course, these strategies are not exclusive.)

The first strategy is somewhat attractive for the case of “tonk”. Perhaps no one can have tonk-ish thoughts, because “tonk” fails to stand for a genuine concept. (Perhaps this is so because its purported constitutive rules are inconsistent, or because adding the rules to a consistent language does not yield a conservative extension of that language, or because of some other reason.¹⁶) But it is difficult to see how this strategy can be made at all plausible for the other three examples. Do racist statements employing words such as “boche” fail to have senses? Has no one ever had a derogatory thought employing such a concept?¹⁷ And the point generalizes to other “thick concepts”¹⁸ that involve false evaluative judgments: Has no one ever believed that someone was chic, geeky, blasphemous, or chaste? Moreover, there seems no non-question-begging reason to think that “aqua” and “flurg” fail to stand for genuine concepts. Indeed, now that we know the

¹⁵ A related strategy is to argue that CRS does not apply to the ranges of concepts into which some of these examples fit. This strategy is problematic for a number of reasons, perhaps most importantly the troubling lack of uniformity it introduces into semantics, the lack of an obvious principled basis for such a restriction, and the fact that restricting the scope of CRS substantially reduces the explanatory power of The Link.

¹⁶ The first line is adopted by Stevenson (1961) and Peacocke (1987). The second line is endorsed by Belnap (1961). We are skeptical of both. Consider the term “naïve set” with the conceptual role provided by the full comprehension schema. It is plausible that this term stands for a genuine concept – for instance, we attribute beliefs about naïve sets to the early developers of set theory – although its conceptual role is neither consistent nor conservative. See Hale and Wright (2000, pp. 301–2) for discussion of a related example.

¹⁷ See Brandom (2000, p. 70) and Boghossian (2003) for related points.

¹⁸ The term “thick concept” is due to Williams (1985).

relevant facts about water and elliptical equations, there seems to be nothing that prevents us from introducing these terms into our language, and succeeding in using them to stand for genuine concepts. The suggestion that none of these apparent concepts is genuine is a non-starter. It would be committed to far too much content failure.¹⁹

So much, then, for the strategy of ruling out certain purported concepts as spurious. Let us now consider ways to adopt the second strategy, that of restricting The Link.

One natural idea is to restrict The Link to those meaning-constituting rules of inference that are truth-preserving, or perhaps to those that necessarily are truth-preserving. But this restriction is too weak, for the rules stipulated to be meaning-constituting of “aqua” and “flurg” necessarily are truth-preserving. Indeed, restricting The Link further to methods that *a priori* are truth-preserving is insufficient, for the rules stipulated to be meaning-constituting of “flurg” *a priori* preserve truth.²⁰

There is reason to think that the “flurg” example can be addressed if we restrict The Link to meaning-constituting methods that are, in some intuitive sense, *directly a*

¹⁹ Williamson (2003, section II) argues that the pejorative nature of “boche” and the like is a matter of their pragmatic, not semantic, features, and that in fact the sense and reference of “boche” is the same as that of “German”. If this is so, then the meaning of “boche” does not involve one in a commitment to unacceptable inferences (such as that from *Hans is German* to *Hans is cruel*), and so “boche” does not provide a clear counterexample to the unrestricted Link. Now, there are reasons to doubt that the pejorative nature of “boche” is a matter of its pragmatic features. And even if it is, there are reasons to doubt that appealing to the pragmatics/semantics distinction can save the unrestricted Link from the counterexample. But whatever the status of “boche”, the examples of “aqua” and “flurg” suffice to discredit the strategy of ruling out purported counterexamples to The Link as involving spurious concepts.

Later in his paper, Williamson (2003, section III) argues that Boghossian’s stipulated concepts – *aqua*, *flurg*, and the like – are not genuine concepts. But his principal arguments for this claim depend on general objections to CRS, and need not concern us here.

Horwich (2005, pp. 153–4) argues that the rules listed for *boche* and *flurg* cannot be followed as cognitively basic rules, since doing so would violate our basic rules for other concepts, such as *cruel*. Therefore the rules do not constitute genuine concepts. We do not see why accepting the rules for *boche* is incompatible with accepting our rules for *cruel*.

²⁰ This point is due to Boghossian (2003).

priori.²¹ Surely, the meaning-constituting methods for “flurg” are not directly *a priori* for any thinker. But this suggestion is unsuccessful, too. The notion of a directly *a priori* belief-forming method is dangerously unclear. If directness is meant to be a psychological notion, so that a directly *a priori* method is an *a priori* truth-preserving method that a thinker is disposed to follow and is not so disposed on the basis of any more fundamental considerations, then the restriction is insufficient. It is certainly possible for there to be thinkers for whom the meaning-constituting methods of *flurg* are directly *a priori* in this sense. Such thinkers will not be justified in following the inference rules constitutive of *flurg*. (Further restricting The Link to those methods that we are *actually* disposed to follow would be hopelessly *ad hoc*.) Perhaps, then, directness should be understood as a normative, and not a psychological, notion. For instance, a directly *a priori* method might be an *a priori* truth-preserving method that is epistemically permissible for us to adopt as basic. In this case, The Link (so restricted) may achieve extensional adequacy, but only at the price of being dangerously close to vacuity. The account fails to supply a satisfactory explanation of our justification, but merely replaces the need for an explanation of one normative epistemological status with the need for an explanation of another, closely related, one. Perhaps, instead, directness should be understood as a metaphysical notion. For instance, a directly *a priori* method might be an *a priori* truth-preserving method that is metaphysically haloed in some way. In this case, a fuller and more precise account of the directly *a priori* is badly needed. It is reasonable to suspect that no such metaphysical account is possible that preserves the intuitive plausibility of the restriction to the directly *a priori*.

²¹ This suggestion loosely follows Peacocke’s (1992, p. 6) restriction to inference rules that are “primitively compelling” for a thinker.

What if we restrict The Link to logical concepts, such as *the conditional*?²² This restriction is subject to the well-known difficulty of providing a principled demarcation of the logical.²³ Moreover, with very few exceptions, available accounts of logical concepthood seem not to be of epistemic relevance.²⁴ An additional problem is that for many accounts of the logical, there are logical concepts resembling *flurg* that serve as clear counterexamples to The Link.²⁵ Finally, and most importantly, as we will argue below, there is the need for a unified account of our justification – one that solves the problem of justifying basic belief-forming methods across a wide range of cases. The narrower the restriction, the more pressing the objection from the need for a unified account becomes, and a restriction to logical concepts is a very narrow restriction indeed.

Many other restrictions of The Link can undoubtedly be thought of. What is needed, however, is not some unmotivated qualification introduced purely to deal with counterexamples, but a principled way of addressing them. And to be principled in this way, it seems that the restriction must be connected with the central normative notions of epistemology, such as those of justification, permissibility and obligation, blame and responsibility. The natural idea is that certain concepts have an epistemically privileged status, and thinkers are justified in employing the meaning-constitutive methods of such concepts.

²² We thank an anonymous referee for pressing us on our response to this view.

²³ See Gómez-Torrente (2002) for an overview of the difficulties in providing a principled demarcation of the logical.

