

Dear Seminar Participants,

The paper is quite rough in several ways. For one thing, the paper was originally concerned only with issues about reliability. I have recently decided to revise the paper to address issues concerning justification. In this draft, these issues remain mostly unexplored.

I also probably should do more to explore the possibility of what I call “inferences with the wrong structure”. For example, I should better develop the discussion in the legal domain of the possibility of a kind of induction from a proper subset of the determining facts (perhaps plus other facts) to the target facts.

Another point: I need to clarify the modal strength of the tracking thesis with respect to law.

Here are two questions that I would love to get your feedback on. First, in section 2, I distinguish three senses of tracking. I make clear that, in my central sense, one need not think of the more basic determining facts *qua* more basic determining facts. I wonder whether I need to address the further question whether there is a certain way in which one must think of the more basic determining facts, e.g. under some canonical description or at least under an accurate one. For example, if semantic facts are determined by use facts, in order for a method of ascertaining the semantic facts by inferring them from the use facts to qualify as tracking, must one think of the use facts as *use facts*?

Second, when I wrote the legal section, I tried for the most part to avoid substantive questions about the metaphysics. I wonder whether it would be helpful at various points in the discussion to explicitly assume various views about the metaphysics of legal facts for the sake of argument. I do a bit of this, but I could do more to argue, for example, that my thesis is plausible on several different leading views of the metaphysics.

I look forward to the seminar.

Best regards,

Mark

NONBASIC EPISTEMOLOGY: MUST THE EPISTEMOLOGY OF A NONBASIC DOMAIN TRACK ITS METAPHYSICS?

Mark Greenberg, UCLA¹

1. Introduction

The paper introduces a topic, which it labels loosely *nonbasic epistemology*. Somewhat more precisely, the topic concerns whether, and if so, under what conditions, the epistemology of a nonbasic domain must *track* its metaphysics.² The explication of this central notion of tracking is an important part of the project.

Some domains are nonbasic in the sense that the facts of those domains obtain in virtue of more basic facts. The mental, the semantic, the legal, and, more controversially, the moral are potential examples of nonbasic domains. Microphysics, mathematics, and perhaps morality are candidates for basic domains, but plainly much of our knowledge is of nonbasic domains.

My topic concerns the relation between how the facts of nonbasic domains are (metaphysically) determined and how we can ascertain these facts. (To avoid confusion, I use “determine” and its cognates exclusively metaphysically. When I want to talk epistemically, I use “ascertain,” “figure out,” or other unambiguously epistemic terms.) More specifically, I

¹ I am grateful to Tyler Burge, Sam Chilovi, Sinan Dogramaci, Michael Rescorla, Scott Shapiro, Seana Shiffrin, David Sosa, and Tim Williamson for valuable discussion. I am especially indebted to Tyler Burge for detailed written comments on an ancestor of this paper. I would also like to thank participants in the Annual Analytic Philosophy Symposium (at the University of Texas, Austin), the UCLA philosophy colloquium, and the UC Santa Barbara philosophy colloquium for helpful comments.

² This label is not intended to evoke Nozickian tracking. As explicated in the text, what is being tracked is the metaphysics (though plausibly a method that successfully tracks the metaphysics ensures tracking in Nozick's sense).

focus on a cluster of questions concerning whether and under what conditions we need to infer the facts of a domain from the more basic determining facts.

There are some relatively clear instances of non-inferential access to the facts of a domain. One's access to one's own mental states and perception of the external world are good examples. In these cases, tracking is unnecessary: because of our non-inferential access to the target facts, we don't have to infer the target facts from the more basic facts that determine them. But there are difficult questions about other domains. In this paper, after clarifying the issues, I take the legal domain as a case study.

A parenthetical: The topic introduced in the paper is, I hope, of intrinsic interest. But, in addition, the topic has great practical import in the legal case. Law has tremendous importance in the United States in part because of the pervasive role that it plays in settling controversies in almost every area of life. I argue elsewhere that the thesis that the epistemology of law must track its metaphysics has extensive implications for debates over legal interpretation. The thesis yields a much-needed framework for evaluating theories of legal interpretation; given that framework, many influential theories can be shown to be on shaky ground.

Section 2 introduces and clarifies the question of whether the epistemology of a nonbasic domain must track its metaphysics. I explicate tracking, distinguishing the relevant notion from related ones. And I classify and explore alternatives to tracking. Section 3 turns to the legal domain and argues that it is unusual in that its epistemology must indeed track its metaphysics.

2. Explicating the topic

a. Tracking

On a common view, mental facts obtain in virtue of physical facts. Semantic facts are often thought to obtain in virtue of facts about the use of words. Legal facts obtain in virtue of, among other things, facts about what various people have done and decided in the past, facts about the meaning of certain words or utterances, and, more controversially, moral facts such as facts about democracy and fairness.

It should be clear that, though these determining facts may be more basic than the facts of the target domain, many or all of them are themselves not among the most basic facts of the universe. (I use *target domain* as a shorthand for the nonbasic domain access to which is under discussion and *target facts* for the facts of that domain.) In this paper, talk of “more basic facts” refers to facts that are more basic than the relatively high-level target facts under discussion, not to the most basic facts of the universe.

The topic that I introduce in this paper concerns how our access to the facts of non-basic domains relates to the metaphysics of those domains. In *Being Known* (1999), Christopher Peacocke fruitfully explored questions about the responsibility of epistemology to metaphysics. My question concerns a specific kind of responsibility of epistemology to metaphysics that I will now explicate. I will say a little below about how my project relates to Peacocke’s.

Because the target facts are, by hypothesis, determined by more basic facts, one obvious route to ascertaining the target facts goes via the relevant more basic facts (and the way in which

they determine the target facts).³ Indeed, on a straightforward line of thought, one way to ascertain the facts of a nonbasic domain follows willy nilly from the way in which the facts of that domain are determined. Suppose that the X facts are determined by more basic A facts, B facts, and C facts. To put it schematically, let's say that the X facts are a particular function f of those facts -- $f(A,B,C)$. That is, f maps sets of A, B, and C facts to target facts. On the straightforward line of thought, then, one way to detect the X facts is simply to detect the A, B, and C facts and infer the target facts in a way that conforms to the mapping given by f .

