

Jeffrey Daniel Torborg

+1 516-640-9283

jdt2166

Prof. Justin Clarke-Doane

Inference and Entitlement

03/26/22

I. The Problem, and an Outline

In virtue of what are we justified in employing certain deductive rules like Modus Ponens (MP)? Put differently, owing to what are we justified in following a rule, MP, which licenses us to infer q from the justified beliefs p and *if p then q* ? Presumably—on some intuitive understanding of ‘basic’—MP is a basic deductive rule for us. Hence, in employing MP, our justification won’t come from our being justified in employing other, more basic deductive rules. And though I assume in what follows that MP is a basic deductive rule, if we suppose otherwise—such that we can explain our justification for employing MP by our being justified in employing this more basic rule—we can simply rerun the question with *this* rule in place of MP: in virtue of what are we justified in employing *it*? Justifications come to an end somewhere, and in the context of deductive rules like MP (or whichever rules we suppose basic), our task is explaining in virtue of what we are justified in following them.

One approach appeals to considerations of *meaning*. Boghossian (2020b), for instance, pursues this line by appealing to an inferentialist stance on our grasp of the logical constants, among them the concept of the conditional¹. On this view, our grasp of the logical constants

¹ In what follows, I will use “the concept of the conditional” and “*if*” interchangeably.

amounts to our defeasible willingness to assent to certain inferences using them. To the extent that certain inferential transitions make an appearance in the possession-conditions of the concept of some logical constant, finding those inferential transitions compelling is a precondition on our grasp of that concept. In the case of MP, then, the thought is that our very grasp of the conditional requires that we be defeasibly disposed to reason according to MP. Whence comes our justification for employing MP? This proposal suggests that we are so justified, since required for our grasp of the conditional is our defeasible willingness to follow MP.

It is this proposal that the present paper adopts as its starting point. Specifically, I develop the meaning-based response to our initial question in the fashion of Boghossian's thoughts in his "Blind Reasoning" (2020b). Yet, after sketching some ways in which it seems objections to this view ought *not* proceed, I raise some critical worries for this style of meaning-based approach. Echoing earlier worries about the effectiveness of transcendental arguments, raised by Stroud (1968), I suggest that facts about concept-possession alone won't justify us in reasoning according to certain deductive rules, because left unsettled is whether we are justified in reasoning with the concept in question. That is, even if we show that reasoning according to MP is a necessary condition on our grasping the conditional, this fact in tandem with the fact that we grasp the conditional cannot alone justify us in employing MP, for that would seem to require a justification to reason with the concept of the conditional, rather than a related concept conditional*, grasping which does not require we assent to MP.

Moving on from considerations of meaning, I next discuss a different style of approach to our initial question, a broadly pragmatist proposal developed in Enoch and Schechter (2008). On this view, there are certain projects which we (as rational agents, cognizers, or the like) cannot

help but engage in: call these *rationally required projects*. To successfully engage in a rationally-required project, certain belief-forming methods—deductive rules, possibly—must be effective. Owing to the status of these projects as rationally required, we cannot help but regard the relevant indispensable methods as effective; hence, we are justified in employing them. How does this account for our justification in employing MP? We are to understand ourselves justified in employing MP insofar as we see the project of planning for future contingencies as rationally-required, and the effectiveness of MP as necessary for our successfully engaging in this project: in short, we are justified in employing MP owing to the fact that it is indispensable to a rationally-required project. Yet, this proposal too suffers weaknesses. Owing to critical ambiguities in the definition of a rationally-required project, it emerges that the proposal is subject to four crucial worries, among them that the fundamental justificatory work is accomplished not by a deductive rule's indispensability to a project, but by whatever normative status the relevant project carries.

The structure of this paper, then, occupies itself mainly with the treatment of two proposals. I develop the meaning-based account found in Boghossian (2020b) and the pragmatist account rendered in Enoch and Schechter (2008). I suggest, in turn, that each of these suffer weaknesses, though at first pass, I stop short of the claim that *any* meaning-based or pragmatist account suffers the weaknesses I point out; only an account after the manner of the proposals discussed here, in the relevant respects, will be subject to the worries raised.

In closing, however, I take up the question directly, asking in virtue of what we are justified in employing MP. I suggest that the question, construed realistically, is misguided. Here, I *do* raise broader considerations, telling against the aptness of meaning-based or pragmatist accounts generally: the proposals of Boghossian and Schechter and Enoch settle

certain kinds of facts about MP (that it's constitutive of our grasp of the conditional, that it's indispensable to a rationally-required project, etc.), settling which, nonetheless, fails to settle the question of whether to employ it; but, I close with a broadly non-cognitivist moral, advancing the stronger recommendation that settling *any* fact about MP, even the facts about justification, fails to settle the question whether to employ it.

II. A Meaning-Based Approach

Begin with some ground-laying, by way of two related points. First, I follow Enoch and Schechter (2008) in denying that a thinker need believe that they are employing a method (or entertain any beliefs about the method) for them to count as employing it. And a similar point holds for those cases in which we want to say we're justified in employing a method. I adopt a way of framing things owing to Burge (1995), on which a person's being justified doesn't require that they have reflective access to in what their justification consists: call a person "entitled" when, though from the theorist's perspective they are justified, they nonetheless lack an articulated account of their justification. Certainly, to count someone as justified (or entitled), a theorist ought to possess such an account: indeed, this intuition is the motivation for the proposals discussed here. But in what follows, I will often opt for the language of entitlement, mainly to focus the claim at which Boghossian's argument drives: certain facts about concept-possession in place, whether they are aware of these facts or not a person is entitled to follow some deductive rule, owing merely to their grasp of the relevant concept.

Now, how is this central claim reached? To begin, Boghossian appeals to implicit definitions, and an inferentialist stance on our grasp of the logical constants. Broadly, an implicit definition of a term is held to work by our regarding certain sentences containing it as true, and

by our supposing that this decision bestows a meaning (Horwich 1997, p. 424). More exactly, a sentence “#F” is an implicit definition of some term F if all the other terms of “#F” are already understood, and our decision to regard “#F” as true fixes F with whatever meaning makes “#F” true (p. 424). Analogously, we can speak of implicit definitions as proceeding from treating certain inference rules as valid. In this case, we implicitly define F when we stipulatively decide that certain inference-patterns containing F are to be regarded as valid, and we take it that F has whatever meaning it must have for the inferences containing it to be valid.

Implicit definition affords us a plausible account of our knowledge of the meanings of the logical constants: they are implicitly defined by our stipulation of their standard introduction and elimination rules, and they are thereby given whatever meaning they must have for those rules to be correct². In the case of the logical constants, that the work of implicit definition is being done by rules of inference—and that the meanings of the constants is given³ by their appearance in certain inferences and not others—amounts to an *inferentialist* account of the semantics of the logical constants. Importantly, this stance furnishes us with a picture of what it takes to *grasp* the concept of some logical constant. Call the inference-patterns or deductive rules, appearing in which a logical constant has its meaning fixed, *constitutive* inference-patterns. Then, on this view, our grasp of a logical constant consists in our willingness to reason according to those constitutive inference-patterns. To possess the concept of some logical constant, it must be that a thinker possess a defeasible disposition to treat the relevant inferential transitions as correct, or to find them “primitively compelling” (Peacocke 1992, p. 6). On this view, I cannot be said to grasp

² See both Hale and Wright (2000, p. 287) and Giovannini and Schiemer (2019, p. 1685) for a discussion of the history and prospects of this proposal.

³ Or, more strongly, that the meanings the constants acquire are to be *identified* with the rules of inference that define them: “According to logical inferentialism, the meaning of logical constants can be equated with the usual rules of inference that govern their correct use” (Giovannini and Schiemer, p. 1685).

some concept if I am not defeasibly prepared to reason according to its constitutive inference-patterns, for my possessing this concept consists in just this disposition to so reason.

It is this picture of meaning and, especially, concept-possession that will figure prominently in what follows. For it is to this generally inferentialist account of the meanings of the logical constants (2020a, pp. 14, 19), and of concept-possession (2020b, p. 37), that Boghossian finally arrives en route to answering our initial question: in virtue of what are we entitled to employ MP⁴?