²⁴ Tarskian accounts of logical concepthood, which involve notions such as that of permutation invariance, seem to have little relevance to epistemic justification. Accounts that restrict logical concepts to those concepts that are general or ubiquitous in thought face a similar problem. Accounts that make use of epistemic notions, such as that of *a priori* justification, threaten to collapse the proposed explanation of justification into vacuity. Accounts of logical concepthood that involve a conservative extension requirement are subject to the problems presented in the discussion of conservativeness below.

²⁵ Consider, for instance, the one-place operator “!”, whose conceptual role consists of the following two rules: From *p*, infer *!p*; From *!p*, infer *p* and *C*, where *C* is a specific highly complex logical truth.

A possible line of thought is that some concepts are, in some sense, epistemically harmless for thinkers to add to their conceptual repertoire, and The Link ought to be restricted to such harmless concepts. One way to flesh out this idea is to restrict The Link to concepts that are conservative, in the sense that their meaning-constituting rules conservatively extend the appropriate background derivability relation.²⁶ But such a restriction is subject to several difficulties. First, considerations of conservativeness only apply to a concept given some inferential background; whether a rule is a conservative extension can only be evaluated with respect to some presupposed derivability relation. What should be taken to be in the inferential background? Many possible answers can be envisioned, but each seems subject to the same general problem: Whether a concept is conservative may turn out to be a function of features of the inferential background that are, it seems, irrelevant to its epistemic status. Second, the restriction seems too liberal. There are examples of concepts that are each conservative over a given background, but that are jointly inconsistent. Which of these concepts should we take to have meaning-constituting rules that are vindicated by The Link? Third, the restriction seems too severe. The meaning-constituting rules for (classical) *negation* are not conservative over those for *the conditional* and *disjunction*. The meaning-constituting methods for many of our observational and theoretical concepts plausibly also are not conservative. This, of course, is not a problem with these concepts, but a strength; possessing them allows us to

²⁶ Conservativeness is usually presented as a constraint on which conceptual roles stand for genuine concepts, rather than as a restriction on which concepts give rise to the justified employment of their constitutive methods. That view was implicitly addressed in the discussion of why Boghossian's counterexamples to The unrestricted Link are not spurious. The proposal to impose a conservativeness requirement stems from Belnap's (1961) suggestion that a conceptual role only bestows a genuine meaning on a logical constant if it is conservative. Dummett (1973, p. 454) suggests that the problem with the concept *boche* is that its meaning-constituting rules do not satisfy a conservative extension requirement. Hacking (1979) demarcates the logical, in part, by appealing to conservativeness.

develop increasingly powerful sets of beliefs about the world.²⁷ If we think that the adoption of these concepts gives rise to the justified employment of their constitutive belief-forming methods, then restricting The Link to conservative concepts rules out too much.²⁸ Finally, the restriction to conservative concepts has difficulty handling the example of *flurg*, since *flurg* is conservative over the derivability relation generated by the concepts of classical mathematics.²⁹

Let us turn, then, to what will be our main target in the rest of this paper. This is the meaning-based account developed in a series of papers by Boghossian, culminating in Boghossian (2003). His account is also an attempt to capture the intuition that some concepts are epistemically harmless, and that we are justified in employing the methods constituting such harmless concepts.

Boghossian introduces his view by considering the case of *flurg*. Suppose we're trying to decide whether to use the term "flurg", and are unaware of the fact that all elliptical equations can be correlated with modular forms. Introducing a concept by stipulating the conceptual role for "flurg" described above would be epistemically blameworthy. As far as we know, there may be some elliptical equations that cannot be

²⁷ This point is due to Brandom (2000, pp. 70–72).

²⁸ Proponents of a conservativeness requirement may reply that, for all that has been said, the methods constituting the meaning of such concepts may be vindicated by some other non-meaning-based considerations, so there is no problem with their violation of conservativeness. But meaning-based theorists can only endorse such a reply at the price of restricting the scope of their meaning-based account, thereby making it more vulnerable to the objection from the need for a unified account we develop in section 4.2.

²⁹ An alternative way to try to flesh out the idea that certain concepts are epistemically harmless is to appeal to what Dummett (1991) calls "intrinsic harmony". Very roughly speaking, a concept is intrinsically harmonious if the methods that specify how it is to be applied match the methods that specify what is entailed by applications of the concept in some appropriate way. (This idea stems from the Gentzen-Prawitz tradition of characterizing the logical constants using constraints on their introduction and elimination rules.) However, the detailed accounts of intrinsic harmony that have been developed only apply to logical concepts, and it is difficult to see how they may be generalized much further. Moreover, restricting The Link to just those concepts that are intrinsically harmonious may very well suffer – depending on the exact details of such an account – from the same difficulties as apply to a restriction to

correlated with a modular form, and so, as far as we know, introducing a concept with the described conceptual role may lead us astray. Surely, the responsible thing to do would be to be more tentative in introducing the concept. One should not adopt the inference rules licensed by the conceptual role of “flurg” unless one is willing to be committed to the claim that all elliptical equations can be correlated with modular forms. (And, even if one were so committed, Boghossian would add, introducing a concept that forecloses on the possibility that the commitment is false is itself epistemically irresponsible.) Hence, when we introduce a conceptual role for the term “flurg”, we should do so by conditionalizing on the claim that the given inference rules are truth-preserving. Namely, we should introduce a conceptual role for “flurg” with the following two rules of inference: from *there is a property P such that (e is an elliptical equation entails e is P, e is P entails e can be correlated with a modular form, and x is P)* infer *x is flurg*, and the converse inference.³⁰

Now, not every concept can be introduced in such a conditional form.³¹ *The conditional*, for instance, cannot be so introduced, because in order to conditionalize it is necessary to possess the concept of the conditional. A thinker is therefore not epistemically blameworthy for introducing *the conditional* without conditionalizing, because introducing it conditionally is impossible. A thinker is epistemically blameworthy for introducing a concept without conditionalizing if it is possible to

conservative concepts. For sympathetic general discussions of harmony and its relation to justification, see Tennant (1987) and Dummett (1991).

³⁰ There are subtle nuances concerning exactly how the conditionalization ought to go. There are also worries that the concepts ordinary thinkers typically possess are not conditionalized in anything like the way envisaged by Boghossian, even under the appropriate idealization. See Horwich (2000) and (2005) and Williamson (2003) for relevant discussion. There are interesting connections – ones we cannot discuss here – between Boghossian’s view and Brandom’s claim that the role of logical concepts is to make explicit a thinker’s implicit inferential commitments.

conditionalize and thus avoid a commitment the thinker is not entitled to undertake.

Concepts that can be introduced conditionally ought to be so introduced.

When, then, is a thinker justified in employing a belief-forming method – such as MP – purely on the basis of possessing a concept? Boghossian’s answer is that a thinker is so justified only if the method is meaning-constituting of a concept that the thinker is epistemically blameless in possessing. A thinker is epistemically blameless in possessing a concept if the concept was either introduced – explicitly or implicitly – in the tentative, conditional way described above, or is a concept that could not have been so introduced.³²

Notice that Boghossian’s view nicely deals with the four examples: All four conceptual roles relevant to the examples could have been introduced conditionally, but were not, and so their meaning-constituting rules of inference are not justified for thinkers to employ. Indeed, had the concepts been introduced conditionally, no problem would have ensued. Thinkers possessing the conditional version of *flurg*, for instance, would only have been justified in believing that a given elliptical equation can be correlated with a modular form if they had been antecedently justified in believing that there was a property that satisfied the unconditional inference rules; if, that is, they had already justifiably believed that all elliptical equations could be correlated with modular forms.³³

³¹ For ease of exposition, we sometimes write as if the unconditional concept and its conditional variant are actually the same concept.