To take a simple example, suppose that we want to know the U.S. trade deficit for 2016. Trade deficit facts are not basic facts about the universe. In order to follow the straightforward approach, we need to know the more basic facts that determine the U.S. trade deficit. The trade deficit is the excess of imports over exports. Thus, in order to figure out what the trade deficit is for 2016, we need to figure out the total of U.S. imports and the total of U.S. exports for 2016 and then subtract the latter from the former.

This is a bit of a toy example because the stipulated account of trade deficit facts is not a serious metaphysical account of a domain. Here's another simplistic example. Suppose that what makes it the case that one is dreaming is that one's alpha fibers are firing at 30 hertz. In that case, on the straightforward approach, we can infer from the fact that someone's alpha fibers are firing at 30 hertz that they are dreaming.

Here's a more serious example. Suppose for the sake of argument that HLA Hart's well-known theory of the metaphysics of legal facts is true. (Let a legal fact be a fact about the content of the law, for example the fact that, in California, contracts for the sale of land are not valid

³ As noted in the introduction, I use *determine* (*determining*, *determination*, and so on) in its metaphysical sense throughout.

unless they are in writing.⁴) On Hart's theory, roughly, the convergent practices of judges and other legal officials determine both which legal materials are relevant and exactly how those legal materials together determine the content of the law. For example, if judges, say, treat customs as relevant to the content of the law in a particular way, then customs determine legal facts in exactly that way.⁵

On the straightforward approach, one would ascertain the legal facts simply by inferring the legal facts from the facts about the way in which judges treat legal materials as bearing on the content of the law. Suppose, for example, that in the local jurisdiction judges treat the semantic content of statutes as constituting legal norms. In that case, we would infer from this practice of judges and from the fact that there is a statute with the semantic content that the sale of marijuana is prohibited to the conclusion that there is a legal prohibition on the sale of marijuana.

Part of the project is to elucidate the intuitive notion of tracking illustrated in the foregoing examples. We need to distinguish at least three possible understandings of tracking. There is, first, what I will call the *trivial* understanding. Given that mental facts, say, are determined by more basic facts, there is a sense in which any method that correctly identifies the mental facts is *ipso facto* tracking the more basic facts that determine the mental facts. More generally, given that the facts of any nonbasic domain are determined by more basic facts, there is a sense in which any method that accurately identifies the nonbasic facts is *ipso facto* tracking the determining facts. This sense of tracking is trivial because, in this sense, to track the determining facts just *is* to ascertain the target facts.

⁴ Legal facts need to be indexed to a jurisdiction and a time, but I will generally omit this qualification for ease of exposition.

⁵ For brief elaboration of Hart's account, see Greenberg (2017).

On the second understanding of tracking, for a method of ascertaining the target facts (of a nonbasic domain) to track the metaphysics of the domain is for the method to involve *inferring* the target facts *from* the more basic facts. (Immediately below, I further refine this characterization to make clear that the method of inference must be such that each target fact is inferred from all and only the more basic facts that determine it.) On the second understanding of tracking, however, the method of ascertaining the target facts need not make reference to the more basic facts *under that description* or *as such*. In other words, one who employs the method need not think of the more basic facts *qua* more basic determining facts (and more generally need not think of the metaphysics of the target facts *qua* metaphysics). Indeed, one can employ the method without even having concepts necessary to think of the more basic facts in this way, such as the concept of metaphysically more basic.

The third understanding of tracking – *strong* tracking – differs from the second precisely in that it requires that the method of ascertaining the target facts involve inferring the target facts from the determining facts *under that description*.

The second and third notions of tracking are the interesting ones. My general question about when and to what extent the epistemology of a non-basic domain must track its metaphysics could be asked with respect to either notion. My focus will be primarily on the second notion of tracking, however. (When I use the term *tracking* without qualification, intend the second notion of tracking.) For example, when I turn in the next section to a case study of the legal domain, my thesis is that a method of ascertaining the legal facts must involve inferring the legal facts from the determining facts, but not that the inferences must involve thinking of the determining facts *as* determining facts. At at least one point, however, the argument touches on the stronger thesis that the inferences must involve strong tracking. (In brief, issues concerning

reliability support the need for tracking, while considerations of (internalist) justification push in the direction of strong tracking.)

I need to add an important clarification of the idea that tracking requires inferring the target facts from the determining facts. Sometimes it is convenient to say simply that the target facts are determined by facts of particular types or domains. For example, one might claim that semantic facts are determined by facts about use. The clarification is that, in order to qualify as tracking the metaphysics, it is not enough that one infers the target facts from some facts or other of the determining type or types. For example, assuming that semantic facts are determined by use facts, not just any inference from use facts to semantic facts qualifies as tracking. Rather, one must infer semantic facts from all and only the facts that determine those semantic facts. I will sometimes refer to this requirement as the requirement that inferences to the target facts *have the right structure*.

Let me elaborate. According to a causal/informational theory of meaning, very roughly, a symbol refers to a property in virtue of a causal link between the property and the symbol's tokening. If such a theory is true, in order to track the metaphysics, inferences from the use of a symbol to its meaning would have to have the following structure: from the fact that the tokening of a symbol "s" is reliably caused by a property p, one would infer that "s" means P.

By contrast, consider Paul Horwich's (2008) Use Theory of Meaning. That theory claims that the meaning of a word is determined by a certain aspect of the use of the word – a "basic acceptance property." Horwich offers examples of the basic acceptance property for various words. For example, according to his theory: "'true' means what it does to us in virtue of the fact that the law governing its use is that we are prepared to provisionally accept any instance of the schema, '<p> is true ↔ p'." Again: "'bachelor's' meaning is engendered by the fact that its

basic regularity of use is our acceptance of the sentence, ‘the bachelors are the unmarried men.’” (Horwich 2008, 26.) According to the theory, however, nothing general can be said about the relation between the basic acceptance properties and meanings. Each basic acceptance property simply determines the appropriate meaning.