The proposal is put succinctly as the *meaning-entitlement connection*: this holds that “Any inferential transitions built into the possession conditions for a concept are *eo ipso* entitling” (p. 38). In other words, the proposal states that we are entitled to just those inference-patterns that are constitutive of the meaning of the logical constant—that is, to the inference-patterns we must find compelling if we are to possess the relevant concept.

Why should we suppose that we are entitled to the deductive rules that are constitutive of our grasp of a concept? Boghossian’s specific brief is to defend a kind of entitlement on which we are not epistemically blameworthy in following an inference-pattern, but nor is it required, for an inference-pattern to be entitling, that we have reflective access to in what our entitlement

⁴ There are, to be clear, at least two distinct developments of this claim—the claim that an inferential understanding of our grasp of the logical constants secures us an entitlement to follow certain deductive rules—in Boghossian’s work: in 2020a, and 2020b. Insofar as each purports to accomplish the same (vindicating the entitlements that arise from concept-possession), I feel comfortable developing only one of them here. But my choice, in what follows, of laying out Boghossian’s final position in 2020b, rather than 2020a, owes itself to the following: the formulation in 2020a is both slightly more elaborate and, accordingly, vulnerable to objection on what I take to be less-than-essential grounds. Though space doesn’t permit an involved discussion of this formulation here, Ebert (2005) discusses some of the controversy unique to it. Relatedly, as mentioned in the introduction, I focus on the account in Boghossian 2020b rather than on similar meaning-based approaches pursued in Peacocke (1993), and in Hale and Wright (2000); the somewhat wider focus of Enoch and Schechter (2006) raises concerns that are directed at Boghossian (2020b), as mine are, but also at these other approaches. It will emerge in concluding, however, that my local critiques of Boghossian (2020b) lead naturally to some sweeping claims about justification, which I understand to impact any meaning-based approach, these others included.

consists⁵. This notion of entitlement in mind, consider our current view of concept-possession, on which we cannot be understood to grasp a concept without finding certain inferential transitions compelling. If I cannot so much as possess a concept without employing certain inference-patterns—if I can't grasp the ingredient concepts in an inferential transition without regarding that transition as correct—then Boghossian suggests that we are epistemically blameless in employing the inference-pattern: “If inferring from A to B is required, if I am to be able to think the ingredient propositions, then it looks as though so inferring cannot be held against me, even if the inference is blind” (2020b, p. 37). And ostensibly, if we cannot grasp a concept embedded in an inference, neither can we understand the inference itself; even entertaining the inference requires that we grasp the ingredient concept, and *this* requires that we be disposed to make the inference: to truly entertain the inference, we have no option but to think of it as correct. Though we have no antecedent reason to suppose an inference-pattern correct, so long as our grasp of some concept could not come about without our taking the inference-pattern to be correct, we are blameless in so regarding it, and thus in employing it.

The full development of Boghossian's proposal, then, is the following: if it is true that the meanings of our logical constants is understood inferentially—which is to say that their meanings are fixed implicitly by certain rules of inference in which they appear, and accordingly that possessing the concept of a logical constant requires a thinker to be defeasibly disposed to reason according to certain inference-patterns—then grasping the concept entitles us to employ just those 'constitutive' inference-patterns, because we could not have grasped the concept in the

⁵ Boghossian (2020b) eschews, in accounting for our entitlement to reason according to certain basic inferences like modus ponens, both a naïve reliabilism and a simple internalism—for the former allows that someone might be deemed justified who we would regard as epistemically irresponsible (blameworthy), and the latter suffers problems of Carroll-like regress, among other problems (p. 33). Hence the working notion of entitlement reached in (2020b), in terms of which Boghossian frames his thesis concerning entitlements from concept-possession, is one on which we can be *blameless* in our employment of some inference-pattern, but reflectively *blind* as to why we are entitled in so inferring.

first place without doing so. In the case of MP, the claim is that our grasp of the concept of the conditional requires that we be defeasibly willing to conclude q from beliefs p and *if p , then q* (i.e., that we be willing to follow MP). To our initial question, ‘in virtue of what are we are entitled to employ MP?’, the response is the following: given that MP is an inference-pattern constitutive of our grasp of the conditional, we are entitled to follow it owing to the fact that we do grasp the conditional.

It is the above picture which I’ll regard as the archetype of a meaning-based response to our question, one which I hope to suggest in a moment is misguided. But first, I would like briefly to address two styles of objection to meaning-based accounts like Boghossian’s, which I think miss the mark, and offer candidate routes by which I think a successful objection ought *not* proceed.

The first of these questions whether, in fact, the meanings of the logical constants are fixed in the way described⁶. But this style of resistance to Boghossian’s proposal, insofar as it strives to some such claim as “grasping a concept doesn’t in fact work in the way Boghossian describes”, in any case misconstrues the proposal’s force. That is, Boghossian’s proposal is at heart a conditional: *if* our grasp of the logical constants is understood inferentially, then entitlements to certain inference-patterns are had merely from concept-possession⁷. The striking epistemic claims of this proposal—i.e., the thought that on a certain view of concept-possession, we *could* get entitlements to some inferences for free so long as we grasp the relevant concept—survives the effort to show merely that our view of concept-possession is mistaken.

⁶ I take it that the force of Williamson (2020) and Casalegno (2004), in showing that links between concept-possession and willingness to assent to inferences is shallow, lies in pursuing this style of objection.

⁷ “*If* a conceptual-role semantics is true, and if A is indeed constitutive of C’s meaning what it does, then those facts by themselves constitute a warrant for A” (Boghossian 2020a, p. 23, italics mine).

A second kind of objection is expressed in Horwich's discussion (1997) of the meanings of our scientific terms. Suppose we want to say that some theoretical term is implicitly defined by the sentences of a scientific theory, such that our term has whatever meaning renders the relevant sentences true. But it seems that the relevant terms possess a meaning even if the allegedly defining theories one day turn out to be *false*. If we want to say that "phlogiston" was implicitly defined by the relevant sentences of 18th-century chemistry, the puzzle is this: if implicitly defined, "phlogiston" was given whatever meaning it would need for the relevant sentences to turn out true; it seems indeed that "phlogiston" was given a meaning; but, we now understand the sentences of the theory to be false!

This worry can be interpreted as motivating one of two things, a *semantic* or an *epistemic* thesis. Regarding the first of these, we might take it to be denying that implicit definition is actually how our theoretical terms in science get defined—and the worry generalizes to any term defined by a sentence which could, in principle, turn out false. Now again, Boghossian argues that, *given* the inferentialist story about meanings, we get certain entitlements. For this reason, worries about whether our terms in fact get their meanings in this way seem to leave inadequately challenged the interesting epistemic claims at issue.

Now, the *epistemic* reading of this objection raises a worry about whether grasping an implicitly-defined meaning can really entitle us to anything. For we can read the worry as stating: "grasping an implicitly-defined meaning can't entitle us to believe the defining sentences, because it could very well be that these sentences one day turn out false (or that the defining inference rules turn out invalid), as the empirical science example shows". Yet my response is that, to the thesis that entitlements flow from grasping implicitly-defined terms, it should be no threat that these entitlements are defeasible. If we suppose that, grasping some

concept, we are entitled to just those inferences that are constitutive of our grasp of it, we certainly should not suppose that these entitlements cannot be defeated by later evidence, or that the only candidate entitlements we can read off of concept-possession are indefeasible ones. Boghossian certainly does not expect indefeasible entitlements to arise from concept-possession, and it is perfectly compatible, with the view that concept-possession entitles us to regard something as true or valid, that we someday come to view it as false or invalid⁸.

III. Transcendental Arguments and Concerns

It is here, then, that I hope to sketch a new worry for Boghossian's proposal, one which—crucially—avoids the pitfalls of the objections just mentioned.

To begin, assume the inferentialist stance towards the meanings of our logical constants, and our grasp of the relevant concepts, is correct: their meanings are fixed implicitly by certain rules of inference in which they appear, and relatedly, possessing the concept of a logical constant requires a thinker to be disposed to employing certain inference-patterns. Then the proposal states the following: if being so disposed is a necessary condition on our grasp of some logical constant, then merely grasping the concept, we are entitled to employ the relevant inference-patterns. Importantly, this entitlement need not be an indefeasible entitlement: compatible with the proposal that entitlements arise from concept-possession is the possibility that the entitlement is defeated by theoretical reflection.