³² It is not clear if Boghossian thinks of this as a necessary and sufficient condition or a merely sufficient condition for blameless concept-possession. In either case, presumably, a “no underminers” clause must be added to handle the case of thinkers who justifiably but falsely believe that a concept is somehow defective.

³³ Notice that Boghossian’s view nicely explains our ability to understand thinkers with racist and other problematic unconditional concepts. We can use our conditional concepts to interpret their beliefs involving the relevant unconditional concepts. Notice, too, that to use Boghossian’s account of conditional concepts for this purpose, it is not necessary to endorse his claims about justification.

Notice, too, that Boghossian's view seems to provide a satisfying explanation of our justification for employing MP. There are certain concepts that we are epistemically blameless in possessing. Doing anything necessary to possess such a concept is itself blameless. *The conditional* is a blameless concept for us to possess. Employing MP is needed to possess *the conditional*. Hence, employing MP is epistemically blameless. We are therefore justified in employing MP.³⁴

4. The Explanatory Deficiencies of Meaning-Based Accounts

Despite the attractiveness of meaning-based accounts of justification, particularly ones resembling Boghossian's version, meaning-based accounts of the justification of MP cannot be understood as providing the fundamental explanation of our justification.³⁵ In this section, we provide three reasons to accept this claim – one based on the need to restrict The Link if it is to be extensionally correct, another based on the need to provide a unified account of the justified employment of basic belief-forming methods, and a third grounded in the fact that adequate accounts should explain thinkers' epistemic obligations as well as their epistemic permissions. These considerations suggest that a more fundamental account of our justification for employing MP is needed. It is this more fundamental account that does the real explanatory work.³⁶

³⁴ Boghossian's discussion differs in a few ways from our restatement. He frames the debate as concerning the conditions under which an inference transfers justification. He does not directly consider the question of the justification for employing the relevant rules. He also only considers rules of inference, not belief-forming methods more generally. The first difference is not important for our purposes here. The importance of the second will be made clear below.

³⁵ In a series of papers, Horwich has argued against meaning-based accounts. See Horwich (1997), (2000), and (2005). Also see Williamson (2003). Our thinking on these matters has benefited from their discussions.

4.1. *Restricting The Link*

Let us grant, provisionally, that some restricted version of The Link provides a sufficient condition for the justifiable employment of belief-forming methods, and so some meaning-based account can provide a true generalization about justification. But, as already emphasized, this is not enough. What we are after is an *explanation* of why it is that we are justified in employing MP and not, say, Affirming the Consequent. We now want to argue that merely appealing to a theory of concept-possession and The Link, suitably restricted, cannot provide the fundamental explanation of our justification.

It is important to be clear regarding the role of The Link in meaning-based accounts. Theories of concept-possession are entirely descriptive.³⁷ They provide necessary and sufficient conditions for thinkers to be able to think certain thoughts. In contrast, a theory of justification must be thoroughly normative. The theory of concept-possession therefore cannot play the central normative role in the explanation of our justification for employing MP. It can only play a role analogous to that of a minor – descriptive – premise in a practical syllogism. General normative principles must be provided to bridge the *is/ought* gap. This is the work The Link is intended to perform.

The Link, then, is intended to help explain the justified employment of certain belief-forming methods by connecting justification to meaning-constitution. Can The Link be thought of as a normative principle featuring in the fundamental explanation of this justification? If an unrestricted version of The Link were acceptable, perhaps it could be so considered. But, as we have seen, an unrestricted version of The Link is not

³⁶ Horwich (2005, pp. 156 and 162) suggests that meaning-constitution cannot do any real explanatory work.

³⁷ Even if content is normative – whatever exactly that may mean – its normativity does not seem to undermine the point in the text.

acceptable. And a restricted version of The Link cannot plausibly be taken to be the fundamental normative principle at work. Since not every instance of meaning-constitution gives rise to justification, that some instances do is certainly a fact that ought to be explained. Such an explanation will have to rely on additional – and more basic – normative principles.

Examining Boghossian's view helps to buttress this point. On his account, a thinker is justified in employing a belief-forming method that constitutes a concept if it is either conditional or unconditional and unconditionalizeable. But why is this so? Boghossian suggests the following: Conditional concepts do not foreclose on factual possibilities, and hence are blamelessly possessed. A thinker cannot be blamed for not employing the conditional form of an unconditionalizeable concept, and hence, unconditionalizeable concepts are blamelessly possessed as well. Anything needed for behaving blamelessly is itself blameless. Blameless belief-forming methods are justified.

We see, then, that Boghossian's view implicitly relies on several additional normative principles: a principle relating foreclosing on inquiry with blameworthiness, an ought-implies-can principle relating the inability to act otherwise with blamelessness, and a principle relating blamelessness and justification. It is only by relying on such principles that Boghossian's view is at all plausible. Indeed, Boghossian's version of The Link is only a derivative normative principle; it is derived from these more fundamental normative principles.

Moreover, there is no apparent reason why Boghossian's view should not be generalized beyond the case of meaning-constituting methods. The more basic normative principles Boghossian's view relies upon do not appeal to considerations of meaning.

Rather, considerations of meaning play the role analogous to that of the minor premise in a practical syllogism, determining how the general normative principles are to be applied to the cases at hand. And it seems highly unlikely that the only cases to which the general principles apply are ones in which the relevant methods are meaning-constituting. If they apply elsewhere, the crux of Boghossian's account can straightforwardly be generalized to other cases, where meaning-constitution is simply irrelevant. Even on Boghossian's view, then, considerations of meaning seem to play a relatively minor role in the explanation of our justification.

What does the fundamental normative work on Boghossian's view, then, is not his version of The Link, but rather meaning-independent normative principles concerning notions such as blame and justification. The same holds true of other meaning-based accounts. Indeed, our discussion can be made fully general. To be extensionally adequate, The Link must be restricted. To be well motivated, the restriction of The Link must be tightly connected with the notions of blame, responsibility, justification, or other central normative notions of epistemology. For a meaning-based account to be plausible, it must be supplemented with additional normative principles concerning some of these notions. And it is this separate apparatus – not grounded in considerations of meaning – that does the fundamental explanatory work.

4.2. *Uniformity*

Meaning-based accounts, even if successful, can only solve a few instances of a very general problem. The general problem is that of showing how thinkers are justified in employing basic rules of inference and other basic belief-forming methods. (The case of

non-basic methods can presumably be reduced to the case of basic ones.) One instance of this is the problem of explaining our justification for employing MP. But there are many other instances, concerning, perhaps, Conjunction Introduction and Conjunction Elimination; versions of Enumerative Induction and Inference to the Best Explanation; reliance on perceptual faculties, memory, introspection, moral and other normative intuitions, and intuitions of conceivability and possibility; and perhaps other cases as well. What explains our justification for employing these belief-forming methods?

