The Value and Normative Role of Knowledge

JULIEN DUTANT

Over a decade as my supervisor and mentor, Pascal has skilfully managed to keep me under the impression that I carried my research freely while quietly guiding my attention towards deeper and unfamiliar issues. He thus introduced me to knowledge-first epistemology, the debate over pragmatic encroachment and the role of norms, values and reasons in epistemology before I quite grasped their significance. Each time the new thoughts would slowly make their way into my own and I would eventually find myself intensely preoccupied with the issues that were central to Pascal's seminars a few years back. The present paper is another instance of this phenomenon. I dedicate it to Pascal, with all my respect and gratitude, and with apologies for my *esprit d'escalier*.

In his (2009), Pascal connects two topics that epistemologists have mostly kept apart. One is the debate over pragmatic encroachment, namely, the idea that whether one knows (or believes, or justifiably believes) partly depends on practical factors such as one's interests and the stakes one faces. Central to this debate is the claim that *knowledge* is the norm of action: that is, that in some salient sense of "ought", one ought to act in view of what one knows. The other issue is the question of the *value* of *knowledge*, namely, whether and why knowledge is a good thing, and in particular, a better thing than mere true belief. Most of Pascal's paper deals with pragmatic encroachment. It was what attracted all my attention then. But it also raised a second question that epistemologists seldom discuss: could the idea that knowledge is the norm of action explain the value of knowledge? Pascal's answer was negative. At

the time I could hardly get my mind around the question. It is not *prima facie* clear how one could even try to derive one claim from the other. One may sketch some paths; for instance, if you should act only on what you know, then when you know you have "more" to act on than when you merely have a true belief. But these hardly constitute a suitable basis for discussion. Pascal's negative answer supported the widespread attitude of keeping the two ideas apart. I went along and forgot all about it. A few years later, I feel I have finally reached a perspective from which I can take up Pascal's question. My views on the matter are far from settled, so this is more of a progress report. The option that I currently find the most appealing differs from Pascal's. Like him, it denies that we can explain the (alleged) value of knowledge by its normative role, but unlike him, it does take the normative role of knowledge to shed light on its value, by showing why it need not have value at all. Before I get to this, however, I will lay out the perspective from which I take up Pascal's question.

1. Why knowledge matters

Epistemology in the second half of the XXth century enjoyed a spectacular revival. But when it came to knowledge it focused almost exclusively on two questions: what it is and whether we have it. It also asked when a belief was "justified", which, to some at least, was the same as asking what one ought to believe. But the latter question was mostly treated as independent of, and prior to, questions about knowledge. For a variety of reasons, by the 1990s epistemologists increasingly wondered whether and why knowledge mattered. It is well to ask what knowledge is and whether we have it, but, they began to ask, why should we care?

The question is pressing if we distinguish knowledge from justified true belief — as is common in post-Gettier times. Suppose I come to the conclusion that I do not know whether there are lions. That may be an unsettling conclusion to reach. But why, exactly? I may start wondering whether there really are lions. And I may be unsettled at the thought that there are no lions. But all that suggests is that whether there are lions matter, not whether I know that they are. Similarly, I may be unsettled at the thought that I may have been mistaken on that matter and others. But all that suggests is that whether my beliefs are true matters. Another reaction I may have is to judge that I ought not to believe that there are lions. But all that suggests is that whether my belief that there are lions is justified — whether I ought to believe — matters. If whether I ought to believe it does not whether I know it, then again that does not entail

that knowledge itself matters. Putting it all together, it may matter whether p, whether I have a true belief that p, whether I am justified in believing that p; but if I can have a justified true belief that p without knowing p, it is unclear whether and why knowing itself matters.¹