Now, the next step comes in construing our initial proposal as a transcendental argument, to the extent that it accounts for an entitlement of ours in the same way a transcendental argument would. A transcendental argument, roughly, is an argument that purports to justify our

⁸ Indeed, Boghossian is especially clear on this point, though in different contexts: see (2020b), pp. 38-39.

belief that X, on the grounds that X's obtaining is a necessary condition on the possibility of Y, where we know already that Y obtains; since Y obtains, then, we are entitled to the belief that X. Boghossian's argument, similarly, strives to explain our entitlement to employ some inference-pattern, on the grounds that our doing so is a necessary condition on the possibility of our grasp of some concept, and we *do in fact* grasp the concept; since we do grasp the concept, then, we are entitled to employ the relevant inference-pattern. For support of so construing Boghossian's argument as a transcendental argument, see the following:

Suppose it's true that my taking A to be a warrant for believing B is constitutive of my being able to have B-thoughts (or A-thoughts, or both, it doesn't matter) in the first place. Then doesn't it follow that I could not have been epistemically blameworthy in taking A to be a reason for believing B?...If inferring from A to B is required, if I am to be able to think the ingredient propositions, then it looks as though so inferring cannot be held against me (2020ab, 37).

Here, it is proposed that making certain inferential transitions (for instance, from A to B) is a necessary condition on my coming to understand or entertain thoughts of either. For, as spelled out earlier, I cannot understand an inference without understanding the concepts ingredient in it, but I cannot do *that* (understand an ingredient concept) unless I am willing to make the inferential transition. So we have an especially striking statement of the necessary conditions on concept-possession: to even entertain a concept, it is necessary that I be willing to employ some inference-pattern in which it appears. And, as in the final move of a transcendental argument, since we *do* have thoughts about concepts—we do grasp them—it must be the case that we are entitled to the necessary conditions of such a grasping; namely, we are entitled to reason

according to the relevant inference-pattern: we cannot be epistemically blameworthy in employing it.

There is an important condition to be met by any transcendental argument, however, if it is to successfully account for an entitlement of ours. As Stroud (1968) observes, for transcendental arguments to establish that we are entitled to believe X on the basis that it is necessary for the obtaining of Y, it is insufficient that Y be merely a certain area of our thought and talk, or merely the way we think and talk as we do *now*, with the concepts we now use. For even if we were to show that X is a necessary condition on our having the concept Y, absent some further claim that we *must* or should reason with Y, it is open to us to simply give up reasoning with Y: “we could simply give up our present ways of thinking and speaking (of which they [the X’s] are the necessary conditions) and adopt others (of which they are not)” (1968, p. 252). To illustrate this, consult the obvious weakness of the following: libertarian free will is true, since without libertarian free will, we could not use the concepts of praise and blame, and we actually *do* use the concepts of praise and blame. For, insofar as we can simply cease to use the concepts of praise and blame, we have not shown that we are entitled to believe anything about free will. Absent some claim that we *couldn’t help* thinking and talking and having the experience we actually do now⁹, or that are ones that we couldn’t do without, we cannot justify our belief in their necessary conditions, whatever those might be.

The worry, then, for the original proposal, is this. Recall our treatment of the concept of the conditional: since reasoning according to MP is a necessary condition on our grasp of the concept of *if*, and since we do in fact grasp *if*, we are entitled to reason according to MP. But consider now a concept intimately related with *if*: *if**. *If** participates in nearly all the same

⁹ As I take it is Kant’s claim, for instance, in the Transcendental Aesthetic of the First Critique, hence its force in justifying our entitlement to the concept of space.

inferences as *if*, and the only salient difference between them is the following: while to grasp *if* it is required that we defeasibly find instances of MP primitively compelling, to grasp *if*^{*} it is *not* required that we defeasibly find all instances of MP compelling. Consider a set of inferences, inferences of the MP schema among them, which we must be willing to assent to, if we are to grasp *if*. Though the possession-conditions for *if*^{*} are the same in every other respect, it is not the case that our grasp of *if*^{*} requires we be willing to assent to just those inferences in the set. (Postpone for the moment worries about what *if*^{*} looks like; these I'll address below).

Now, it's agreed that a defeasible willingness to reason according to MP is a necessary condition on grasping *if*. Have we thereby shown, grasping *if*, that we are entitled to some MP inference? It would seem, following the previous discussion of transcendental arguments, that we haven't. For showing that a defeasible willingness to reason according to MP is a necessary condition on grasping *if* does not settle whether to reason with *if*, rather than *if*^{*}. Absent some argument to the effect that we must or ought to use *if* rather than *if*^{*}—which is *not* an argument furnished by our spelling out the necessary conditions on our grasp of either—we have not shown that we are entitled to the necessary conditions on *if* rather than *if*^{*}. For we might simply cease to use the concept of *if*, and adopt *if*^{*}, grasping which does *not* require that we defeasibly regard certain instances of the MP schema valid. That we do in fact grasp *if*, and that some MP inference is constitutive of our grasping *if*, doesn't settle whether to reason with *if* rather than *if*^{*} to begin with.

So, does concept-possession *alone* entitle us to reason according to some inference, where so reasoning is a necessary condition on our grasp of the concept? Alone, no. Absent some further argument—showing that we ought to reason with *this* concept rather than any other, or that we cannot help but use this concept—we could simply cease using the concept in question,

at which point the observation that some inference is a necessary condition on our grasping it goes no way towards securing our entitlement to it. That we in fact grasp some concept, and that a willingness to make some inference is constitutive of our grasping it, doesn't settle whether we're to reason with this concept, rather than any other, in the first place.

Some brief worries for this sketch are here worth considering. The first concerns whether the concept of *if** is too awkward a concept for us to plausibly understand it as the meaning of "if". A related concern is that the only plausible referent in the neighborhood for "if" is the logical object *if*, not *if**¹⁰. These worries might very well dissuade us from declaring *if** to be expressed by anyone, with the use of the expression "if"; the worry, accordingly, is that the supposedly pressing question—"which concept to use?"—is not a live deliberative dilemma.

In response, the first thing to note is that there are plausible real-world candidates for concepts that might stand in for *if* and *if**. It is no secret that notions of the conditional vary: there are multiple competitors for the concept expressed in English by the expression "if...then". Theorists disagree, for instance, about whether the material conditional or some other notion—concepts that vary in their constitutive inferences—best respect English usage. And, on an inferentialist understanding especially, it seems reasonable to suppose that the relevant participants in the debate *do* grasp different concepts, each with different inferential transitions written into their possession conditions¹¹.

There are, of course, reasons to doubt the thesis that we're dealing in the above with different concepts out of different people's mouths, notwithstanding the appeal to the idea's plausibility just sketched. But critically, *wherever* one stands on the issue of what concepts we

¹⁰ Owing, perhaps, to the relative "naturalness" of *if* rather than *if**. For an elaboration of metasemantic concerns like this one, see Sider (forthcoming).

¹¹ Indeed, Boghossian himself takes this line in his 2020c.

actually do express by our use of “if...then”, such a verdict seems not to foreclose on the question of which concepts we are able to engineer, think, or entertain. For there seems nothing about the concept *if** that should prevent us from asking theoretical questions about what inferences are constitutive of our grasp of it, and practical questions about whether we ought to use it; it strikes me as no different, in these respects, from the “Tonk” connective cooked up by Prior. And so long as it is possible to wonder whether we ought to reason with *if**, as opposed to *if*, the objection sketched above goes through, whether or not *if** is ever actually expressed by our use of the word “if”: so long as it is open to us to deliberate over whether to use *if**, rather than *if*, showing that some inference is a necessary condition on our grasping *if* will not secure our entitlement to it, absent more basic justificatory work as to which concepts to reason with.