If Horwich’s theory is true, then to track the metaphysics, inferences from use facts to semantic facts would have to mirror the mapping from basic acceptance properties to meaning facts specified by Horwich’s theory. For example, from the fact that a word has the basic acceptance property for “red” specified above, we would infer that the word means RED. And so on for the other basic acceptance property/meaning mappings specified by the theory.

The lack of generality of the relation between determining facts and meaning facts on Horwich’s theory contrasts sharply with the metaphysics of meaning as characterized by a causal/informational theory. According to a causal/informational theory, there is a completely general relation between the determining facts and the meaning facts. Consequently, as we saw, if an informational theory is true, we can specify in general terms what a method of ascertaining the meaning facts must be like in order to track the metaphysics. But even if the metaphysics of a domain is such that nothing general can be said about the way in which the determining facts relate to the target facts, we can say that in order for a method of ascertaining the target facts to track the metaphysics, the inferences must move from all and only the determining facts to the relevant target fact.

Incidentally, one might question whether it is possible to reliably infer the target facts by inferences that lack the right structure. But there could, for example, be an innate or acquired heuristic that, perhaps by exploiting statistical regularities, permits reliable inferences from use facts to meaning facts, though the use facts from which an inference proceeds include only some

of the use facts that determine the meaning fact in question (and perhaps other use facts that are not among the determining facts).

Now there might be different ways of characterizing which facts are the more basic determining facts (and how they determine the target facts) that are merely notational variants. In the schematic example of the X facts obtaining in virtue of the A , B , and C facts, we might be able to describe the A facts' bearing a certain relevant relation to the B facts as the obtaining of the D facts. In that case, it might be possible to equivalently describe the metaphysics of the X facts by saying that the X facts are a function g of the C facts and the D facts. (Among other things, the D facts would have to be more metaphysically basic than the X facts.) A method of inference that reflected the $f(A, B, C)$ account of the metaphysics would be different, however, from a method of inference that reflected the $g(C, D)$ account. One could employ the latter method of inference, for example, without even having access to the A and B facts. What counts as tracking the metaphysics should not depend on mere notational differences. Therefore, if the two descriptions of the metaphysics are mere notational variants, both methods of inference qualify as tracking.⁶

b. Whether tracking is necessary: alternatives to tracking

I take it that tracking is a good method of ascertaining the facts of a nonbasic domain when it is feasible. But *must* the epistemology of a nonbasic domain track its metaphysics? Our ability to obtain knowledge of a domain other than by tracking admits of degrees both because

⁶ This point raises the question, which I can't pursue here, whether there might be a *canonical* account of the metaphysics of a domain; if so, the difference between otherwise equivalent descriptions might not be merely notational, for one account would have a privileged status.

how much knowledge we can acquire without tracking the metaphysics is a matter of degree and because epistemic methods can vary with respect to how much they diverge from tracking, e.g., the extent to which they base inferences on fewer than all of the determining facts. Thus, for each nonbasic domain, we can ask whether and to what extent its epistemology must track its metaphysics. The answer to this question will vary from domain to domain, as we will see. There may be complex intermediate cases; for instance, there may be domains in which we can obtain some knowledge without tracking but further knowledge requires tracking.

In the next section, I will take the legal domain as a case study and argue provisionally that, unlike many nonbasic domains, its epistemology must track its metaphysics. This statement of the thesis is shorthand for a more nuanced claim: the reliability of a method of acquiring knowledge of legal norms depends on the extent to which it tracks the metaphysics. A method of acquiring knowledge of legal norms cannot be a reliable method in general without fully tracking the metaphysics. Partial tracking of the metaphysics may allow limited pockets of knowledge, for example knowledge of the legal facts locally.

The fact that the necessity for tracking varies from domain to domain raises a further theoretical challenge: what is it about a domain (and our relation to it) that has the consequence that its epistemology must track its metaphysics? At the end of the paper, I will offer some tentative suggestions about what it is about the legal domain that makes tracking necessary.

We need a preliminary clarification of the claim that the epistemology of a domain must track its metaphysics. We can illustrate the point with the legal case: There is of course a way of finding out what the law is that does not involve tracking the metaphysics – one can ask an expert or look it up in a treatise or on the Internet. But such testimonial sources of knowledge of legal facts cannot be the *primary method* of discovering legal facts. An expert has to figure out

what the law is in order to be able to tell you, and ultimately the expert needs a way of finding out what the law is that does not rely on testimony. In other words, finding out what the law is by asking an expert is parasitic on a primary method of ascertaining the content of the law. In general, the thesis we are interested in is that the *primary* method of ascertaining the facts of a domain has to involve tracking.

Another preliminary point is that there are several ways in which the issues in this area depend on substantive claims about the metaphysics of the target domain. First, whether a domain is basic obviously depends on its metaphysics. Whether one takes mathematics or morality to be a nonbasic domain, for instance, will depend on one's view about its metaphysics. Second, whether a way of ascertaining facts counts as tracking the metaphysics depends on which facts are the determining ones (and on how they combine to determine the target facts). For example, if facts about intentional mental states were determined by behavioral facts, then inferring facts about others' mental states from their behavior might qualify as tracking. Third, most importantly, whether one finds it plausible that tracking is necessary in a given domain will depend on one's views about the metaphysics of the domain. For example, whether the kind of heuristic that I mentioned earlier (that permits reliable inferences from a proper subset of the determining use facts to meaning facts) is feasible likely depends in part on the metaphysics. The metaphysics of a domain might be such that one could reliably use a kind of induction to infer the target facts from only a proper subset of the determining facts. For these reasons, it's not possible to conduct the discussion in a way that entirely avoids substantive metaphysical questions about the relevant domains, and I will sometimes suppose the truth of particular metaphysical theories for purposes of argument.