A somewhat shallow answer defers to common sense. We think about knowledge a lot. The verb know is currently one of the ten most used verbs in the Oxford English Corpus. It is the second propositional verb (after say) and the most common verb describing a mental state (just before see, think, look and want). It is much more used than believe, true, justified and even more used than ought, should and must.² Since we talk about knowledge a lot, we think about it a lot. Moreover, we take ourselves to know many things and we want to know many things.³ So philosophers can rest assured that knowledge at least matters to us.4 The answer is somewhat shallow, however. First, even on the assumption that knowledge is something we desire, we may still wonder whether and why it is desirable.⁵ Second, similar remarks can be made about other common notions. We take ourselves to do, make and get many things, and we want to do, make and get many things. Yet few philosophers would say that doing, making or getting matter. That is so, I venture, because philosophers take these notions to be too crude to describe the underlying phenomena. There is little useful theory to be made about the making that is common to making a plan, making a present and making a soup. Philosophers found it more useful to theorize about the underlying phenomena in terms of intention, action, causation, ownership and so on. One may worry that knowledge is also too crude a notion for picking up something that matters and that is worth theorizing about. The worry is made more acute by the existence of epistemological traditions that do without the notion altogether, adopting instead notions such as justification, evidence, confirmation and probabilistic notions.⁶

¹This line is forcefully pushed by Mark Kaplan (1985).

²"The OEC: Facts about language", Oxford University Press, http://www.oxforddictionaries.com/words/the-oec-facts-about-the-language, retrieved Jan 4th2014.

³As Aristotle famously notes in *Metaphysics* A, 1.

⁴See Williamson (2000, 31): "For knowing matters; the difference between knowing and not knowing is very important to us. Even unsophisticated curiosity is the desire to *know.*"

⁵Even subjectivists who think that things are at bottom made valuable by our valuing them are not committed to the view that everything we desire is desirable.

⁶Peirce was an early defender of the view that the notion of knowledge is disreputable: "there will remain over no relic of the good old tenth-century infallibilism, except that of the infallible scientists, under which head I include [...] all those respectable and cultivated persons who, having acquired their notions of science from reading, and not from research, have the idea that

Two more substantial answers have been prominent in recent literature.⁷ The first is that knowledge is *good*, or, as philosophers prefer to say, that it *has value*. Good things obviously matter; so if knowledge is good, it matters. The idea famously figures in Plato's *Meno*, where Socrates approvingly reports that "knowledge is prized higher than correct opinion".⁸ The answer only goes so far. We may still wonder why knowledge is good, and in particular why it is better than justified true belief. This has been the subject of much discussion in the last decade. But as long as we grant that knowledge is good, as the recent value of knowledge overwhelmingly does, we already have an answer to why it matters. Call it the *value* answer.

The second answer comes from the idea that knowledge is something we *ought* to have in order to do certain things or have certain attitudes. I will focus on two such claims. The first is that *one ought to act only on the basis of what one knows*. Call it the Knowledge–Action Principle. The second is that *one ought to believe only what one knows*. Call it the Knowledge–Belief Principle. There are alternative or additional claims in the vicinity: one ought to do what is best in view of what one knows; one ought to believe only on the basis of what one knows; one ought to be certain only of what one knows; and so on. It does not matter for our purposes which ones we choose; the two selected above will serve as concrete illustrations. Now what one ought to believe and what one ought to do obviously matters. So if what one knows partly determines what one ought to believe and to do, then knowledge matters. So principle like the two above also offer an answer to the question why knowledge matters. Call it the *normative role* answer.

What the normative role answer amounts to is not entirely clear because *ought* may mean many things. For now we just flag the issue; it will take a central importance later on.

The normative principles have an ancient pedigree as well. Zeno (of Citium, the founder of the Stoa) claimed that the wise only assent to what they have a "grasping impression" of — by which he essentially meant, what they know.⁹ Since he clearly thought that one ought to do what the wise does,

[&]quot;science" means knowledge, while the truth is, it is a misnomer applied to the pursuit of those who are devoured by a desire to find things out." (Peirce, 1950, 3) (It does not seem to occur to Peirce that *finding something out* may come down to *coming to know*.)

⁷They are not the only ones. Another one (inspired by Craig, 1990) is that knowledge is property of people that is useful for us to spot while inquiring: we figure out who knows what to decide who to use as source of information.

⁸98a, trad. G.M.A. Grube.

⁹See e.g. Cicero, Academica 2.77–8, quoted in Long and Sedley (1987) as 40D. Stoics would

that amounts to an endorsement of the second principle. Academic sceptics agreed; but since on their view neither we nor the wise knew anything, they argued that one ought not to assent to anything. As a result they were under pressure to deny the first. Confronted with the Stoic objection that the wise would have to act, Carneades developed the idea that one could act on the basis of merely *convincing* impressions. The fact that Stoics saw that as an objection shows that they endorsed the first principle as well. Just as the claim that knowledge is better than mere true belief is often called *Meno's thesis*, we may call the two principles *Zeno's norms*.