That said, a second worry for my critique raises that Boghossian’s proposal is in fact *not* silent on whichever concepts the meaning-entitlement connection applies to. The proposal is explicit in that we are only entitled to the inference-patterns constitutive of concepts that are *themselves* epistemically harmless, or as Boghossian puts it, concepts that are ‘non-defective’ (2020b, p. 42). For there are certain concepts whose possession-conditions include inferences we clearly would be blameworthy in employing. One such a concept is *aqua*, governed by the following introduction and elimination rules: from x is water infer x is *aqua*, and from x is *aqua* infer x is H_2O (p. 39). Certainly, it shouldn’t be that our grasping the concept settles for us whether to believe that water is H_2O , an empirical claim that might prove false. This in mind, we can introduce a conditionalized form of the concept, such that grasping it requires merely our agreement to some such rule as the following: if there is a property which is such that water has it, and if something has that property then it is H_2O , then if x has that property, x is *aqua*. Such a rule allows for the possibility that a thinker is mistaken, and allows others to disagree about

aqua's existing; it seems that grasping this conditionalized form, we *are* blameless in employing the necessary inference (p. 40). But, can all concepts be conditionalized in this way? It would seem, for one, that the *conditional* cannot be, for we would need the conditional to conditionalize on anything to begin with; but then if we have no hope but to use the unconditional form of the concept, even to entertain it, we certainly could not be blamed for introducing it unconditionalized (p. 42). We are blameless in possessing a concept, then, if it was introduced in the conditional form above, or if so introducing it is not an option (p. 41). The meaning-based proposal, accordingly, states the following: we are blameless in employing the necessary conditions on our grasp of just those concepts we are blameless in possessing; since employing MP is a necessary condition on our grasp of *if*, and since we are blameless in possessing *if*, we are entitled to employ MP.

It seems then that we *do* have an account of which concepts are privileged, in just the sense that we are blameless in employing some but not others. Our grasp of *if* is itself blameless, and accordingly we are blameless in reasoning according to modus ponens. But I would like to suggest that the matter is not so neatly resolved. For in virtue of what, exactly, are we blameless in possessing and using the concept of *if*? It cannot be merely the fact that we *do* happen to possess it, for then the objection I sketched above seems especially salient: we've gone no way, appealing merely to the fact that we have some concept, to the normative result of whether to reason with it. It is put forward, rather, that the blamelessness of our possessing some concept lies just in its being introduced in the conditional form sketched above—so that its possession conditions don't "foreclose on the possible falsity of some particular set of claims about the world"—or in its being impossible to so introduce it (p. 41). But then it cannot be that *mere* facts about meaning account for our entitlement to the relevant inferences, for the above is just to say

that further work needs to be done to show that certain concepts are blameless. And what does *this* are substantive normative principles, those that claim that certain concepts (those that foreclose on inquiry) are blameworthy and others (which we can't help but introduce unconditionally) are blameless. The project of showing that our use of *if* is blameless, and our use of a concept like *if** is not, is a project that leans on normative considerations independent from the fact of which concepts we happen to possess and their necessary conditions. And as Enoch and Schechter (2006) note, what we then understand to be carrying the fundamental justificatory load are these normative principles (p. 697); the facts surrounding our grasp of a concept alone don't entitle us to the relevant inferences, because what will settle the question of which concept to use will be these normative theses¹².

A related point is that my challenge appears to succeed even if we suppose that my construal of our original proposal as a transcendental argument is mistaken, such that Stroud's demand on transcendental arguments can't fairly be applied to Boghossian's. For if we take the foregoing to suggest that we cannot account for our entitlement to employ some deductive rule as arising *merely* from facts about concept-possession, but from considerations bearing on what concepts to use in the first place, then nor can we see ourselves as entitled to employ the relevant inference, *merely* because we possess a certain concept. From the perspective of practical deliberation, facts about the necessary conditions for grasping some concept, and the fact that we grasp one rather than another, won't settle whether to adopt one or another. Again, this is not to say that we never have reason to use *if* rather than *if**, or that we have no theoretical candidates

¹² Though I take it that the crucial point is above, as a tentative aside, I'm unsure that the discussed principles are even strict enough to show that we are entitled to reason with *if* rather than a related notion of the conditional like *if**. If we are blameless in possessing *if* because we cannot help but introduce it in its unconditional form (for doing so would require that we use *if*), then it seems too we are blameless in possessing *if**, for the reason that we could not introduce it conditionally without already possessing *if**. If we think we can use an in some sense more basic notion of the conditional to introduce my gerrymandered *if**, then it seems too that we are blameless in possessing *if**, now in the approved 'conditionalized' form.

for explaining in what our justification to use one over the other consists. But as before, establishing this—that *if* is in this sense privileged relative to *if**—requires non-trivial metaphysical work, which will succeed (if it does) by appealing to principles apart from considerations of meaning and concept-possession. If such an account is persuasive, still it remains that facts about a deductive rule’s status as constitutive of some concept, and our grasp of it, are inadequate to secure our entitlement to the relevant inferences. For this, we need *additionally* to settle whether to reason with the desired concept, and this account will likely take the shape of considerations independent of facts about meaning.

IV. A Pragmatist Approach

It is possible to see the following proposal as a kind of complement to Boghossian’s, in the following sense. We left with the claim that settling one kind of fact about a deductive rule—that a defeasible disposition to reason according to it is a condition on our grasp of some concept—is insufficient to settle the question of our entitlement to follow that rule, insofar as settling such a fact leaves open the question whether to reason with the relevant concept. Assuming the inferentialist picture is correct, and the link between concept-possession and our assent to certain inferences is a deep one, then the broadly pragmatist approach sketched in a moment can be seen to offer an account of which concepts we’re entitled to, by way of vindicating our entitlements to employ certain deductive rules in the first place. For those enamored of the inferentialist picture, and Boghossian’s strategy in particular, there seems no insurmountable obstacle to construing the following as in this way supplementary, so long as we are exceptionally clear about a few things. First, the pragmatist approach strives to show that we are entitled to employ certain deductive rules directly, *not* by way of vindicating the concepts of

which (on an inferentialist picture) defeasibly following those rules are necessary conditions; indeed, the proposal succeeds on *any* view about the relationship of concept-possession to inference-patterns. And in this spirit, the proposal differs from one like Boghossian's insofar as, *whatever* facts about meaning and concept-possession obtain, it is *not* these that are understood to be doing the fundamental justificatory work, as in Boghossian's proposal. For these reasons, though a kind of marriage of the two views is plausible, I will follow Enoch and Schechter (2006) in viewing the proposal as instead a replacement of the meaning-based approach.

The proposal as laid out in Enoch and Schechter (2008) begins by articulating a theoretical motivation for giving an account of our justification for employing MP. Some belief-forming methods, among them the deductive rule MP, we presumably are entitled to follow. Some others, like Affirming the Consequent, ostensibly we are not entitled to follow. An account of our justification, then, must describe things such as to be extensionally adequate (it must be that the account endorses our following MP but not Affirming the Consequent), and explain in virtue of what we are justified in employing some methods but not others (3).

The way this is achieved is by observing, first, that there are certain projects are in which "we rationally ought to engage" (4). These Enoch and Schechter term "*rationally required* projects" (4). And though much more will be said in a moment concerning the status of such projects, some tentative, intuitive candidates include those projects that we can't seem to avoid, or those that as rational agents we ought to engage in: consider, for instance, the explanatory project (the project of coming to true beliefs about the world), or the deliberative project (the project of planning for future contingencies).

Importantly, there are certain belief-forming methods whose effectiveness is necessary for our successful engagement in these projects. We might take it, for instance, that the reliability

of MP is necessary for our successful engagement with the project of planning for future contingencies. And the claim is simply that, in such a case, we are defeasibly entitled to employ MP, owing merely to the fact that its effectiveness is indispensable to our successful engagement with the deliberative project (13). Generally, for any belief-forming method, if it is such that it “is possible to successfully engage in a rationally required project by employing it, and such that it is impossible to successfully engage in the project if the method is ineffective, then we are *prima facie* epistemically justified in employing that method as basic, even in the absence of a justified belief concerning the method” (10). With this, ostensibly, we realize our starting aspiration: we can see that we are entitled to follow MP, as its effectiveness is necessary for our successful engagement with a rationally required project, but not Affirming the Consequent, as there is no project which is rationally required for us, successful engagement with which requires that Affirming the Consequent be effective. In virtue of what are we justified in employing MP? The fact that MP is indispensable to a rationally required project is understood to accomplish the fundamental justifying work, and it is owing to this fact alone (recall, MP is understood to be a basic method) that we are justified in following it.