In none of these cases is it plausible that thinkers possess explicit justifications for the employment of the relevant belief-forming method. Not only is the claim that we possess explicit justifications false to the phenomenology, but there is reason to think that no explicit non-circular justifications are even possible. Any justification for the belief that IBE is a reliable rule of inference, for example, would seem to require the use of an ampliative rule of inference such as IBE. Any explicit justification for belief in the reliability of memory plausibly requires reliance on the memory of successful past correspondence between remembered facts and the world. And so on.³⁸ So the problem of showing how we are justified in employing these methods is structurally very similar across all of them. This suggests that there should be a uniform solution for all of them, or at least that we ought not to despair too quickly of the possibility of finding such a uniform solution.

³⁸ Boghossian has provided us with the following objection: An explicit justification for a belief-forming method is an argument to the proposition that the method is justified, not to the proposition that the method is reliable. Assuming the denial of reliabilism, it is not clear that such arguments must be rule-circular, except in the case of MP. Our response is as follows: A justification of a belief is an argument to the belief, not to the claim that we are justified in believing it. Similarly, a justification of a belief-forming method is an argument to the proposition that the method is truth-preserving, or tends to be truth-preserving, or the like. It may be that such arguments can sometimes go by way of arguing to the claim that the belief or method is justified for us to adopt, such as on some construals of transcendental arguments and

However, no uniform account can be based on meaning. For the case of MP, there is a compelling candidate, *the conditional*, for the concept in virtue of which we are justified in employing the rule. In the case of certain perceptual methods, observational concepts may provide plausible candidates for concepts in virtue of which we are justified in employing them.³⁹ Yet, there are no analogous candidates for the other cases. In virtue of possessing which concept are we justified, for instance, in employing IBE? The only apparent candidate is the concept *explanation*. But it does not seem to follow from the meaning of “explanation” that we are justified in inferring to the best – simplest, most predictively powerful, most easily testable, etc. – explanation. A thinker could perfectly well possess the concept *explanation* without being disposed to infer to the best explanation. That is, IBE does not seem to be meaning-constituting of *explanation*. Moreover, it is plausible that a thinker can employ IBE – and be justified in so doing – without possessing the concept. Possessing the concept *explanation* would thus seem to be irrelevant to the justification of IBE.

A meaning-based account is even less plausible for other basic belief-forming methods. In virtue of which concepts are we justified in relying upon our memory, our normative intuitions, or our modal intuitions? The only obvious candidates – the concepts *past* and *remembrance*, the concept *experience*, the concepts *good* and *right*, the concepts *possible* and *conceivable*, and the like – are unlikely to be of use. Thinkers can possess all of these concepts without being disposed to employ the respective belief-forming

indispensability arguments. But these are rare cases, there are restrictions on when this strategy is appropriate, and it does not seem to be appropriate here.

³⁹ For instance, it might be maintained that it is in virtue of the constitution of the concept *red* that we are justified in following the method: (Red) If you have a perception as of something being red, believe that it is red. See Peacocke (1992, pp. 7–8) for discussion of a related view. Accounts resembling Boghossian’s and accounts imposing a conservativeness restriction have difficulty in allowing that a concept with a meaning-constituting method such as the one above is harmless in the relevant sense.

methods. And thinkers can be justified in employing the methods without possessing the respective concepts. This can be seen, for instance, from the fact that some children and animals are justified in employing certain of these methods – such as relying on memory and perception – without possessing the relevant concepts. The methods are in no sense constitutive of, or central to, or justified purely by grasp of, these concepts.

Meaning-based accounts thus cannot provide the needed unified explanation, and this counts heavily against taking them to provide the fundamental explanation of our justification even where they apply.⁴⁰ It would certainly be an odd state of affairs if the fundamental explanation of our justification for employing MP was based in meanings, the fundamental explanation of our justification for relying on IBE was, say, pragmatic, and the fundamental explanation of our justification for relying on memory involved still different considerations. At least until there is good reason to abandon the more ambitious – and intellectually more satisfying – project of coming up with a reasonably unified account of the justification of basic belief-forming methods, there is strong reason to reject meaning-based accounts as providing the fundamental story.

4.3. *Permission and Obligation*

Intuitively, in some circumstances it is epistemically permissible but not obligatory for a thinker to infer q from p and *if p then q* . For instance, the thinker might not care whether q is true, and it might be of no importance to her to correctly believe q . In such a case, the thinker would be justified in inferring q , but would not violate any epistemic requirement in not so inferring. In other circumstances, it seems epistemically obligatory for a thinker

to draw an MP inference. The thinker is required to draw the inference; withholding belief in the conclusion would not be permissible.⁴¹

Any complete explanation of our justification for employing MP ought to account for, or be easily extendible to account for, the epistemic permissibility and the epistemic requiredness of drawing the relevant MP inferences. And this leads to a difficulty for meaning-based accounts. Since justification is plausibly tied to epistemic permissibility, The Link may be able to explain the epistemic permissibility of drawing certain MP inferences. But it is difficult to see how The Link could be modified to account for our epistemic obligations. Additional – quite different – normative principles seem required. This provides further evidence that The Link and a theory of concept-possession are by themselves insufficient to fully explain what ought to be explained.⁴²

This difficulty arises in a sharp form for views that restrict The Link to concepts that are “harmless”, such as views imposing a conservativeness restriction and views resembling Boghossian’s. According to Boghossian’s view, possession of the concept *the conditional* is blameless, employing MP is necessary to possess this concept, and so we are justified in employing MP and in drawing MP inferences. Yet, if this is the essence of Boghossian’s view, notice how limited the conclusion must be. All that could follow from a thinker being blameless in possessing *the conditional* is that it is permissible to draw MP inferences. It does not follow that it is obligatory to draw any MP inference,

⁴⁰ Horwich (2005, pp. 158 and 162–3) argues that meaning-based accounts resembling Boghossian’s are insufficiently general since they cannot explain the epistemic status of MP, Hume’s Principle, and Enumerate Induction, among other rules.

⁴¹ There may be other, related, patterns of inference, so that the content of the relevant epistemic obligation may not be to draw the MP-inference, but rather to engage in one out of a number of related patterns of inference. The points in the text are not affected by this complication.

⁴² Horwich (2000, p. 161) argues that the typical motivations for meaning-based accounts only explain our right to accept meaning-constituting beliefs, and not our obligation to do so. He develops this objection in his (2005).

because as far as the view goes, not drawing the inferences may very well be blameless, too.⁴³

In reply to this difficulty, a meaning-based theorist may simply agree that her account does not show why we are epistemically required to draw certain MP inferences. She can argue that the account does make some progress; it shows how it is epistemically permissible to draw the inferences. Perhaps a very different account can be provided to answer the separate question why it is that particular inferences are epistemically obligatory. The mere fact that an account doesn't give all that is wanted shouldn't count as a reason for rejecting it altogether.

Indeed, there is no general flaw in achieving philosophical progress one step at a time. But this reply will not do here. What is intuitively clear is not just that it is epistemically permissible to draw certain MP inferences, and that it is also epistemically obligatory to draw certain MP inferences. What is intuitively clear, rather, is that – at least for MP inferences that are both permissible and obligatory – the reason why those inferences are permissible and the reason why they are obligatory are the very same reason, or that the two reasons are very tightly related. We can rule out in advance the possibility that an MP inference is permissible for one reason and required for a very different reason. At the very least, we can rule out the possibility that the fundamental explanation of our epistemic permission is very different from the fundamental explanation of our epistemic obligation. If meaning-based accounts cannot explain our

⁴³ Indeed, a more general point can be made against views that restrict The Link to concepts that are “harmless” in some way. Drawing an MP inference does not merely lack negative epistemic value, but rather is of positive epistemic value. Invoking the harmless nature of *the conditional* will not explain the positive value that these inferences have.

epistemic obligations, then they cannot be understood as providing the fundamental explanation of our justification.