The two answers are mutually compatible: it may be that knowledge matters both because it is a good thing *and* because it plays a certain normative role. But it is tempting to see whether one could be used to derive the other. A *Value to Norm* derivation would derive Zeno's norms from Meno's thesis and plausible. A *Norm to Value* derivation would derive Meno's thesis from Zeno's norms. In both cases we would allow the use of plausible background assumptions. Since there is little about norm or value that is uncontroversial, we may also generously allow the use of controversial claims about norms or values in general. As I understand "derivation" no order of priority is required: it may be that both Meno's thesis follows from Zeno's norms and the other way round. ¹⁰

This paper discusses both derivations. Section 2. discusses the Value to Norm route. I am not optimist for it. Some reasons for pessimism come from Firth (1998a) and Berker (2013a, 2013b). They argue that "consequentialist" or "teleological" ways of deriving epistemic norms from epistemic values fail because they result in norms allowing for trade-offs that the correct epistemic norms for belief forbid. I do not find the objection decisive, however. It leaves some "teleological" derivations standing, as well as non-teleological ones. A more serious problem seems to me to be the impossibility of deriving anything

say that what we would now describe as paradigm situations of perceptual knowledge (seeing that an apple is on a table) involve "having a grasping impression". But they would not call it "knowledge" (episteme) yet, for they thought that knowledge required resistance to dialectical cross-examination. Once we set aside this inflated view of knowledge — or once we read Stoics' notion of episteme as denoting something like science or scientific understanding —, we can take their theory of "grasping impressions" as a theory of knowledge. (Commentators sometimes do so without further ado, e.g. Long and Sedley 1987 and Frede 1983/1987; see Frede 1999 and Hankinson 2003 for more guarded statements.)

¹⁰Some characterize "consequential theories" as those that explain the right in terms of the good and "deontological theories" as those that explain the good in terms of the right. Because these terms have many associations and because they are meant to be exclusive of each other, I think it would be misleading to use them to label the Value to Norm and Norm to Value derivations, respectively.

like the Knowledge–Action Principle. In a nutshell, the problem is that Meno's Thesis cannot ground a difference in value between a case where one has a good and a bad belief, but acts on the good one, and a case where one has a good and a bad belief, but acts on the bad one. Thus the problem arises because the Knowledge–Action Principle is not simply a norm about belief but about the coordination of belief and actions.

Section 3. discusses the Norm to Value route. It is *prima facie* more promising. Saying so goes against a strong trend in current epistemology: while there has been much debate over why knowledge is good, little of it has explored the idea that it is good because of its normative role. Epistemologists appear to have assumed that something like the Value to Norm derivation will in turn explain the normative role of knowledge, and that it would therefore be illicit to appeal to the normative role of knowledge to explain its value. That being said, I have doubts about it as well. As far as I can tell, the derivation would have to rely on the idea that knowledge is good because it is required to be allowed to believe or act on one's belief. But in general it is not the case that necessary conditions for being allowed to do something are good — not even that necessary conditions for being allowed to do something good are themselves good. The problem does not show that the derivation fails, but it indicates that more needs to be say.

Section 4. turns the apparent failure of the Norms to Value derivation into a virtue. For once we assume that knowledge plays a central normative role, it becomes unclear what is left of the motivation for the idea that it is a good thing. For instance, the fact that it plays a central normative role is sufficient to explain why knowledge matters. There is no need to make the additional claim that it is a good thing. So we may try to use Zeno's norms to explain away Meno's Thesis. I have put the suggestion forward elsewhere (Dutant 2012, forthcoming). Here I want to discuss two problems for it. The first is that the proposal has a hard time explaining why knowledge is something worth aiming at, for Zeno's norms themselves do not prescribe acquiring knowledge. In reply I argue that such prescriptions follow from Zeno's norms in conjunction with other aims and other norms of action. Another is that the proposal requires a strong primitive, namely a layer of normativity distinct from the usual "objective" and "subjective" ought that are commonly accepted. I will put forward a few considerations in its favour.