However, the proposal, happily simple to state, begins to suffer the moment we try to say more about what exactly is involved in a project’s being rationally required. To get things rolling, it is worth noting that the modifier ‘rationally required’—or a project’s being “not rationally optional” (10)—pulls intuitively in two different directions: a project’s being ‘required’ or ‘not optional’ suggests something like its unavoidability, while its being ‘rational’ suggests something like a normatively loaded *goal*. Indeed, it is along these two directions that the seemingly available interpretations of Enoch and Schechter’s central claim appear to pull, all of which, when disambiguated, display serious weaknesses.

To show this, I'll begin with the following, clear statement of the proposal: "We are justified in employing any belief-forming method needed for successfully engaging in a rationally required project" (4). Taking our candidate rationally required project to be the deliberative project, we get the following: we cannot successfully engage in the deliberative project unless MP is effective, and it is in virtue of this that we are justified in employing MP. What we might take this thesis to be expressing—again, owing to the ambiguity of the status of 'rationally required'—I'll spell out in the following ways.

[a] Least convincingly, we might take this to express that we are justified in employing MP because the deliberative project would fail if MP weren't effective, and *we do in fact* successfully engage in this project (hence, MP must be effective). Generalized, this way of framing the proposal has it that we are justified in employing any belief-forming method necessary for some rationally required project, where 'rationally required' makes no claim to the status of the project other than the fact that it is one we *do in fact* engage in (like the deliberative project). But of course, for reasons already familiar to us, this won't do. For the Stroud-like worries covered earlier may be re-run here. It may be that a belief-forming method's effectiveness is necessary for our engagement with some project, but if it is merely the case that we simply *do* engage in the project—that is, if it is open to us to abstain from engaging in this project, and to engage in others—then it seems we've gone no way to showing our entitlement to the necessary belief-forming method. The thought is: "who cares if MP is necessary for the deliberative project, if we could just as well adopt the deliberative* project, for which MP is not necessary?". The mere fact that some method is indispensable to a project won't achieve the justifying result we're hoping for, which is not to say that we have no reason available to us to engage in one project over another: but any reason we have to engage in some projects rather

than others will then *itself* be seen to do the fundamental justifying work, not the indispensability claim (as Enoch and Schechter proposed was the case). If a rationally-required project picks out a project that we simply do engage in, then demonstrating the necessity of some method for our successful engagement with it won't show by itself that we're justified in following the method.

[b] I say, however, that the above reading of the proposal is unsatisfying, for the reason that it doesn't do proper service to the 'rationally required' nature of projects, which does more than serve as a name for projects we happen to engage in. Privileging the 'required' sense of the term, we might treat the proposal as claiming that we are justified in employing any belief-forming method necessary for any project we *cannot help* but engage in: it is in this sense that a project is a required one for us. This is a project which we do engage in, but not one that we could neglect to: it is one that is in some sense cognitively unavoidable for us, in the way supported textually by such claims that a project is one without which we can't "successfully make the world intelligible" (10), or that some projects are ones "that the relevant agent (psychologically, or perhaps metaphysically) cannot discard (16). We might treat this as saying that the effectiveness of some method is necessary for our having experience simpliciter, and given our constitution, such experience is something we can't help but have (again, Kant's discussion of the pure concept of space in the Transcendental Aesthetic strikes me as the clearest model); or, we might have it as saying that the effectiveness of the method is necessary for our engagement in a project, which for some reason or other we can't help but engage with. In either case, the upshot is crucially the same: this style of justification is *not* subject to the Stroud-like complaint, that the necessary conditions on a project we could easily discard are not ones we are entitled to. For it seems, in this case, that the fact of MP's indispensability for some project, a project that we have no choice but to engage in, *does* by itself show that we are entitled to follow

it: for we certainly can't be blamed in employing the method; it's suggested we couldn't do otherwise, given our constitution and the status of the relevant project as unavoidable. If the effectiveness of a belief-forming method is indispensable to our success in a project we can't opt out of, it seems owing to this fact alone we are entitled to use the method.

The difficulty for this reading, however, comes in plausibly construing a project as unavoidable, but also such that the effectiveness of a method is necessary for our engaging with it. Take MP: my claim here is that the proposal simply accomplishes too little, if it says that we can see our justification for following MP as arising from its indispensability to a project we can't help but follow. Crucially, following the textual prompts above, the sense in which a project is required here is psychological or metaphysical, as Enoch and Schechter put it; we are dealing with something like a requirement for cognition or experience, not an overtly normative requirement (a project's being required in the sense that we really ought to do it, or that we have some kind of moral obligation to engage in it, will be taken up shortly). But intuitively, it's hard to conceive of a construal of cognition or experience on which success of MP is necessary for it: to say so would seem to render baffling how logicians like Vann McGee (1985) can appear to have discovered counterexamples to modus ponens, unless we are willing to make some such claim as that they are wildly off-base in their interpretation of their own experience. Even if we hold that the effectiveness of MP is necessary for our engagement with the deliberative project (or more narrowly, the project of practical reasoning), still this project seems like one we could abandon: for sure, it would be deeply disadvantageous were we to give up planning for future contingencies—it might be that we *really ought not* give up the project—but there seems nothing psychologically unavoidable about the deliberative project, in the sense required for securing our justification for employing MP.

Will appealing to ‘success’ help? Enoch and Schechter clarify that the relevant criteria is not mere engagement in a project, but *successful* engagement in a project: only a belief-forming method’s indispensability to our *successful* engagement in a project is sufficient to show that we are entitled to employ it. And indeed, they make no such claim that the effectiveness of MP is necessary for mere engagement with the deliberative project: crucial is that the effectiveness of MP is necessary for our successful engagement in the deliberative project (18). The thought here, then, is that if our construal of a project as truly cognitively unavoidable is not robust enough to be such that MP is necessary for our engagement with it, then successful engagement with such a cognitively unavoidable project will be. But whatever successful engagement in a required project looks like, it seems that in adding the condition that we successfully engage in a project, we’ve introduced a normative aspiration or goal, such that one can “achieve success in a sense that supports the claim that one *rationaly ought* to pursue it” (18, italics mine). The worry here is just that the effectiveness of a method might be necessary for our successful engagement in a rationally required (i.e., unavoidable) project, but not necessary for our successful* engagement in the same. Introducing the normative goal of successful engagement with a project (normative in the sense that success is characterized as something we *ought* to pursue) robs the ‘unavoidable projects reading’ of its chief advantage: its immunity from Stroud-like worries, holding that we are not justified in the necessary conditions of a project we can drop. The full dilemma for this reading, then, is as follows: if a project is in some sense cognitively unavoidable, it is dubious that the effectiveness of MP is necessary for our engagement with it; now, though it is more plausible that the effectiveness of MP is necessary for our *successful* engagement with the same project, it seems like “successful engagement with a project” is itself a project we can cease to engage in. For we might characterize a different notion of success in this project, success*,

where it is success* that we ought* to pursue: though the effectiveness of MP is necessary for our realizing the first, it is not necessary for realizing the second. The question of which to realize is left open, and it seems our purported justification of our employment of MP is left wanting.

[c] Interpretively, we have been pulling in the ‘required’ direction in our readings of a ‘rationally required project’. Shifting focus, we might now understand the claim that we are justified in employing any method whose effectiveness is necessary for our engagement with some project—not a project that do or cannot help but engage in, but which is constitutive of rationality. Here the normative nature of the construal is explicit: to account for our entitlement to follow MP, we will now lean heavily on a project’s being one which we ought to pursue, which is just what the appeal to what’s constitutive of rationality is communicating. Enoch and Schechter, in their characterizations of a rationally required project which best support the current reading, write: “If a project is such that any rational agent ought to engage in it (and presumably, ought to try to successfully engage in it), then it is rationally required” (16). The relevant project is one which a rational agent ought to engage in: when in response to the claim that MP is necessary for some such project we ask, “who cares?”, the response is that the project is one we ought to pursue. But the problem here is as follows: even if there are certain “projects that a being must engage in to qualify as a rational agent”, those might not be projects that a being must engage in to qualify as a rational* agent. And the question remains whether to be rational or rational*: for though a project may be constitutive of rationality in a sense that a rational agent ought to pursue it, the same project need not be constitutive of rationality*; in this case, a rational* agent need not pursue it¹³. So, to the question, “in virtue of what are we justified

¹³ It is not at all lost on me that Enoch himself employs a strikingly similar critique against attempts to read off normativity from what’s constitutive of agency: settle what’s constitutive of being an agent, and the question

in employing MP?”, the response can no longer be that the effectiveness of MP is necessary for our successful engagement in a project constitutive of rationality. For we could just as well be rational* agents, such that the effectiveness of MP is *not* necessary for our successful engagement in a project constitutive of rationality.