I can now say a bit more about the relation of my project to Peacocke's project in *Being Known*. Peacocke asks how can reconcile our understanding of what is involved in the truth of claims in particular domains with our understanding of how we obtain knowledge in those domains. My topic concerns one specific and straightforward way in which, in a nonbasic domain, the epistemology could be reconciled with the metaphysics. Specifically, the epistemology could simply track the metaphysics.

Peacocke considers the possibility that in a given domain, the only way to reconcile the epistemology with the metaphysics is revisionary. For example, we might be forced to pare down our understanding of what is involved in the truth of claims in the domain. Or, at an extreme, we might be forced to the conclusion that we cannot have knowledge of the facts of the domain or even that the domain is spurious in the sense that there are no truths of the domain. My inquiry will generally begin from the provisional assumptions that a domain is nonbasic and that there are truths of the domain. Moreover, we will generally have some views about what the metaphysics of the domain is like and some views about what kind of abilities to access the target facts we possess. But if those assumptions led, in a particular domain, to the conclusion that tracking was the only way in which we could obtain knowledge of the domain, that conclusion might lead us to go back and revise our initial assumptions. For example, our confidence in our knowledge of a domain and in the infeasibility of tracking in that domain (because, e.g., of our lack of access to the metaphysically determining facts) might be stronger than our confidence in the initial views that led us to the conclusion that tracking was necessary. So the conclusion that (given our initial views) the epistemology must track the metaphysics might lead us to doubt our initial views about the metaphysics of the domain or about the available mechanisms for obtaining knowledge of that domain.

What are the alternatives to tracking? Let's distinguish three broad categories of alternatives. First, we have domains in which we have reliable non-inferential knowledge of the facts of the domain. (I include in this category cases in which we have non-inferential knowledge of *some* facts in the domain and infer knowledge of *further* facts in the domain from the former facts. Morality is a plausible example.) Call such cases ones of *direct* (as opposed to inferential) knowledge.

Second, we have domains in which our knowledge is *inferential*, but the relevant inferences are not inferences from more basic determining facts but from facts of some other sort.⁷ The idea is that we could infer the target facts from *mere* evidence – evidence that does not bear a constitutive relation to the target facts. (I'm tempted to restrict the use of the term *evidence* to mere evidence in this sense.) Call such cases ones of *lateral* knowledge.

Third, there could be domains in which our knowledge is based on inference from more basic determining facts, but the inferential structure does not reflect the way in which the determining facts determine the target facts. That is, the inferences do not move from all and only the facts that determine particular target facts to those target facts.

The four-way classification is exhaustive: knowledge that is not by inference; knowledge by inference from facts other than more basic determining facts; knowledge by inference from more basic determining facts that does not have the right structure; and knowledge by inference from more basic determining facts that has the right structure. I suspect, however, that some will find the classification too neat. There might, for example, be borderline cases. Also, as a

⁷ What about inferring target facts from other facts of the target domain? This cannot be the ultimate method of ascertaining the target facts because it depends on having access to the initial target facts from which the inference proceeds.

classification of domains, the classification is not exclusive. A given domain could be such that knowledge by more than one of these routes is feasible.

It is important to note that, in order for a method to be a genuine alternative to tracking, the method must not be parasitic on tracking. That is, for example, it must not be that we first infer the target facts from the more basic determining facts (by an inference with the right structure) and then use this knowledge of the target facts to validate a putative alternative method.

In the direct category, an example is our first-person access to our own mental states. Humans have more or less reliable non-inferential access to our own mental states, such as beliefs, emotions, and pains. Again, perception gives us reliable non-inferential access to, for example, the properties of middle-sized physical objects.

If mathematics is a non-basic domain, it provides another example. We are adept at amassing mathematical knowledge, despite the fact that the metaphysics of mathematics is highly controversial. Mathematicians use knowledge of mathematical principles, such as mathematical induction, to infer further mathematical truths; they do not begin from metaphysically more basic facts.

The familiar, though much disputed, idea that our grasp of meaning can yield knowledge provides a very different kind of example. Suppose, for example, that our grasp of the meaning of moral terms provides a route to knowledge of basic moral truths. On the natural assumption that this route from grasp of meaning to knowledge is non-inferential, this route would be a direct way of obtaining moral knowledge. (If, on the other hand, the route involved inference from facts about the meaning of moral terms to moral truths, the question would arise whether

the relevant meaning facts were metaphysical determinants of the moral truths. If so, the route could qualify as tracking; if not, the route could qualify as a lateral alternative to tracking.)

It is more difficult to give clear examples of lateral knowledge that are a genuine alternative to tracking. It is certainly common to infer facts from mere evidence, but such inferences are often parasitic on an independent way of verifying their reliability such as perception. One promising kind of case involves inferences from effects to causes. Setting aside the case of testimony, much of our knowledge of the mental states of *other* people is inferred from their behavior (including verbal behavior, which may include a great deal of relevant evidence that is not simply true testimony about the speaker's mental states). Such inferences are plausibly often inferences from effects to causes – from the behavioral effects of mental states to their mental causes.

Now, it is conceivable that some such inferences involve tracking. For example, one might use behavior as evidence of functional facts that determine mental facts and then infer the mental facts from the functional facts. Typically, however, inferences from the behavior of others to mental facts likely proceed without tracking. We probably have a (largely tacit) understanding of the typical causal impact of mental states sufficient to allow us to infer mental facts from behavioral facts without those inferences proceeding via judgments about functional (or other) facts that determine mental facts.

Another option is that we could use a method that might be analogized to Neurath's boat. (Reflective equilibrium in moral philosophy is another parallel). The rough idea would be that we simply find ourselves with beliefs about the target facts, beliefs about our abilities to discern the target facts, and other beliefs about the world. We then engage in a process of reconciling those beliefs as best we can in a quest to maximize coherence, explanatory power, and other

epistemic values. Such a method at least ostensibly relies both on non-inferential access to the target facts – we start with some beliefs about the target facts, which must be somewhat reliable in order for the overall route to provide reliable access to the target facts – and on inferences that are not inferences from the determining facts (the holistic process of deciding how to modify one’s beliefs presumably involves inference). On the other hand, we could press on the source of the warrant for the initial beliefs; for example, are they based on testimony and therefore parasitic on some other source of knowledge, which might involve tracking? I address such issues in the case of law in section 3.