2. From Value to Norms

Let us assume Meno's Thesis and examine whether we can derive Zeno's norms. Meno's Thesis, expressed as the slogan "knowledge is better than mere true belief", is somewhat unspecified. It is unclear whether it is a generic or universal claim and what exactly the bearers of value are supposed to be. For the sake of concreteness we will use on a more precise claim. The claim ascribes values to states of affairs. It states that knowledge is *pro tanto* good and that belief without knowledge is *pro tanto* bad:

(MT) For every S, t, p, the state of affairs of S knowing p at t is ($pro\ tanto$) good, and the state of affairs of S believing p without knowing p at t is ($pro\ tanto$) bad.

It follows from (MT) that knowing p is better than having a true belief in p that does not constitute knowledge. For there is a disvalue in the latter that is absent in the former, namely believing without knowing. It also follows from (MT) that knowledge-constituting belief is better than lack of belief, and that lack of belief is better than belief that does not constitute knowledge. It does not follow from (MT) that knowing p is always *overall* good; the value that it has in virtue of being knowledge can be offset by other considerations. Similarly, it does not follow from (MT) that believing without knowledge is always *overall* bad; its disvalue may be offset by other considerations. (MT) is neutral on whether the *pro tanto* value of knowing p is the same for every p. Perhaps some things are more valuable to know than others.

(MT) is stronger than the claim that *some* or *most* state of affairs of knowing are *pro tanto* good. It is also stronger than its first conjunct alone. If the derivation fails with that strong assumption, it will fail with weaker ones. We may give the derivation its best chance.

Norms are about what we ought to do; values about what is good or bad. How do we derive one from the other? A common paradigm is *consequentialist*: roughly, one ought to do what has or tends to have the best consequences. As Berker (2013a, 351–7) notes, much contemporary epistemology adopts such a framework.¹¹ It is assumed that we have certain epistemic aims

¹¹In Berker (2013a, 342) prefers the term "epistemic teleology", because he thinks that "epistemic consequentialism" will evoke the view that what one ought to, epistemically speaking, is what promotes *practical* (non-epistemic) goods. Firth (1998a) uses the term "epistemic utilitarianism"; others "epistemic instrumentalism" (Kelly, 2003) . "Epistemic consequentialism" is used by Percival (2002), Stalnaker (2002) and Berker (2013b), among others. There are some differences in how these authors characterize the view so labelled. For instance Kelly takes it to include the

— such as having true belief and no false beliefs — and that what we ought to do, epistemically speaking, is what promotes those aims. Berker (2013a, 2013b), building on a problem due to Firth (1998a, 1998b), argues that any such derivation of epistemic norms will fail. That is, any such derivations will misclassify some justified beliefs as unjustified and conversely. While I share the view that consequentialism is unsuited to derive norms of belief and I agree that Firth and Berker's problem shows that many versions of epistemic consequentialism fail, I do not think they rule out all such versions. Be that as it may, Berker's and Firth's problem leave untouched *non-consequentialist* ways of deriving norms from value. So for our purposes, the discussion of epistemic consequentialism is mostly a side-show. Since, however, the paradigm is the most familiar one, it is worth going through it.

Berker (2013a, 344–7) characterizes consequentialist normative theories as having three components. First, a *theory of final value*, which states what things have value in themselves. Second, a *theory of overall* value, which ascribes value to things according to whether and how they promote finally valuable things. Third, a *deontic* theory, which states what one ought to do in terms of overall value. For our purposes we call a belief one ought to have (or is allowed to have) a *justified* belief and a belief one ought not to have a *unjustified* belief. In our attempted derivation, the theory of final value is given by (MT). ¹² To illustrate a complete theory:

```
Theory of final value. For every S, t, p: S's knowing p at t is (pro tanto) finally good, S's believing p without knowing p at t is (pro tanto) bad.
```

We call "final *epistemic* value" the value that things have in virtue of these clauses. We assume that there is some way of summing final values so that the total final epistemic value of a compound state of affairs is the sum of the final epistemic value of its components.

Theory of overall epistemic value (for state of affairs). A state of affairs is epistemically better than another iff the total final epistemic value it brings about (or would bring about if it

idea that epistemic norms and values are contingent upon one's having certain epistemic goals — an idea he objects too. But all share the core idea that what one ought to do, epistemically speaking, is a matter of what promotes the epistemically best consequences.