[d] A final reading in some respect rehearses, though hopefully more clearly, both the formulation and its shortcomings laid out in [c]. That is, we might read Enoch and Schechter’s proposal as claiming the following: we are justified in employing any belief-forming method whose effectiveness is necessary for our successful engagement with a project which—explicit in its normativity—*we really ought to pursue*. But the problem here is one familiar in other clothing, often from its membership in a family of ‘open-question arguments’¹⁴, and which I hope to return to more carefully in conclusion. But for now, it should suffice to observe that settling whether we ought to engage in some project can intelligibly leave open the question of whether to pursue it. For settling that we ought to engage in a project certainly does not settle whether we ought* to engage in it, and whether we regulate our behavior by ought or ought* had better not come in the form of an argument that we ought to use ought: for we ought* to use ought*! At the very least, the familiar moral suggests itself again here: that it is open to us to regulate our behavior according to ought* rather than ought puts pressure on our purported justification. For even if the effectiveness of MP is necessary for some project which we ought to pursue, it is not necessary for some project which we ought* to pursue instead; and it is open whether to pursue the first project or the second. Moreover, it would seem that any account we could come up with, vindicating our pursuit of one project over another, would *itself* be doing the

remains whether to be an agent or a shmagent (2006). Resolving the interpretive unease for which this awareness is responsible, however, is a task for elsewhere.

¹⁴ See, e.g., Dasgupta (2018), Eklund (2017), and, especially here, Clarke-Doane (2020a).

fundamental work of justifying our employment of MP, not the mere claim that MP is indispensable to our successful engagement with such a project.

I take it that the above four readings exhaust, within the space made permissible by concerns for interpretive faithfulness, the plausible readings of the proposal laid out by Enoch and Schechter (2008)¹⁵. Recall, our task was to account for in virtue of what we are justified in employing MP. And the proposed response held that we are justified in employing MP (as we are any belief-forming method), as its effectiveness is indispensable to a rationally-required project; in the case of MP, the relevant project is likely the deliberative project, i.e., the project of planning for future contingencies. But owing to critical ambiguities in what it takes for a project to be rationally-required, it emerged either that the proposal accomplishes too little—i.e., the unavoidable required projects likely don't require that MP be effective for us to engage in them—or they show that the effectiveness of MP is necessary for our successful engagement in a project we might cease to engage in.

V. Conclusion, and a Non-Cognitivist Moral

I have, in the foregoing, laid out two attempts to account for in virtue of what we are justified in employing the deductive rule MP, a meaning-based account and a broadly pragmatist account. The first of these, developed in Boghossian (2020b), suggests that grasping some concept, we are entitled to employ just those inference-patterns that we must defeasibly regard as

¹⁵ To be fair, a fifth plausible reading might be located at p. 17, where an appeal is made to 'primitive normativity', with the claim that a project's status as rationally required might be 'fundamentally' or 'basically' normative. But it seems we can accommodate this in one of the four readings above. Suppose we ask, can there be more than one such primitively normative property? If so, then we still have it that a project may be rationally required but not rationally required* (where the first is primitively normative and the second is primitively normative*), as was the worry for reading [d]. If not—and the only way I can see *this* being the case is if the following is true—then it seems we are claiming that a project's having a primitively normative property amounts to it being unavoidable: but then we land back in reading [b].

compelling if we are to grasp the concept in the first place. I suggested that two styles of objection we ought *not* pursue against the approach above include those that question whether meanings are in fact fixed in the way described, and those which claim that the entitlements read off from concept-possession might one day be defeated. Assuming, then, that the inferentialist picture is correct, are we then to conclude that defeasible entitlements *do* in fact arise from concept-possession? I sketched a reason to suppose that the answer is no. For showing that a defeasible willingness to go along with an inference-pattern is a necessary condition on our grasp of some concept does not settle whether to use that concept, over any other.

A second approach, laid out in Enoch and Schechter (2008), strives to show that we are entitled to employ those belief-forming methods (MP included) whose effectiveness is necessary for our successful engagement in a rationally-required project. But I suggested that resolving the proposal into any one of a series of reasonable interpretations rendered it unpromising, for reasons summarized just above.

Taking a wide view of the foregoing, what I have done is issue two rather local critiques against proposals striving to account for our justification in employing MP, proposals that strive to settle certain facts about it. Grant that the meaning-based account settles one kind of fact about MP: that it's a constitutive inference of a concept we do in fact grasp. I suggested this leaves unsettled the question whether to adopt and reason with the concept for which MP is a necessary condition (hence, it leaves open whether to employ MP). Wondering whether to reason with *if* or *if**, we may find out that MP is an inference constitutive of our grasp of *if*, without having our initial deliberative uncertainty resolved. Likewise, grant that the pragmatist account settles a different kind of fact about MP: that it's an inference the effectiveness of which is necessary for our successful engagement in certain projects. Interpreted such that the relevant project is one we

can chose to adopt or neglect, I suggested that this claim leaves unsettled the question whether to engage with the relevant project, our successful engagement with which requires that MP be effective (hence it leaves open whether to follow MP). Deciding whether to engage in one project over another, we may find out that the effectiveness of MP is necessary for our successful engagement with one, likewise without having our deliberative uncertainty resolved.

Settling certain kinds of facts about MP fail to settle whether to employ it; at least this is how I've so far argued. I now want to suggest a somewhat stronger thesis, one that I take my piecemeal complaints above to illustrate. It is that settling *any* fact about MP, *even the facts about justification*, fails to settle the 'central deliberative question'¹⁶ concerning MP: whether to regulate our beliefs by it (whether to employ it). In one sense the argument is familiarly Moorean: take any 'natural property', and suppose that MP has it: it's a constitutive inference of our grasp of the concept of *if*, its effectiveness is indispensable to our success in a project constitutive of rationality, it's an inference-pattern that licenses inferences which best serve our interests, etc. There still seems to remain a question whether to employ it: a subject has made no mistake in reasoning, "I know that MP's effectiveness is indispensable to my engagement in some project, but I'm still unsure whether to engage in it [and accordingly whether to employ MP]". For, as the thinker might plausibly reason, though the above is true of MP, so too is it the case that MP is *not* a constitutive inference of our grasp of *if**, its effectiveness is *not* indispensable to our success in a project constitutive of rationality*, it's an inference-pattern that licenses inferences which *fail* to best serve our interests*, etc. The practical question merely becomes whether to reason with *if* or *if**, to engage in a project constitutive of rationality or

¹⁶ I owe the term Risberg (forthcoming).

rationality*, to draw inferences that best serve our interests or our interests*; accordingly, the deliberative uncertainty whether to employ MP remains.

The thought, however, is even stronger. Take it that there are facts about justification and entitlement, facts about our being entitled to employ MP, though these may be facts of whatever sort: ‘natural’ or ‘non-natural’, in Moore’s terms. *Even these*, however, fail to settle the practical question whether to employ MP. A subject, it seems, can discover that they are justified in reasoning with some concept (and accordingly are justified in employing its constitutive inferences) and still cogently wonder whether to reason with such a concept. Though I might be justified in pursuing some project, I can introduce a term that diverges in extension, ‘justified*’, and wonder whether I am justified* in pursuing the same. And it is important to note that *whatever the layout* of properties might be, the facts about justification will fail to settle the practical question whether to employ the method: for if I find out I am justified but not justified*, or justified* but not justified, I will simply be faced with question “whether to consult these new properties or the old ones” (Clarke-Doane 2020b, p. 3), that is, whether to regulate my beliefs to respect the justification facts or the justification* facts. Even if the layout of properties is such that I am both justified and justified* (or neither justified nor justified*), in regulating my beliefs according to MP, deliberative uncertainty remains: I can still ask the practical question whether to employ MP (perhaps by noting that it remains open whether I am justified**¹⁷).