I don’t have a convincing concrete example of reliable access to the target facts by inferences with the wrong structure. But I suggested above that the metaphysics of a domain might be such that a reliable heuristic could make it possible to infer the target facts from only a proper subset of the determining facts.

It is of course easy to illustrate an *unreliable* method that relies on inferences with the wrong structure. According to the *moral impact theory* of law, roughly speaking, legal obligations are those obligations that are the moral consequence, or *impact*, of the actions of legal institutions. Thus, for example, a statute’s contribution to the content of the law is the impact of its enactment on our obligations in light of fairness, democracy, and other values. Suppose that the moral impact theory is true. And suppose that one formed beliefs about legal facts simply by reading statutes and other legal texts and inferring from the semantic content of the texts of the statute to the legal facts simply by taking any proposition in the semantic content to constitute a legal norm. Such inferences would have the wrong structure because they do not take into account some of the determining facts, namely the moral facts about the impact of the enactment of statutes on our obligations. The example does not illustrate a means of reliable

access to the target facts, however, because such inferences would not be a reliable way of ascertaining the legal facts.

It is perhaps worth noting that nothing in my argument depends on the reader's accepting that the examples I have offered are in fact successful examples of reliable non-tracking methods of ascertaining the relevant target facts. For example, psychology has raised significant doubts about the reliability of our first-person access to at least some of our mental states. To the extent that our first-person access to the mental is not in fact reliable, it does not offer an alternative to tracking in that domain. My purpose in offering the examples is to illustrate potential alternatives to tracking in some domains. In the next section I will argue that such alternatives are not plausible in the legal case.

3. The legal domain

a. Preliminaries

I take the legal domain to comprise facts about the content of the law – *legal facts*, for short – such as the fact that it is not permissible to compel a person to give self-incriminating testimony.⁸ Legal facts are not among the most basic facts of the universe. Rather, they obtain in virtue of other, more basic facts. A core project in philosophy of law is that of giving what I will call a *theory of law* – an account of how the content of the law is determined by more basic facts. (Thus, theories of law, as I use the term, are theories of the metaphysics of legal facts.)

⁸ As I use the term, legal facts are general – e.g., that one is liable for negligently caused harm – not particular, e.g., that John Doe owes \$30 to Jane Smith. *Legal fact* is thus roughly equivalent to *legal norm*. As noted above, legal facts need to be indexed to a jurisdiction and a time, but I generally omit such qualifications.

In this paper, I avoid taking controversial positions on the metaphysics of legal facts, though some of the arguments depend on assumptions about the metaphysics that I hope are widely shared. I take it that an account of legal facts can be given in terms of facts, such as semantic, mental, and normative facts and facts about human behavior, many or all of which are not among the most basic facts of the universe.

But I think it might be helpful just to have in the background a few examples of competing views about the metaphysics of law just so that people have some sense of the territory. I will sketch three. First, as mentioned earlier, on Hart's legal positivist theory, the way in which legal materials contribute to the content of the law is determined by convergent practices of judges and other officials in the very simple way that I described.

By contrast, according to Ronald Dworkin's influential account, the content of the law is the set of principles that best *justify* the past practices.

And, finally, on my own view – the *moral impact theory* of law – roughly speaking, legal obligations are those obligations that are the moral consequence, or *impact*, of the actions of legal institutions. Thus, for example, a statute's contribution to the content of the law is the impact of its enactment on our obligations in light of fairness, democracy, and other values.

One more bit of terminology: I use the term *legal interpretation* for the process of ascertaining legal facts. (I argue elsewhere for the substantive claim that legal interpretation is best understood in this way, but for present purposes, I will just stipulate this usage.)

The central thesis of this section is that our only reliable method of finding out about legal facts is through the metaphysically more basic facts that determine the legal facts. More carefully, as stated above, the reliability of a method depends on the extent to which it tracks the metaphysics. To the extent that a method diverges from tracking, it can yield only incomplete knowledge, e.g. knowledge of a limited body of legal facts, such those of a local jurisdiction.⁹

Before specifically investigating the possibility of non-tracking methods of ascertaining legal facts, I want to make a few preliminary remarks about why the legal domain seems unpromising territory in which to seek an alternative to tracking. First, far from there being a method of legal interpretation that is widely agreed to be reliable, there is widespread debate over the correct method of legal interpretation. Competing methods of legal interpretation take inconsistent positions about the role of the literal meanings of statutory texts, so-called “public meaning,” communicative intentions of legislatures, intentions of the legislature with respect to how statutory texts are to be applied, moral values, contemporary understandings, historical practices, legislative history, and much more.

⁹ If legal interpretation requires tracking the more basic facts that determine the legal facts, will ascertaining those determining facts require, in turn, tracking still more basic facts that determine them – and, therefore, ultimately tracking the most basic facts of the universe? The answer will depend on the epistemology of the facts that determine the legal facts. For example, if empirical facts that are available through perception are among the more basic facts that determine the legal facts, then ascertaining those facts will not require tracking still more basic facts that determine those empirical facts. Perception gives us a direct route to them. In general, then, the extent to which tracking the more basic facts that determine the legal facts will require tracking still more basic facts will depend on the epistemology of the more basic determining facts.

Second, there is a great deal of disagreement about the legal facts themselves (in part because of the disagreement about methods). And though there are many legal facts on which there is consensus, many of the competing methods of legal interpretation yield all or nearly all of these legal facts (thus rendering problematic an approach to validating a method by showing that the method yields more of the legal facts on which there is consensus than any other method). This situation is unsurprising because the reason for the consensus on certain facts is in significant part that the leading methods of legal interpretation all yield these facts (though no doubt the dependence also goes the other way; a method is not plausible unless it yields the facts on which there is consensus).