¹²Rather, we treat it as final. The theory leaves open the possibility that the value of knowledge is ultimately reduced to something else, *e.g.* the value of true belief or practical value.

obtained) is greater than the total final epistemic value the other brings about (or would bring about if it obtained).

We leave open what exactly counts as *brought about* by a state of affairs: all effects, including long-term ones; proximate effects; constitution; constituents (see Berker, 2013a, 347 for some discussion).

Deontic theory (for beliefs). For every S, p, t: S ought to believe p at t, epistemically speaking, iff S's believing p at t is overall epistemically better than S's not believing p at t.

The qualification "epistemically speaking" leaves room for one's epistemic duties to be overruled by other duties. The *ought* claim we derive here is an 'objective' one: it roughly says that one ought to have the beliefs that *in fact* have the best epistemic consequences, whether or not one is aware of them. As in consequentialist ethics, we may associate a 'subjective' *ought* to the objective one:

S (subjectively) ought to believe p at t, epistemically speaking, iff S's believing p at t is *expectably* overall epistemically better than S's not believing p at t.

Where something is *expectably* overall better iff its expected overall value (to S at t) is higher. ¹³

The crucial feature of consequentialist views, in Berker's characterization, is to ascribe overall value to what *promotes* final value. As a result, overall value typically allows for trade-offs: something may be overall good despite having bad consequences, provided it has many good consequences as well. Berker take these trade-offs to generate mistaken epistemic norms. He does not propose a general argument that it is so, however; rather, he mainly argues by generalizing from cases.

Berker's prediction seems borne out when we consider a *direct* deontic theory. I call a *direct* deontic theory one that prescribes a belief directly as a function of its overall value. Their form is along the lines of:

Believe p iff the (expected) overall value of doing so is above a certain threshold.

 $^{^{13}}$ The principle assumes that a notion of expected overall value is defined — e.g., a sum of values of possible outcomes weighted by their probability. It leaves open what sort of expectation is relevant, e.g. what degrees of belief the subject has, or what degrees of belief she should have in view of her evidence, and so on.

The theory given above is an illustration. Now take a case of unjustified belief — say, a belief based on reading tea leaves, while one knows very well that tea leaves do not indicate anything. We can alter the case so that the belief has many epistemically good consequences — for pretty much any notion of consequence and any notion of epistemic good. With enough good consequences, the belief will be counted as overall good. We can even pile up the good consequences until any desired threshold of overall value. By the direct deontic theory, the belief will be counted as justified, *contra hypothesi*. So the theory is false.

Firth (1998c) has put forward cases along those lines (see also Berker 2013b, 369). A brilliant set theorist is on the verge of a ground-breaking discovery, but she is suffering from a serious illness and the doctors give us less than two months' time. Against all evidence, she clinches to the conviction that she will live one full year. The belief in fact raises the chances that she survives long enough to complete her work. Her present belief that she will live has good epistemic consequences: it is a means for her to acquire further knowledge. However, it is not a belief she ought to have, epistemically speaking; it is unjustified. So the theory stated above misclassifies it.¹⁴

It is less clear that Berker's prediction holds good when we consider *indirect* deontic theories. Broadly, we may call "indirect" deontic theories those that prescribe beliefs in virtue of a *relation* to something of overall value. But more precisely, prominent indirect theories all prescribe beliefs in virtue of the overall value of the *process*, *disposition* or *rule* they result from. ¹⁵ These deontic theories are along the lines:

Believe *p* on basis *X* iff the (expected) overall value of *X* is above a

¹⁴The example targets the 'objective' ought claim, but we can adapt it to 'subjective' ones. We may suppose, for instance, that the set theorist knows that if she somehow manages to convince herself that she will live *ten more years*, that will keep her alive for the *six months* needed to complete her work and acquire much new knowledge. Hence the theorist may *expect* the belief to have good epistemic consequences; yet it is not a belief she ought to have, epistemically speaking, since everything indicates that she will *not* survive ten years.

¹⁵There is a rough parallel between act- *vs.* rule-utilitarianism and direct *vs.* indirect theories. The theories require a theory of overall value for process types, dispositions or rules. Typically it is characterized in terms of effects of (actual and possible) instances of the process, manifestations of the dispositions or applications of the rule.