A meaning-based theorist will want, then, to argue that her account can yield the epistemic requiredness of drawing certain MP inferences.⁴⁴ The natural thought for her to pursue is that what is intuitively clear is not that thinkers have an absolute obligation to draw certain MP inferences, but that they have the conditional obligation to draw certain MP inferences given that they possess the concept *the conditional*. This seems plausible, since a thinker who does not possess *the conditional* does not have the belief *if p then q* to start with. So perhaps we should think that a requirement to draw certain MP inferences is somehow built into the possession conditions of the concept *the conditional*. A part of what it takes to possess *the conditional* is to be required to draw certain MP inferences.⁴⁵

It is important to recognize that such a response relies on a very different view of content than what we hitherto have been assuming. Before, possessing *the conditional* was a matter of being *disposed* to reason in certain ways. According to the view now under consideration, possessing *the conditional* is in part a matter of being *required* to reason in certain ways (and, presumably, being permitted to reason in others). On this view, then, the possession of a concept is a normative matter, and this is already a

⁴⁴ An alternative reaction is to claim that the correct meaning-based account does not explain either the permissibility or the requiredness of drawing MP inferences. What is meaning-constituting of *the conditional* is not MP, but rather a rule like: (CON) Don't believe *p*, believe *if p then q*, and reject *q*. The meaning-based account explains our justification for employing this rule. Some very different account explains the permissibility and requiredness of drawing particular MP inferences. In response, we would like to make two points. First, this view is subject to the uniformity objection presented in section 4.2 in an extremely strong form. Second, intuitively, not only are we permitted to conform to CON, but we are also obliged to conform to CON, at least if we possess *the conditional*. The meaning-based account has difficulty in accounting for this obligation. We thank Paul Boghossian for pressing us on this issue.

⁴⁵ What if, instead, we import the obligation into the content of the rule? The claim would be that it is meaning-constituting of the conditional that we employ a rule like: Under conditions C, from the beliefs *p* and *if p then q*, it is obligatory to infer *q*. It is difficult to understand exactly what it is to employ a rule of this sort. Moreover, the response is insufficient to explain our obligations. A meaning-based account could

theoretical commitment badly in need of support.⁴⁶ But even if such support can be had, such a view cannot rescue the meaning-based theorist. For her primary interest, remember, was in providing an explanation of our justification for employing MP. On the view under consideration, however, a meaning-based account of justification can only “explain” our justification by appeal to the normative requirements and permissions built into the possession conditions of concepts. And such an explanation is vacuous, or dangerously close to it. What is explained – our justification – is very close to what is supposed to do the explanatory work – our epistemic permissions and requirements.⁴⁷

But perhaps we have been mischaracterizing the intuitive issues all along. Perhaps what is intuitively clear is not that we are (conditionally) required to draw certain MP inferences but that we are required to reason suppositionally in certain ways. The idea is that in some circumstances, when we accept a proposition under some supposition and then come to accept the supposition, we are epistemically required to discharge the supposition and accept the proposition outright. This is not the same, or not obviously the same, as being required to draw MP inferences and to reason with the concept *the conditional*.⁴⁸

To fully address this worry, we would need to defend a general view about the relationship between suppositional reasoning and *the conditional*. We cannot, of course, do that here. Let us settle, then, for making two quick remarks. First, it is difficult to see how the suggestion gets meaning-based theorists off the hook. On this suggestion, meaning-based theorists no longer owe an account of why we are required to draw MP

only explain why it is permissible for us to employ the rule – that is, why it is permissible for us to *treat* the relevant inferences as obligatory. It could not explain why the inferences are genuinely obligatory.

⁴⁶ This proposal is close to Brandom’s view. See his (2000) for a summary.

inferences, but they do owe an account of why we are required to engage in the analogous suppositional reasoning. And it is not clear why addressing this second challenge would be easier than addressing the first one. Second, the suggestion causes trouble for meaning-based theorists even in accounting for our epistemic justification. For it is plausible that MP is a rule that helps constitute the meaning of *the conditional*. But it is not plausible that the rule governing the discharge of suppositions is meaning-constituting of *the conditional*. It is difficult to see how a meaning-based theorist could account for our justification for performing such reasoning, let alone our obligation to do so.

Before concluding this section, we want tentatively to suggest that meaning-based accounts face a difficulty in addition to their inability to account for our conditional obligation to draw MP inferences. They also fail to account for certain unconditional obligations thinkers have. And this again demonstrates that they cannot be taken to provide the fundamental explanation of our justification.

Since, as we have already noted, to believe p and *if p then q* a thinker must possess *the conditional*, it is difficult to motivate the claim that thinkers are under an unconditional obligation to draw certain MP inferences. (We do, however, suspect that this claim is true.) It is somewhat easier to motivate the claim that thinkers have an unconditional obligation to draw certain inferences in accord with another rule, the rule of Conditional Introduction (CI): From the belief q under the supposition p , infer *if p then q* . A thinker who never draws CI inferences intuitively is violating her epistemic obligations. The thinker is epistemically required to form certain conditional beliefs given

⁴⁷ In a somewhat different context, Williamson (2003, p. 291) argues that modifying CRS by importing normative notions into the conditions for possessing a concept cannot be used to explain justification.

her beliefs under suppositions. And intuitively, this obligation is not conditional on the thinker's possession of the concept *the conditional*.

Presumably, a meaning-based theorist will want to account for the justification we have for employing CI (or some rule like it) in much the same way as she will want to account for our justification for employing MP. Both are meaning-constituting rules of *the conditional*, and both are justified via The Link. Now, any complete explanation of our justification for employing CI ought to account for, or be easily extendible to account for, our unconditional obligation to draw certain CI inferences. Given a meaning-based account, how is this obligation to be explained? A view that relies only on it being "harmless" or blameless to possess *the conditional* can account at best for an unconditional permission to draw the inferences. It cannot account for an unconditional obligation.

The meaning-based theorist presumably will claim that our unconditional obligation to draw certain CI inferences stems from our unconditional obligation to employ the rule CI (or some rule like it). The only apparent strategy for her to pursue is to argue that there is some range of concepts, including *the conditional*, such that we are epistemically required to employ their meaning-constituting methods. On this view, there is a normative principle analogous to The Link that governs the unconditional epistemic obligations thinkers have to employ certain belief-forming methods. And since we do not have an unconditional epistemic obligation to employ every meaning-constituting belief-forming method, to be extensionally adequate, the normative principle will need to be restricted.

⁴⁸ We thank an anonymous referee for drawing our attention to this intriguing set of issues.

To explain our epistemic obligation to employ CI, then, a meaning-based theorist must provide an account of which concepts have unconditionally required meaning-constituting methods. These concepts presumably will be the concepts that we are epistemically required to possess. The meaning-based theorist, then, owes a principled demarcation of the concepts that are epistemically required, and an explanation of why these concepts are required. It is by no means clear whether and how a meaning-based theorist can provide such an explanation.