More fundamentally, the consensus is likely the result of partial tracking. All of the leading methods of legal interpretation involve inferring legal facts from facts that are strong candidates to be among the determining facts. Moreover, the leading methods agree about the importance of certain basic determinants of legal norms (though not qua basic determinants), such as the meanings of certain texts or utterances. In other words, leading methods of legal interpretation all partially track the determining facts. When the determining facts on which the leading methods converge all point in the same direction and nothing else points in a different direction, we get a consensus about the legal facts. If this suggestion is right, then an appeal to legal facts on which there is consensus to validate a method would not avoid dependence on tracking.

In sum, we lack an independent way of validating a candidate method of ascertaining legal facts. Without some way of checking whether the deliverances of a method are correct, we can't use induction to move from the fact that a method has yielded the right answers so far to the conclusion that its answers to new questions are correct. In other areas, we can hypothesize

that a particular kind of inductive inference is reliable. And we can then see whether it bumps up against reality – whether it yields false predictions. But in the legal case, it's hard to see how we can ever get such a reality check.

This situation raises a general worry: Even if a certain non-tracking method happens to be reliable, there would be a serious question about justification. How it could be justified to rely on the method in the face of pervasive disagreement about which method is correct and no independent way of validating the method? The situation is not like one in which a creature has a reliable method that it has no choice but to rely on (despite having no access to the reliability of the method). Here there is a choice to rely on one method rather than competing methods without a basis for believing the method more reliable. A related point is that the truth of beliefs formed through the method seems merely lucky, despite the reliability of the method, because the *choice* of this method over others is merely lucky.

These considerations about justification push toward the claim not merely that a method must track the determining facts, but that it must track them *as such*. For the argument applies to any method, regardless of how reliable, unless the method can be independently validated. And the argument purports to show that, in the circumstances, the only way of validating a method would be to show that it tracks the determining facts. I will for the most part set aside such issues in this paper, focusing instead on issues concerning reliability.

b. A prima facie case that tracking is necessary in the legal domain

I turn now to specific exploration of possible alternatives to tracking. The possibility of direct – non-inferential – knowledge of legal facts seems wild. It seems obvious that the primary route to legal facts (in the sense of “primary” explicated above) is via inference from, among other things, facts about what various texts say or mean and what various people have done or decided. In the case of legal facts, we have nothing analogous to perception of the external world, first-person access to the mental, or a priori knowledge of fundamental moral principles.

[It might be pointed out, however, that if we can reach a conclusion by explicit inferential reasoning, presumably the reasoning can become automated. Chessmasters can look at a chessboard and reach a conclusion without being aware of the premises or of the inferential route. In order to make my thesis about the legal domain stronger than a claim about what abilities humans now happen to have, I could explicate the notion of tracking in a way that allows the inferences from determining facts to target facts to be tacit. Now this might be thought to raise a worry about my way of elucidating tracking by contrasting it with non-inferential methods such as perception. A quick answer is that though perception probably involves tacit inferences – think for example of the processes involved in depth perception – it does not involve tacit inference *from the more basic determining facts*. Perception does not begin from premises about micro-physical facts.]

I return later to potential objections that purport to vindicate the possibility of direct legal knowledge – i.e., legal knowledge that is not based on any kind of inference from the determining facts.

Lateral knowledge of legal facts also seems problematic. There would have to be a domain in which the facts – call *them* the X facts – are correlated with legal facts in a lawlike way – but not because the X facts are determinants of legal facts. One possibility is that the relationship is the other way around: the X facts are *downstream* of the legal facts – most likely, causally downstream. A second possibility is a domain that runs in parallel to the legal domain, neither determined by nor determining of the legal. The most plausible candidate would be the moral domain or some other normative domain, perhaps the domain of practical reasoning. I address the most promising versions of these lateral possibilities immediately below when I respond to objections.¹⁰

The possibility of inferring legal facts from their *effects* faces a special obstacle. The phenomena that are studied by physics and the special sciences are embedded in causal networks. Consequently, there are indefinitely many ways of finding out about these phenomena. By contrast, legal facts seem to have their principal causal impact through only one pathway – people’s *beliefs* about those legal facts. For example, the fact that the law prohibits insider trading affects behavior mainly by way of people’s beliefs that the law does so. (Suppose there is a jurisdiction in which the only source of law relevant to insider trading is an utterly unknown statutory provision that, properly understood, has the effect that insider trading is criminal. In this jurisdiction, insider trading is, by hypothesis, prohibited, but that prohibition will have no

¹⁰ Note that we encounter again the problem of how to validate any such putative method. As things stand, there isn’t an uncontroversially reliable method of inferring legal facts from facts of another domain. We don’t find ourselves endowed with something like our almost automatic ability to infer others’ mental states from their behavior. Any proposed method of lateral inference would have to be validated by showing that it in fact accurately identified legal facts (beyond the area of consensus in which many methods give the right results). But, again, without a reliable method of ascertaining legal facts, it’s hard to see how such validation could be accomplished. So, for example, suppose it were proposed that we could infer legal facts from moral facts. How could it be justified to rely on such inferences without showing, e.g., that moral facts and legal facts are determined in a parallel way by certain more fundamental facts, and that this common determination structure makes available a reliable inferential path from moral facts to legal facts?

effect on behavior as long as no one is aware of it.) Similarly, if there were a widespread misunderstanding about the legal norms governing some issue, we would expect to observe the effects that correspond to the mistaken view, not the ones that would correspond to the actual legal norm.) Now it is an interesting question why the effects of legal norms are mediated primarily by our beliefs. It is not in general true that the effects of nonbasic facts are mediated primarily by our beliefs. Perhaps the answer is that normative facts are special in this respect – they have their causal impact via humans’ beliefs about them.¹¹

I won’t try to pursue this issue here. But the bottom line is that because the effects of legal facts are mediated primarily by our beliefs, the effects of legal facts do not provide us with a primary way of ascertaining legal facts. The beliefs about legal facts that mediate the effects of the legal facts must themselves be formed in a way other than by inference from the effects of the legal facts.