Berker's (2013a, 347) characterization of indirect theories is slightly different. On his account an indirect deontic theory comprises (a) a norm that directly prescribes *processes* (rules etc.) on the basis of their overall value, and (b) a norm that prescribes *beliefs* depending on whether they result from allowed processes (rules etc.). My characterization leaves (a) out and replaces "allowed" by "overall good enough" in (b). They are equivalent for present purposes.

certain threshold.

The argument sketched above does not apply to those theories. By the trade-off aspect of overall value, there will be *bases* with some bad consequences — but many good ones — that will have an above-threshold overall value. But why think that we will find unjustified beliefs with such bases? Considering a few concrete cases will help.

If we are liberal about consequences, we will certainly find such cases. Take a case of unjustified belief resulting from a certain process *X*. Modify the case so that uses of *X* regularly but indirectly bring about good epistemic consequences. For instance, whenever one reads tea leaves, one gets in a good mood that greatly increases one's inferential abilities. In the resulting case process *X* has overall good epistemic consequences. But the belief based on it is still unjustified. As Berker (2013b, 374) notes, these cases are avoided by indirect theories that restrict the consequences relevant to overall value to *proximate* ones.

Berker (2013b, 374) puts forward a further type of case. The case targets epistemic consequentialist theories that use true belief as final value. It goes as follows: a man has a single process to evaluate whether a number is prime: namely, when presented with any number, he forms the belief that it is not prime. The process is quite dumb, but it produces a high ratio of true beliefs, given the relative rarity of primes among numbers. Hence it has an overall good value; by the indirect deontic theory, the beliefs it produces are justified. But they are not. I agree, but it is not easy to generalize the example to theories that take knowledge as final value. For the envisaged process does not produce any knowledge: when the man forms the true belief that 8 is not prime, he is merely guessing. So knowledge-based indirect epistemic consequentialism does not have to ascribe the process any overall value. To get a parallel example with knowledge, we need a process or disposition that typically produces knowledge, but on one occasion produces an unjustified belief. It is not clear that there are such cases. Prima facie, if one forms one's belief in a manner that would typically yield knowledge, that one's belief would seem justified. To discuss precise examples would get us into unclear debates about what counts as "the" process by which a belief is formed. For present purposes it suffices to register the worry that Berker's case reveals a problem with the idea that reliability — in a sense relevant to justification — is merely a matter of ratio of true beliefs.

A perhaps more serious problem with knowledge-centred indirect consequentialism is the following. Suppose that a process typically fails to produce knowledge, but sometimes does. An indirect consequentialist account may count the overall value of the process bad, and as a result the belief it produces as unjustified. In particular, those that constitute knowledge would nevertheless be unjustified. That goes against the common idea that knowledge entails justification. However, there may be independent reasons to reject it (Lasonen-Aarnio, 2010). Alternatively, one may as before doubt whether such cases are possible.

So while Firth and Berker's problem rule out direct epistemic consequentialist theories and truth-belief-centred ones, it is less clear that it arises for knowledge-centred indirect theories restricted to proximate effects. Such theories need not even adhere to the "separateness of propositions" (the idea that final epistemic value with respect to one proposition cannot be aggregated with final epistemic value with respect to another proposition) and the "separateness of times" (the idea that overall value at a time is only a matter of promoting final value at that time) that Berker takes to be necessary to avoid certain trade-offs.

Whatever we think of epistemic consequentialism, there are *non-consequentialist* ways of deriving epistemic norms from values. A simple one is :

S ought to believe p at t iff S's believing p at t has (would have) final epistemic value.

In conjunction with (MT) it follows that *S* ought to believe *p* if *S* knows *p*, or if *S* would know *p* were they to form the belief. So the theory gives us the Knowledge–Belief Principle.

But I fail to see how to derive the Knowledge–Action Principle from Meno's Thesis alone. Meno's thesis does imply that acting on knowledge has more value than acting on a belief that is not knowledge. For the first entails having knowledge, which is good, and the second entails having a belief that is not knowledge, which is bad. But consider the following pair of states of affairs:

- (a) one knows that p, has a mere belief that q, and acts on p.
- (b) one knows that p, has a mere belief that q, and acts on q.