Even if an explanation of which concepts we are required to possess can be provided, considerations parallel to those of section 4.1 arise. The explanation of why certain concepts are required would seem to rely on general normative principles concerning epistemic obligations and other normative notions. These principles do not plausibly depend on considerations particular to the case of concept-possession. And so it is these principles, rather than facts about concept-possession, that provide the fundamental explanation of our epistemic obligations to employ MP and other belief-forming methods.

5. The Lack of Motivation for Meaning-Based Accounts

Meaning-based accounts cannot provide the fundamental explanation of our justification for employing MP. But such accounts are not without philosophical appeal. In this section we consider some of the motivations underlying meaning-based accounts and argue that whatever kernel of truth they contain can be accommodated without endorsing a meaning-based account.

Let us begin by considering some of the considerations in favor of meaning-based accounts that arose in our presentation of the view. One motivation was grounded in the observation that if a thinker appears to believe that employing MP is unjustified, this provides strong *prima facie* evidence that she does not understand the meaning of “if ... then”. But this observation can be explained consistently with the rejection of meaning-based accounts. If, as seems likely, the fact that we are justified in employing MP is widely accepted (at least implicitly), then someone who denies this is likely to be making a verbal mistake. This general correlation between correctly understanding “if ... then” and accepting that MP is justified suffices to explain the observation. No stronger tie between justification and meaning need be invoked.

A second motivation was the naturalness of the thought, for instance, that the source of our justification for employing MP is the nature of the concept *the conditional*, and the source of our justification for believing *all bachelors are male* is the nature of the concept *bachelor*. But given the need to restrict The Link in view of examples such as *boche*, such thoughts should be treated with suspicion. And furthermore, to an extent such intuitions may be accommodated by alternative accounts, accounts in which meaning-based normative principles are merely derivative and the fundamental normative work is done by meaning-independent normative principles. (In the next section, we discuss one such view.)

A third motivation was that meaning-based accounts deliver the intuitively right result in a wide class of cases. But as we have seen, to avoid counterexample, meaning-based accounts must be restricted so as not to apply to all concepts. Some complication must be introduced. Moreover, meaning-based accounts do not explain all cases of the

justified employment of basic belief-forming methods. The extension of meaning-based accounts, then, is not one of their strengths.

Finally, it was suggested that meaning-based accounts fit well with Conceptual Role Semantics (CRS), and that this counts in favor of the view. In response, it should be noted that CRS faces serious worries, even in the case of logical concepts such as *the conditional*. Until these worries are satisfactorily dealt with, the relation between meaning-based accounts and CRS may count against, and not in favor of, meaning-based accounts. Furthermore, whereas meaning-based accounts may require some form of CRS to be at all plausible, the converse implication fails. Alternative accounts of justification may very well be compatible with CRS.

So much, then, for the more obvious apparent attractions of meaning-based accounts. Let us now briefly consider some additional motivations that have been explicitly articulated in discussions of meaning-based accounts of justification.⁴⁹

Peacocke (1993) motivates his meaning-based account by tying it to truth-preservation. Peacocke argues that there are purely meaning-theoretic reasons to believe that any concept has a semantic value that makes the constitutive rules of inference for that concept truth-preserving. If this view is correct, then those rules of inference that help to constitute acceptable conceptual roles are guaranteed to preserve truth, and so employing a rule of inference that is meaning-constituting cannot lead a thinker astray. As Peacocke (p. 190) writes, “[T]his is as good as justification can get.”

⁴⁹ The motivations due to Peacocke and Boghossian are also discussed in Horwich (2000) and (2005). Horwich also discusses a further motivation, one depending on our presumed right to mean whatever we like. Horwich argues that none of these motivations support the claim that facts concerning meaning-constitution play a role in the explanation of our justification.

Peacocke's proposal is subject to three independent problems. First, as we have seen with the example of *boche* and other thick concepts involving false normative commitments, there is reason to think that meaning-constituting methods of genuine concepts need not be truth-preserving. Second, even if meaning-constituting methods were truth-preserving, still a move from this to conclusions about justification would have to bridge the gap between truth-preservation and justification. On Peacocke's suggested view, it is not required that the thinker recognize that the relevant method is truth-preserving. Rather, it is sufficient that it is truth-preserving. But it seems that the only way to connect truth-preservation and justification so directly is to adopt a form of reliabilism about justification, such that a belief-forming method is justified in virtue of it reliably preserving truth. And there are good reasons to reject reliabilism.⁵⁰ Finally, even if meaning-constituting methods were truth-preserving, and even if some form of reliabilism provided the correct account of justification, still the role meaning-constitution plays would be secondary at best. The fundamental explanation of our justification for employing MP would simply be that it was a truth-preserving method. Considerations of meaning would not do any real explanatory work.⁵¹

Boghossian (2000, p. 250) suggests a very different motivation for meaning-based accounts:⁵²

⁵⁰ See Bonjour (1985) for the well-known clairvoyance cases in which a thinker is epistemically irresponsible – and so presumably unjustified – in employing certain belief-forming methods despite the fact that the methods are reliable.

⁵¹ The last objection can also be found in Horwich (2005).

⁵² A similar argument may be found in Boghossian (2003) as an “initially helpful thought”. He there indicates that he no longer endorses it as stated. Boghossian leaves it unclear whether he endorses a modified version.

If it is true that certain of our inferential dispositions fix what we mean by our logical words (in the language of thought), then it is very plausible that we should be entitled to act on those inferential dispositions prior to, and independently of, having supplied an explicit justification for the general claim that they are truth preserving, for *without* those dispositions there is nothing about whose justification we can intelligibly raise a question about: Without those dispositions we could not even *have* the general belief whose justification is supposed to be in question.

The thought seems to be this: Suppose that CRS is correct for *the conditional*. Then, a thinker cannot have any beliefs involving *the conditional* without already employing certain belief-forming methods involving *the conditional*, since employing those methods is part of what it takes to possess the concept. And the thinker can only have an explicit justification for employing such a method if she already possesses the concept *the conditional*, since the concept features in the claim being justified. Hence, the thinker is not epistemically blameworthy for – and so is justified in – employing the method without having an antecedent justification for it.

Thus put, the argument depends on a normative principle that can be stated roughly as follows: If a thinker cannot explicitly justify a belief-forming method, then the thinker need not provide an explicit justification to be justified in employing it. And this principle, as it stands, is not at all plausible. Clearly, many belief-forming methods cannot be explicitly justified, because they are unjustifiable for any thinker to employ.

The inability of a thinker to provide a justification does not entail that the method may be justifiably employed.

For the motivation to be at all plausible, some more restricted normative principle must be invoked. It is not clear what this more restricted principle could be. Boghossian seems to suggest that the principle will depend heavily on it being impossible to even “intelligibly raise a question” about the justification of the relevant methods without relying on those very same methods.⁵³ But even apart from the unclarity about what exact principle is intended, there seems little reason to accept any principle of this sort. The fact that a thinker must accept a method antecedently to being able to raise the issue of its justification does not demonstrate that an explicit justification of the method is unneeded once the thinker employs the method. It is not impossible for a thinker to present an explicit justification of a method she employs. So it is not clear on what grounds Boghossian can reject the requirement that a thinker must present an explicit justification in the range of cases at issue.