Inferences with the wrong structure – in particular, inferences based on only some of the determining facts – seem a more promising avenue to legal knowledge than the direct or lateral possibilities. But it is consistent with my thesis that a method that partially tracks the

¹¹ Now one might question this suggestion. It might be doubted whether legal facts are genuinely normative. And it might be pointed out that the facts that determine the legal facts include facts, for example about human decisions, that have causal effects. Now on my own theory of law, the moral impact theory, the legal facts depend on moral facts so the legal facts are constitutively normative. And on Hart’s view, legal facts can depend on moral facts, depending on the role that judges give to moral facts.

But what about on views on which legal facts depend only on social practices? Here’s a very quick sketch of what I would want to say about this. The kinds that figure in high-level sciences classify things together for reasons of causal explanatory power. Think of Putnam’s example of the square peg and the round hole. But legal kinds classify things in a way that is utterly gerrymandered from the point of view of causal explanation. So even if the basic determining facts are causally potent, legal norms will not have systematic effects from which the norms can be inferred – other than via beliefs about those norms.

metaphysics makes it possible to acquire legal knowledge that is limited in scope. The worry is that, given that metaphysical determinants are nonredundant – if they did not do metaphysical work, they would not be metaphysical determinants – inferences based on only a proper subset of the determining facts will not provide a generally reliable method.¹² In the discussion of objections below, I consider more concretely methods that rely on inferences from a proper subset of the determining facts.

c. Responses to objections

I turn now to a few objections to my thesis that the primary method must involve tracking. First, it has sometimes been suggested to me that a natural law or non-positivist view of law might make it possible to ascertain legal facts without tracking. The idea seems to be that, to the extent that the law is constrained or determined by morality (or divine command), we can know the content of the law without consulting legal texts or practices. The details will differ depending on the precise nature of the non-positivist view in question. (For brevity, I use *non-positivist* to encompass both *natural law* and *non-positivist* positions.) If, for example, the view is that anything sufficiently morally wrong cannot be a valid legal norm, then we can know immediately that the law does not require fugitive slaves to be returned -- indeed, that humans cannot be property. Far from *avoiding* tracking, however, such a method of ascertaining legal facts relies squarely on tracking. According to the non-positivist view in question, facts about what is morally wrong are more basic than legal facts and play a specific role in determining the

¹² Perhaps this would be the place to argue that plausible assumptions about the metaphysics of law make it difficult to see how induction from a proper subset of the determining facts could be generally reliable – that is, how some of the determining facts could be epistemically, though not metaphysically, redundant.

content of the law. Thus, the proposed method of ascertaining legal facts straightforwardly appeals to more basic facts and their role in determining the legal facts to infer the appropriate conclusion about the legal facts. Parallel points apply to analogous proposals based on other non-positivist views. The key point is that non-positivist views make it possible to infer legal facts from moral facts precisely because the moral facts are, according to the views in question, determinants of legal facts.

A second objection is that, without any reliance on non-positivist assumptions, we can know that the law will, in many ways, be what it would be wise or sensible for it to be. We know, for example, that murder, assault, robbery, and kidnapping are legally prohibited. If we go to a new jurisdiction, we can take such legal facts for granted without consulting any texts or practices. Perhaps similar points can be made for legal facts outside of the *mala in se* prohibitions of the criminal law, such as nearly ubiquitous features of property, contract, and tort law.

As the suggested method is not supposed to rely on non-positivist assumptions, the line of thought has to be that, because legislators, judges, and other legal officials are reasonable and share widespread human goals, they will take actions and make decisions that will yield the expected legal norms. This suggestion does not bear close examination. To begin with, the method would not be reliable, as predictions about what legal officials will do based solely on generalizations about human psychology are shaky. The content of the law varies widely across jurisdictions and is often far from what would be wise. More importantly, the proposed method tacitly relies on tracking, for it moves from empirical assumptions about the actions and decisions of legal officials – important determinants of legal facts – to conclusions about legal facts. For example, from the plausible assumption that legislators have enacted statutes

specifying criminal punishments for murder (or that other legal officials have taken actions that with similar legal impact), the method concludes that murder is criminally prohibited.

A third objection is that a method of inferring the legal facts with the wrong structure – in particular a method that infers the legal facts from only a proper subset of the determining facts – could be locally reliable. In section 2, I gave the example of a method of inferring legal facts by simply reading the authoritative legal texts and inferring the communicative intentions of the authors – without attending to other factors such as moral facts and the practices of judges.¹³ If Hart's theory of the metaphysics of law were correct, this method could be locally reliable (despite not having the right structure) if judges in the local jurisdiction happen to treat authoritative legal texts as contributing to the law in accordance with the communicative intentions of the authors. Or suppose that the moral impact theory of law is true. In that case, the method could be locally reliable if, because of contingent facts about the local situation, the moral impact of enactments turns out to be whatever the authors intended to communicate.

Intuitively, it is questionable that such a locally reliable method gives rise to knowledge. We can articulate the intuitive concerns in terms of the absence of luck and in terms of safety. It seems just lucky that the method works, as it would not work in other jurisdictions or perhaps even in this jurisdiction at other times. Indeed, if things had been somewhat different, it would not have worked in this jurisdiction now. Similarly, the method is plausibly unsafe, as it seems that beliefs based on the method could easily have been false. For example, in the illustration based on the supposition that Hart's theory is true, the beliefs about legal norms would be false in a situation in which judges treated statutes differently.

¹³ For purposes of this example, I set aside a host of familiar problems about the nature and existence of communicative intentions of legislatures and other lawmaking bodies.

Furthermore, even if the method gives rise to knowledge in the local jurisdiction at the present time, it is not a reliable method of obtaining legal knowledge in general. Recall that my thesis allows that a method that partially tracks the determining facts can be a way of acquiring knowledge that is limited in scope. Thus, it is consistent with my thesis that a method that infers the legal facts from only a proper subset of the determining facts could be locally reliable.