Now, some more things need to be said here, if briefly. The first concerns what I mean when I say that settling some fact “fails to settle the relevant practical question”, or “leaves open a practical question”. The most natural way to spell this out is as staking a claim about

¹⁷ This argument is ambitious in another way, too: I intend similar remarks to apply (in a way that I hope is evident enough) to *any* concept or status familiar to epistemology of which it is true that “there is *some* normatively significant, positive, epistemic status a belief can have that is independent of its truth” (Greco 2015, p. 516). These include not only “entitlement” and “justification”, but also “knowledge”, “evidence”, and so on.

conceptual distinctness, after the manner of the following: a thinker has made a distinct kind of mistake in wondering whether “water is water”, but not in wondering whether “water is H₂O”¹⁸. And this distinction need not show anything about *properties*, but does suggest that the concepts of “water” and that of “H₂O” are different. Likewise, a thinker can settle a factual epistemic question without having their practical question settled, as in the now-familiar example: “I know that I am justified in believing X, but I don’t know whether to believe X”; and if we suppose any kind of mistake has been made here, at least it is not the kind of mistake involved in the thought, “I know that I am justified in believing X, but I don’t know whether I am justified in believing X”. It certainly seems as if there are different concepts at work here (and so two strictly different questions). Now, this is all highly intuitive, and there are certainly worries to be pursued and details to be filled out, regarding the kinds of appeals to conceptual distinctness and conceptual mistakes¹⁹ I’ve been making. But for now, the above should at least give a clearer picture of what I intend by the claim that settling some factual question fails to resolve a practical one.

Assuming it is at least plausible that these open-question-type considerations gesture at the presence of two distinct questions, what more can be said about them? The first of these, what I’ve been calling the factual question, is somewhat clear. It is a cognitive question concerning whether something (our holding a belief, perhaps) is justified, where what settles this question is some fact: the fact that we are, or are not, justified. The second question I’ve termed “practical”, owing to the presence of some concept like “the belief-forming method to employ”,

¹⁸ See Gibbard (2003), especially Chapter 1, appeal to the same sorts of considerations, in making a similar argument.

¹⁹ One of these might be a concern that the picture leans on a worrisome analytic-synthetic distinction, insofar as we suppose that there are strict boundaries dictating when we’re using the same concept or a different one. Gibbard (2003) addresses this, with the suggestion that were we unable to distinguish thoughts or concepts from one another at all, it is unclear how we could ever keep track of our thoughts (and ostensibly, we *do* track our thoughts). See also the mention made of this worry in Greco (2015).

or, in what amounts to the same, “the thing to believe”²⁰. Given the question, “is this the thing to believe?”, what settles it? I’ve suggested throughout that it won’t be a fact, or, in a characterization more consistent with the perspective of practical deliberation, it won’t be anything like a belief in a proposition about what one is justified in believing. What *will* settle this question, if anything, is an *intention* or a *plan*. And hence I suggest that the second question, the question that remains when the relevant epistemic facts are settled, is not itself expressing a question about some fact: i.e., it ought not be interpreted cognitively. The proper understanding of this question is characteristically expressivist: it is to be understood as expressing a kind of deliberative uncertainty, whose resolution will be not a belief, but again, an intention or plan. This in mind, I take it that Gibbard (2003) has offered compelling motivation for supposing that a practical concept like something’s being “the thing to do” is crucial for our deliberative lives, that the proper treatment of it is expressivist, and that the concept behaves in ways that we need not find objectionable²¹. The claim, then, is this: the deliberative question that remains after the settling of the epistemic facts involves a concept of “what to believe”, a concept that expresses a plan or intention to believe, and which behaves in all possible ways like Gibbard’s concept of “the thing to do”²².

²⁰ Whether to employ MP is amount to the question whether, believing *p* and *if p, then q*, to believe *q*.

²¹ E.g., that it can be embedded in conditionals and work in inferences, contra the ‘Frege-Geach’ problem.

²² Why not, as Gibbard (2003) goes on to do with some familiar evaluative terms, offer expressivist treatments of our terms like “justification” and “entitlement” themselves—why retain the cognitivist understanding of these and posit the existence of a strictly different question (as Risberg (forthcoming) does with ethical questions)? For this, my motivation has little to do with any particular stance I entertain about how we are to interpret our epistemic terms. That is, it might very well be that the best understanding we can achieve as to the semantics of our epistemic terms, or that the best characterization we can offer of the relevant mental states, is expressivist, cast in the language of plans and intentions rather than beliefs. If this is so, then not much in my account is lost: it will just be the case that we will just be using terms like “justification” and “entitlement” to do the work of the concept, “what to believe”. My point is simply that, *even if* the proper construal of our epistemic terms is cognitivist (and even if we take those beliefs to be true; i.e., even if we take there to be epistemic facts), there will still remain a practical question that remains, whether or not we use our familiar epistemic language to express it. Whatever language we actually use to express them, the two questions: the factual question, and the practical one, are distinct.

Most of the objections that raise themselves here strike me as easily transposable to the key of expressivism about certain ethical concepts familiar from Gibbard (2003), surrounding which there is discussion robust enough to pass beyond the scope of this paper. Yet, tabling concerns familiar to expressivists, there remains still a biting worry for my proposal. That is, I have argued that the concept expressed by the deliberative question is a question concerning “what to believe”, on analogy with, in the ethical case, the practical question of “what to do”? But to show *this* requires reason to suppose that beliefs are akin to actions, in the sense that we have (at least in some cases) control over what to do²³. Do we have control over what we believe? In fact, the problem for my proposal is very deep. For I’ve gestured at a deliberative uncertainty that survives the settling of certain epistemic facts: the fact that MP is indispensable for our engagement with the deliberative project, I’ve argued, fails to settle the question whether to employ MP (that is, whether to conclude q beliefs from the beliefs p and *if p , then q* ²⁴). But if we take it that we have no control over what beliefs we hold (or, no control over what propositions we believe), then it seems that there can be no deliberative uncertainty to begin with, no question whether to believe or not. Another way of putting this is that the question of which norms to consult by which to regulate our beliefs—e.g., the question whether to consult justified or justified*—is not up to us.

The familiar name for this is the question of where some form of “doxastic voluntarism” is true, where this is simply the thesis that we can believe that q by intending to believe that q . If

²³ Risberg (2022) notes this as a problem for an expressivist approach to epistemic terms on which, along the lines of Gibbard (2003) and Clarke-Doane (2020a), the relevant mental state is an intention or plan.

²⁴ The worry is live, too, for the kind of deliberative uncertainty (“whether to use *if* or *if**”) that I allege remains when certain meaning-facts are settled. Can I decide which concept to grasp? Recall that on the inferentialist framework, concept-possession requires a defeasible disposition to find certain inferential transitions like MP compelling, where this consists in concluding q from p and *if p , then q* . If believing that q is not in some respect a matter of voluntary control, then neither is the question of which concept to grasp. The deliberative uncertainty I claim remains seems impossible.

false, it suggests a serious worry for the plausibility of the deliberative questions I've suggested survive our settling the epistemic facts; at least, this threatens the view that we can resolve a question what to believe (if there is such a live question) by forming an intention to believe, as I've suggested is the case. Now, my response to this will actually remain agnostic on the question of whether this thesis is true, and questions merely what we are to conclude about my remarks if it is. For either doxastic voluntarism is true or it isn't. If true, then naturally the worries are dissolved. If it can be a live deliberative question for us whether to believe q or not, and the question whether to believe q is one resolved by our intentions to do so, then in just the sense that we can resolve to do something we can resolve to believe something; the practical uncertainty in the epistemic case is akin to that in the ethical case, and accordingly, a similar expressivist treatment recommends itself for both. And this seems compelling even if only some form of *indirect* doxastic voluntarism is true: on this view, though it may not be the case that our intention to believe that q by itself allows us to believe that q , we have voluntary control over whether we believe that q , if, e.g., our intentions lead us to seek out evidence to confirm q and if we form our beliefs according to the evidence. For again, any voluntary control we have over our beliefs allows for the question whether to believe q to be a live one for us, for it's at least up to us whether to form the intention to believe q .