There is an additional difficulty with the suggested motivation for meaning-based accounts. Even if successful, Boghossian’s motivation only directly supports the claim that in the relevant range of cases no explicit justification is needed. It does not follow that what ultimately accounts for the justification is a meaning-based account. Some alternative account of the justification of the relevant methods might do just as well or better.

⁵³ Such a principle seems to apply equally well to the rules constitutive of *boche* and *flurg*, and so is subject to counterexample. Boghossian’s discussion does provide materials that might be used to address this difficulty. He could argue that thinkers possessing unconditional concepts can be interpreted by thinkers possessing their conditional variants, and so the issue of the justification of a meaning-constituting method for an unconditional concept can be raised by a thinker possessing only its conditional variant. However, making this modification has a cost; it reduces some of the intuitive appeal the motivation originally seemed to have.

Boghossian also motivates meaning-based accounts by emphasizing one of their advantages.⁵⁴ They occupy an intuitively appealing middle ground between extreme versions of both internalism and externalism about justification. It is unreasonably strict to require that to be justified in employing a belief-forming method, a thinker must possess an explicit justification for employing it, and so extreme versions of internalism about justification ought to be rejected. On the other hand, the idea that the reliability of a belief-forming method is sufficient for a thinker to be justified in employing it, regardless of whether the thinker does, can, or should know that the rule is reliable, is just as counterintuitive. So extreme versions of externalism about justification also ought to be rejected. The hope is that by appealing to meaning-constitution one may adopt an intermediate position. According to The Link, there are cases in which one is justified in employing a belief-forming method despite having no available explicit justification, in contrast with extreme internalism. And the explanation of this justification does not depend on the method being reliable, in contrast with extreme externalism.

We agree that this is an advantage of meaning-based accounts (if, that is, their details can be filled in to make good on the promise of supplying such a middle ground). But as a motivation for meaning-based accounts, this consideration should be rejected. Other accounts may satisfy this desideratum just as well, and perhaps even better. In particular, the account we sketch in the next and final section does just that.

6. An Alternative Account of Justification

If the arguments in section 4 are correct, meaning-based accounts of justification cannot be understood as providing the fundamental explanation of our justification. And if the

⁵⁴ See Boghossian (2000), (2001), and especially (2003).

arguments in section 5 are correct, the motivations typically adduced for meaning-based accounts are unpersuasive. This naturally invites the question: If not a meaning-based account, then what? Indeed, if there is no acceptable alternative account in the offing, one might be forced to accept a meaning-based account as the best of a bad lot. To address this challenge, in this section we provide a sketch of an alternative account of justification, an account that we call “the pragmatic account” of justification. We believe that this account – which we elaborate elsewhere⁵⁵ – holds promise as providing the fundamental explanation of our justification for employing MP and other basic belief-forming methods.

Recall that in section 4 we emphasized the need to provide a unified account of justification, one that explains our justification for employing MP, IBE, relying on perception and memory, and other belief-forming methods. These methods have very little in common, other than the important role they play in our lives. This role seems relevant to the explanation of their justification. Consider, for instance, IBE. Employing IBE is central to how we go about explaining the world. And this explanatory project is of tremendous importance to us. We are explaining creatures, creatures that try to make sense of ourselves and the world around us. Perhaps we are essentially explaining creatures. Perhaps we cannot disengage from the explanatory project so long as we are to count as rational agents. But even if we can disengage, it doesn't seem that we should. In an intuitive sense, the explanatory project is a project in which we are rationally required to engage.

This line of thought suggests an explanation of our justification for employing IBE: We are rationally required to engage in the explanatory project. This counts in favor

⁵⁵ See Enoch and Schechter (forthcoming). See Schechter (2006) for discussion focused on the case of MP.

of employing whatever methods are necessary to engage successfully in the explanatory project. IBE is plausibly such a method, at least for beings like us. And this explains our justification for employing IBE.

Of course, the needs of the explanatory project cannot explain our justification to employ a wide range of belief-forming methods. But it is straightforward to generalize the main idea, since there are many projects in which we are rationally required to engage. Describing the world, understanding the world, rationally deciding what to do, and having thoughts at all plausibly are such projects. If employing a certain belief-forming method is required to engage successfully in one of these projects, then we seem to be justified in employing it.

Here, then, is the general idea behind the pragmatic account: Given a project that we are rationally required to engage in, and given a belief-forming method that we must employ if we are to engage successfully in that project, we are justified in employing the method.

The statement of the account requires several refinements. Here are three: First, the sort of justification a belief-forming method gains is only *prima facie* justification, the kind of justification that can be overridden and perhaps undermined by other considerations. Second, the pragmatic account ought to be restricted to *basic* belief-forming methods. Non-basic methods are justified (if at all) by appealing to other, more fundamental, belief-forming methods. Third, for the pragmatic account to vindicate a belief-forming method it should not be required that employing the method is the only way that we may engage successfully in a given project. Rather, if a belief-forming

method is such that if any method would yield success it would too, we seem to be justified in employing the method.

Of course, many additional details need to be filled in for the pragmatic account to be fully convincing. What is it for a project to be such that we are rationally required to engage in it? Which projects are rationally required? What is it to engage successfully in a project? What are the relevant modalities involved in the statement of the account? And even with these details in place, several important objections remain to be answered. Nevertheless, the sketched account is intuitively appealing, and we have reason to hope – and to believe – that these details can be satisfactorily filled in, and the relevant objections adequately addressed.⁵⁶

Notice that the pragmatic account provides answers to the analogues of the three challenges to meaning-based accounts discussed in section 4. The pragmatic account offers a principled way of distinguishing methods that are justified from methods that are unjustified; all epistemic justification ultimately traces back to the rational requirements we have to engage in certain projects. The pragmatic account also provides a general, unified answer to the question of in virtue of what many of our basic belief-forming methods are justified for us to employ. Employing IBE is presumably needed for success in the explanatory project. Relying on moral and normative intuitions is presumably needed for success in deliberation.⁵⁷ Assuming that modal notions are indispensable to science, relying on modal intuitions is indispensable to explanation. And so on. Finally, the pragmatic account can explain the epistemic obligations we have to employ certain

⁵⁶ We fill in many of these details in Enoch and Schechter (forthcoming). We also address several worries about the pragmatic account, perhaps most notably that it conflates epistemic and pragmatic justification, that it falls prey to the same problem that bedeviled Reichenbach's account of the justification of induction, and that it fails to account for the objectivity of our justification.

belief-forming methods and to draw certain inferences. If employing certain methods is needed to engage successfully in projects that we are required to engage in, then it is obligatory – not merely permissible – for us to employ those methods.

The sketched pragmatic account is consistent with a meaning-based account of the justification for employing MP. (Of course, if both are adopted, the pragmatic account is what does the most fundamental explanatory work.) So long as possessing *the conditional* is needed for successfully engaging in some rationally required project, its meaning-constituting methods are justified for us to employ. And it is somewhat plausible that possessing *the conditional* is indeed necessary for engaging in some rationally required projects. It is difficult to imagine how a thinker could successfully inquire, deliberate, or reason at all without having the ability to have conditional beliefs.

Indeed, it is plausible that Boghossian's discussion presupposes a version of the pragmatic account.⁵⁸ Boghossian claims that belief-forming methods that help constitute conditional concepts are justified because possessing such concepts does not rule out possibilities that should remain open for inquiry. Belief-forming methods that constitute unconditionalizeable concepts are justified, on the other hand, because such concepts are needed to conditionalize, and hence are needed to engage in inquiry. This suggests that on Boghossian's view the justification in all of these cases ultimately springs from the needs of rational inquiry, presumably a rationally required project.