A fourth objection is that a theorist could use a method analogous to reflective equilibrium in ethics to ascertain the content of the law.¹⁴ A theorist could begin with, on the one hand, firm convictions about the way in which some concrete interpretive questions must come out and, on the other, intuitive convictions about the content of the law at a more abstract level. As an example of the former, one might take it as a fixed point that the equal protection clause of the 14th Amendment of the U.S. Constitution must have the consequence that segregated public schools are unconstitutional. With respect to the latter, a theorist might take it to be intuitively clear that the U.S. law has a principle that no one may be a judge in his or her own cause, a principle that no one is above the law, or a principle of separation of powers. The suggestion is that, from such starting points, a theorist could make progress in figuring out the content of the law on more controversial issues, for example by using convictions about interpretive outcomes to test and modify convictions about more abstract principles and, at the same time, adjusting convictions about interpretive outcomes in the light of intuitions about abstract principles.

A different, but related, idea is that a holistic method analogous to reflective equilibrium could be used *not* to work out first-order legal facts directly but to make progress on figuring out the correct *method* of legal interpretation. In addition to firm convictions about the outcome of

¹⁴ I am grateful to Tyler Burge, Jeffrey Helmreich, and Seana Shiffrin for discussion of this point.

some interpretive questions, a theorist would appeal to intuitively attractive theoretical principles *about legal interpretation*. Some possible examples: a good theory of legal interpretation should avoid absurd results; should give importance to the linguistic meaning of authoritative legal texts; should have the consequence that the content of the law is accessible; should respect democratic values. The suggestion is that, from such starting points, a theorist could make progress on figuring out the correct method of legal interpretation through a holistic process.

I am dubious about both versions of the reflective equilibrium suggestion. The problem is the source of the convictions that the method draws upon. How can they be reliable? And what is the warrant for these convictions? One possibility is that the convictions are based on testimony – all one's life, one has been told or has read that certain norms are part of U.S. law, so one firmly believes it. To the extent that it is based on testimony, the reflective equilibrium method cannot bear much weight. Testimony about the content of the law, however often repeated, is only as good as the primary method of ascertaining the legal facts on which it is based.

The second possibility is that the convictions are inferred from partial and perhaps tacit views about the basic determining facts (though not necessarily under that description) -- in other words, that the convictions are tacitly relying on tracking. For example, the conviction that the 14th Amendment must have the consequence that segregated public schools are unconstitutional might be based on a view that American constitutional norms are based on, or limited by, fundamental truths about justice. The important point for present purposes is that, to the extent that the convictions derive from inferences from determining facts, then the proposed method is not really an alternative to tracking.

Could the convictions be reliable in virtue of the role of examples in determining the meaning of relevant terms, e.g., “legal” or “law”? Actual examples plausibly play an important role in the meta-semantics of terms and concepts for physical kinds and substances. For example, putative examples provided by ostension may help to determine the meaning of “water,” “paper,” and “tree”. Because of this role, such examples, though defeasible, may play an important role in a holistic confirmation process, providing a source of convictions about particular cases that is not based on tracking.

There does not seem to be anything comparable in the case of legal facts, however. It is crucial to distinguish statutes, constitutions, and other such legal texts or materials from legal norms. Our question concerns not our access to legal texts, but our access to legal norms. Perhaps examples play a role in the meta-semantics of terms for legal texts like “statute” or “law.” But it is difficult to see what could play a comparable role with respect to legal norms. Simply being told that there is a particular legal norm, say, against self-incrimination gives one no way of recognizing whether that legal norm obtains in another jurisdiction or at another time. And watching someone, for example, read an ordinance and conclude that a given legal norm obtains does not provide a factor that is independent of tracking.

It might be suggested that convictions about interpretive outcomes are the consequence of situations in which all credible methods of legal interpretation yield the same answers. For example, one’s confidence that a statutory provision has the effect that two witnesses are required for a will to be valid might derive from the fact that all remotely plausible methods of legal interpretation yield that result. Cases in which all credible methods of legal interpretation yield the same answers are likely cases in which all plausible candidates for the more basic facts

that determine the content of the law point in the same direction (or have nothing to say). (If some factors pointed in another direction, the issues would not be so free from controversy.)

Consider an issue on which all of the plausibly relevant factors – the literal meaning of the relevant statute, what the legislature intended to communicate, what legal norm the legislature intended to enact, what the audience would reasonably have understood the text to mean, what the audience would reasonably have understood the legislature to have meant, the best justification of the enactment of the statute, the consequences of the enactment of the statute for our moral obligations, and so on – all converge. On such issues, we are inferring legal facts from the best candidates for more basic, determining facts – it’s just that we don’t need to resolve any controversial metaphysical questions in order to make the inferences: platitudes about legal interpretation or about how legal facts are determined are sufficient in the circumstances. Thus, the beliefs in question are based on tracking. At any rate, the beliefs in question will not help us to make progress on any controversial issues of legal interpretation. By hypothesis all credible methods of legal interpretation yield these same outcomes. Consequently, they will not provide us with any help in choosing between different methods of legal interpretation.

Thus, though I am generally sympathetic to the methodological point that we have more confidence in certain first-order beliefs than in any philosophical arguments or theories, I am skeptical that this point takes us very far in the legal case. A method that takes certain first-order beliefs as a starting point is reliable only to the extent that these beliefs are reliable. In the legal domain, plausible explanations of the reliability of the relevant first-order beliefs appeal to tracking.

4. Conclusion

[To be written. The section will pull together some of the distinctive features of the legal domain that make the tracking thesis plausible in that domain: that legal facts have their causal impact only or principally via people's beliefs about those facts; the absence of any plausible analogue of perception or of anything like a priori knowledge of moral or mathematical facts; pervasive disagreement over both the first-order facts and over methods of ascertaining them; that examples do not play an important role in the meta-semantics; perhaps certain plausible assumptions about the metaphysics of legal facts.

This section will also suggest other domains where it might be worth exploring the necessity of tracking: normative domain such as etiquette, religion, and games; domains in which the facts are grounded in conventions.]