Now, suppose doxastic voluntarism (direct or indirect) is false: it is not the case that by forming an intention to believe that q , we can come to believe that q . The consequences, again, are that an intention expressed by the concept of "the thing to believe" will never resolve a question of what to believe, and as I've suggested, there's a sense in which there's no real deliberative question for us anyhow (if whether to believe is not up to us). My only comment is that, if this view of belief is accurate, then it is not just my own picture that suffers, but even the

factualist approaches to epistemology reflected in the efforts of Boghossian (2020b) and Enoch and Schechter (2008). For it is not just that the question of what norm to consult to regulate our beliefs—justified or justified*—is not a question we can ask. There is also the question of whether we can even consult the norm that a meaning-based theorist thinks is authoritative: if what beliefs we hold is not a matter of our voluntary control, then there is hardly a sense in which we are consulting anything in coming to a belief. Certainly, on this picture, it's still the case that some norm governs our coming to a belief, but it is not the norm characteristic of practical deliberation: it might be a descriptive norm that accounts for what actually occurs when we arrive at a belief, but is not something that I can consult in coming to a decision. Is this a problem? Perhaps not, if we think that epistemology is merely concerned with these kinds of descriptive norms. But I take it that our interest in justification or entitlement stems from their putative connection with questions of 'policy'²⁵. I take it that part of the motivation for discovering that we are justified in employing MP, even on a realist construal of the justification facts, is to resolve a question of how to go on: discovering we are or are not justified in employing MP is important, because this is a fact which we, as deliberative agents, consult in deciding whether to believe q from p and *if p , then q* ; it is a fact we consult as a *reason* to go on in a certain way. Which is all to say that the picture of epistemology as in any style *normative*, a picture that I take even most realist or cognitivist proposals to endorse (in the sense that the concern for facts of justification, realistically construed, flow from a concern for giving us reasons to go on in a certain way), seems seriously threatened by the view that we have no voluntary control over what beliefs we come to. If this view is correct, then there is in a way no recognizable sense in which an epistemic deliberation ("whether to believe q ", or "whether we

²⁵ The sense of this term I take from Clarke-Doane (2020b).

are justified in believing q ”, on some normative construal of ‘justified’) is live for us. Whether we take it that settling a fact does or does not succeed in settling the questions of policy—which is how I distinguish the proposals discussed from my own—the sense in which we settle questions of policy at all strikes me as worrisome, if doxastic voluntarism is false.

In a similar spirit, allow me to close with the following. I have made as a central thesis of the earlier portion of this paper the following remark: certain meaning-based and pragmatist efforts to justify our employment of MP are thwarted by the remainder of the kind of deliberative uncertainty I’ve mentioned, after facts of a certain kind are settled. An objection to this is that I’ve shown no such thing. For a certain stripe of cognitivist about our epistemic thought and talk might very well say, contra the gestures I’ve made above, that whether or not an epistemic fact settles a practical question is not the business of epistemology. It might very well be that the layout of facts is such that we are justified in engaging in projects constitutive of rationality, and justified* in engaging in projects constitutive of rationality*; if true that the effectiveness of MP is necessary for our engagement with the first but not the second, the proposal from Schechter and Enoch is still successful, so this view has it: for it has shown that we are justified (even though not justified*) in employing MP, and any questions expressing the kind of deliberative uncertainty I’ve been discussing are a change of subject. To this, my response is, “have ‘epistemology’”. Even if we think we ought to be realist about matters of justification, I have suggested that there remains a pressing deliberative question that the proposals above (and any like them) do not resolve by settling the facts, even facts about justification. And as I suggested above, I take it that the resolution of these practical questions is not only pressing for our deliberative lives construed broadly, but is a reason for our interest in epistemic questions to begin with: it seems that a hope of investigating and consulting the epistemic facts has been that

we will thereby achieve some guidance for the ‘how to go on’ question of what to believe²⁶. If we nonetheless refrain from counting this a question of proper epistemology, so be it. The upshot is still that a substantial deliberative question—“what to do”, “what method to employ”, “how to regulate my beliefs”—survives settling the facts, even the facts about justification on a realist construal of those.

²⁶ “It would be bewildering if an epistemologist were to say: ‘the evidential support relation is R, but do not worry about having beliefs that bear R to the evidence’” (Clarke-Doane 2020b, p. 2).

References:

Boghossian, P. (1994). Inferential role semantics and the analytic/synthetic distinction.

Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition,
73(2/3), 109-122.

Boghossian, P. (1997). Analyticity. In B. Hale & C. Wright (eds.), *A Companion to the
Philosophy of Language* (pp. 331-386). Blackwell.

Boghossian, P. (2020a). Analyticity reconsidered. In *Debating the A Priori* (pp. 1-27). Oxford
University Press. (Reprinted from *Nous*, 1996, 30(3), 360-91).

Boghossian, P. (2020b). Blind reasoning. In *Debating the A Priori* (pp. 28-45). Oxford
University Press. (Reprinted from *Aristotelian Society Supplementary Volume*, 2003,
77(1), 225-48.).

Boghossian, P. (2020c). Williamson on the a priori and the analytic. In *Debating the A Priori*
(pp. 78-85). Oxford University Press. (Reprinted from *Philosophy and
Phenomenological Research*, 2010, 82(2), 488-97).

Burge, T. (1995). Content preservation. *Philosophical Issues*, 6(Content), 271-300.

Casalegno, P. (2004). Logical concepts and logical inferences. *Dialectica*, 58(1), 395-411.

Clarke-Doane. (2020a). *Morality and Mathematics*. Oxford University Press.

Clarke-Doane. (2020b). Epistemic non-factualism and methodology. In (ed.), M. Klenk, *Higher
Order Evidence and Moral Epistemology* (pp. 256-265). Routledge.

Dasgupta, S. (2018). Realism and the absence of value. *Philosophical Review*, 127(3), 279-322.

Ebert, P. (2005). Transmission of warrant-failure and the notion of epistemic analyticity.

Australasian Journal of Philosophy, 83(1), 505-521.

Eklund, M. (2017). *Choosing normative concepts*. Oxford University Press.

- Enoch, D. (2006). Agency, shmagency: Why normativity won't come from what is constitutive of action. *The Philosophical Review*, 115(2), pp. 169-198.
- Enoch, D. & Schechter, J. (2006). Meaning and justification: The case of modus ponens. *Noûs*, 40(4), 687-715.
- Enoch, D. & Schechter, J. (2008). How are basic belief-forming methods justified? *Philosophy and Phenomenological Research*, 74(3), 547-579.
- Gibbard, A. (2003). *Thinking How to Live*. Harvard University Press.
- Giovannini, E. N. & Schiemer, G. (2021). What are implicit definitions? *Erkenntnis*, 86(1), 1661-1691.
- Greco, D. (2015). Epistemological open questions. *Australasian Journal of Philosophy*, 93(3), 509-523.
- Hale, B. & Wright, C. (2000). Implicit definition and the a priori. In P. Boghossian & C. Peacocke (eds.), *New Essays on the a Priori* (pp. 286-319). Oxford University Press.
- Horwich, P. (1997). Implicit definition, analytic truth, and aprior knowledge. *Nous*, 31(4), 423-440.
- McGee, V. (1985). A counterexample to modus ponens. *Journal of Philosophy*, 82(9), 462-471.
- Peacocke, C. (1992). *A Study of Concepts*. MIT Press.
- Peacocke, C. (1993). How are a priori truths possible? *European Journal of Philosophy*, 1(2), 175-199.
- Risberg, O. (2022). Meta-skepticism. *Philosophy and Phenomenological Research*, 1-25.
<https://doi.org/10.1111/phpr.12871>.
- Risberg, O. (Forthcoming). Ethics and the question of what to do. *Journal of Ethics and Social Philosophy* (forthcoming).

Sider, T. (Forthcoming). Dasgupta's detonation. *Philosophical Perspectives*.

<https://philpapers.org/rec/SIDDD>.

Stroud, B. (1968). Transcendental arguments. *Journal of Philosophy*, 65(9), 241-256.

Williamson, T. (2020). Understanding and inference. In *Debating the A Priori* (pp. 46-77).

Oxford University Press. (Reprinted from *Aristotelian Society*, 2003, 77, 249-93)