Meaning-based accounts may very well rely on something like the pragmatic account. But with the pragmatic account in hand, an appeal to considerations of meaning

⁵⁷ For an argument that the needs of deliberation justify us in believing in normative truths, see Enoch (2003) and (forthcoming).

seems no longer to be necessary. Consider again the case of MP: It is very plausible that employing MP is directly indispensable to explanation, deliberation, and other rationally required projects. The use of MP in constituting the meaning of *the conditional* need not be mentioned.

Meaning-based accounts, we conclude, are undermotivated (section 5) and inadequate as candidates for the ultimate justificatory story (section 4). The pragmatic account sketched here holds promise as a more adequate alternative. Considered as candidates for the most fundamental account of our justification for employing MP, then, meaning-based accounts ought to be rejected.⁵⁹

⁵⁸ Towards the end of his attack on Boghossian's view, Williamson (2003, p. 289) suggests that a natural fallback position for Boghossian is to accept the "crude pragmatic idea" that we "could not do without logic".

⁵⁹ The authors contributed equally to this paper. The present paper is a descendent of our manuscript, "Epistemic Justification, Pragmatically Justified", written in the spring of 2001 as a response to Boghossian (2000) and (2001). In that paper, we provided objections to meaning-based accounts of justification, argued that a Reichenbach-inspired pragmatic account could answer the objections, and developed our pragmatic account in detail. The detailed discussion of the pragmatic account later became Enoch and Schechter (forthcoming). We are very grateful to Paul Boghossian for his valuable advice and criticism at each stage of this project. We would also like to thank Cian Dorr, Greg Epstein, Hartry Field, Paul Horwich, Anna-Sara Malmgren, Derek Parfit, Christopher Peacocke, Karl Schafer, Crispin Wright, Masahiro Yamada, and two anonymous referees for their helpful feedback.

References

- Belnap, Nuel. (1961) "Tonk, Plonk, and Plink," *Analysis* 22: 130–134.
- Block, Ned. (2000) "Semantics, Conceptual Role," in Edward Craig (ed.), *Routledge Encyclopedia of Philosophy*, London: Routledge.
- Boghossian, Paul. (1996) "Analyticity Reconsidered," *Noûs* 30: 360–391.
- Boghossian, Paul. (2000) "Knowledge of Logic," in Boghossian and Peacocke (2000), pp. 229–254.
- Boghossian, Paul. (2001) "How are Objective Epistemic Reasons Possible?," *Philosophical Studies* 106: 1–40.
- Boghossian, Paul. (2003) "Blind Reasoning," *Proceedings of the Aristotelian Society, Supplementary Volume* 77(1): 225–248.
- Boghossian, Paul and Peacocke, Christopher, eds. (2000) *New Essays on the Apriori*, Oxford: Clarendon Press.
- BonJour, Laurence. (1985) *The Structure of Empirical Knowledge*, Cambridge, Mass.: Harvard University Press.
- Brandom, Robert. (2000) *Articulating Reasons: An Introduction to Inferentialism*, Cambridge, Mass.: Harvard University Press.
- Burge, Tyler. (1993) "Content Preservation," *Philosophical Review* 102: 457–88.
- Dummett, Michael. (1973) *Frege: Philosophy of Language*, New York: Harper and Row. Second edition 1981, Cambridge, Mass.: Harvard University Press.
- Dummett, Michael. (1978) "The Justification of Deduction," in his *Truth and Other Enigmas*, Cambridge, Mass.: Harvard University Press, pp. 290–318.
- Dummett, Michael. (1991) *The Logical Basis of Metaphysics*, Cambridge, Mass.: Harvard University Press.
- Enoch, David. (2003) *An Argument for Robust Metanormative Realism*, Dissertation, New York University.
- Enoch, David. (forthcoming) "An Outline of an Argument for Robust Metanormative Realism," *Oxford Studies in Metaethics* 2.
- Enoch, David and Schechter, Joshua. (forthcoming) "How are Basic Belief-Forming Methods Justified?," *Philosophy and Phenomenological Research*.
- Fodor, Jerry and Lepore, Ernie. (1991) "Why Meaning (Probably) isn't Conceptual Role," *Mind and Language* 6(4): 328–342.
- Goldman, Alvin. (1986) *Epistemology and Cognition*, Cambridge, Mass.: Harvard University Press.
- Gómez-Torrente, Mario. (2002) "The Problem of Logical Concepts," *The Bulletin of Symbolic Logic* 8(1): 1–37.
- Hacking, Ian. "What is Logic?," *The Journal of Philosophy* 76(6): 285–319.
- Hale, Bob and Wright, Crispin. (2000) "Implicit Definition and the *A Priori*," in Boghossian and Peacocke (2000), pp. 286–319.
- Hale, Bob and Wright, Crispin. (2001) *The Reason's Proper Study*, Oxford: Oxford University Press.
- Harman, Gilbert. (1995) "Rationality," in Edward Smith and Daniel Osherson (eds.), *Thinking: An Invitation to Cognitive Science*, second edition, vol. 3, Cambridge, Mass.: The MIT Press, pp. 175–211.

- Horwich, Paul. (1997) "Implicit Definition, Analytic Truth, and Apriori Knowledge," *Noûs* 31(4): 423–440.
- Horwich, Paul. (2000) "Stipulation, Meaning and Apriority," in Boghossian and Peacocke (2000), pp. 150–169.
- Horwich, Paul. (2005) "Meaning Constitution and Epistemic Rationality," chapter six in his *Reflections on Meaning*, Oxford: Oxford University Press.
- Peacocke, Christopher. (1987) "Understanding Logical Constants: A Realist's Account," *Proceedings of the British Academy* 73: 153–200.
- Peacocke, Christopher. (1992) *A Study of Concepts*, Cambridge, Mass.: The MIT Press.
- Peacocke, Christopher. (1993) "How are A Priori Truths Possible?," *European Journal of Philosophy* 1: 175–199.
- Peacocke, Christopher. (2004) *The Realm of Reason*, Oxford: Oxford University Press.
- Prior, Arthur. (1960) "The Runabout Inference Ticket," *Analysis* 21: 38–39.
- Pryor, James. (2000) "The Skeptic and the Dogmatist," *Noûs* 34: 517–549.
- Schechter, Joshua. (2006) *Two Challenges to the Objectivity of Logic*, Dissertation, New York University.
- Stevenson, John. (1961) "Roundabout the Runabout Inference Ticket," *Analysis* 21: 124–128.
- Tennant, Neil. (1987) *Anti-Realism and Logic*, Oxford: Clarendon Press.
- Williams, Bernard. (1985) *Ethics and the Limits of Philosophy*, Cambridge, Mass.: Harvard University Press.
- Williamson, Timothy. (2003) "Blind Reasoning," *Proceedings of the Aristotelian Society, Supplementary Volume* 77(1): 249–293.
- Wright, Crispin. (1999) "Is Hume's Principle Analytic?," *Notre Dame Journal of Formal Logic* 40(1): 6–30.
- Wright, Crispin. (2001) "On Basic Logical Knowledge," *Philosophical Studies* 106: 41–85.