 $^{^{16}\}mathrm{I}\ \mathrm{owe}\ \mathrm{the}\ \mathrm{problem}\ \mathrm{to}\ \mathrm{John}\ \mathrm{Hawthorne}.$

Meno's Thesis cannot count one state of affairs as better than the other. Both include a piece of knowledge and a belief that is not knowledge. The only difference between the two is that the action is caused by the piece of knowledge in one and the mere belief in the other. But that difference is not valued by Meno's Thesis. It need not have effects that are valued by Meno's Thesis either. So from Meno's Thesis alone we cannot derive different values to the two states of affairs. Without different values, it is hard to see how we could derive a deontic theory that prescribes the first and forbids the second. Of course we could simply build the Knowledge–Action Principle in our deontic theory; but that would not be deriving norms from values. ¹⁷

In sum, it appears possible to derive one of Zeno's norms from Meno's Thesis: namely, the Knowledge–Belief principle according to which one ought to believe only what is known. That can be done in a straightforward nonconsequentialist way, and perhaps also in a consequentialist manner. But it does not appear possible to derive Zeno's other norm: the Knowledge–Action principle, according to which one ought to act only on what is known. It is not possible to do so because Meno's Thesis only ascribes value to knowledge, not to relations between one's action and knowledge.

3. From Norms to Value

Let us consider the opposite direction instead. We assume that Zeno's norms hold and try to derive Meno's thesis. But before we do this, it is worth spelling out the norms more carefully. We stated them as follows:

Knowledge-Belief Principle One ought to believe only what one knows.

Knowledge–Action Principle One ought to act only on the basis of what one knows.

But *ought* is a notoriously slippery term. It can be used to mean many things, so the claims above should be clarified. I will distinguish two dimensions of variation in what *ought* claims express. ¹⁸ First, they may vary along normative

¹⁷There are further loops one may go through in this argument, but I do not think they alter the conclusion. One may consider adding more assumptions about value. For instance, we may assume that some actions are good. Insofar as these actions are based on beliefs, the total state of affairs of doing those actions based on those beliefs would be better if the beliefs in questions constitute knowledge. Still, pairs like the one above may still be built.

¹⁸As far as semantics is concerned, we may assume that these variations correspond to various contextually-specified semantic values of "ought". The standard contextualist semantics of

sources. Normative sources are usually put under broad headings such as moral, prudential, legal, aesthetic, all-things-considered, and so on. But I think ought claims may reflect much more fine-grained sources, such as what is prudent for a given task, what is prudent relative to health, and so on. An attractive hypothesis is that normative sources correspond to values: each dimension or aspect of value is a source of ought claims. Second, ought claims vary along normative layers. A typical distinction of normative layer is the one commonly made between 'objective' vs. 'subjective' ought. The distinction is orthogonal to the previous one: if you have mistaken information about the laws, for instance, we can distinguish what you objectively legally ought to do from what you subjectively legally ought to do. The same goes for any other source of value. Thirdly, some ought claims are arguably not normative.¹⁹

The best reading of Zeno's norms, I claim, is that (a) they are normative, though perhaps hypothetically so; (b) that do not express any specific normative *source*, but a normative *layer*; (c) that the normative layer is expressed is distinct both from the traditional 'objective' and 'subjective' ones. Let me detail these points.

Genuinely normative

First, not all *ought* claims are normative. When *ought* is used normatively, there is something amiss with somebody who sincerely accepts that something ought to be so but does not in anyway favour its being so.²⁰ When it is used non-normatively, there is nothing amiss in doing so. In their most natural reading, the sentences below make non-normative claims:

The sky ought to be cloudy tomorrow morning.

Kratzer (2010) could be used. But we need not endorse such semantics here; we can leave open how exactly the various things that *ought* claims express or convey correlate with the semantics of "ought".

¹⁹See *e.g.* Broome (2013, chap. 1).

 $^{^{20}}$ Some characterize an all-things-considered ought as the sense of "ought" which makes the following schema true: it is irrational to believe that you ought to ϕ without intending to ϕ (Broome, 2013, 22). Normative oughts may be characterized by a weaker schema: it is irrational to believe that S ought to ϕ without being to some extent in favour of S ϕ -ing. (As the phrase is used here, one can be to some extent in favour of something without being overall in favour of it.) For some expressivists the schema holds because believing that something ought to be so in the normative sense just is to have a favouring attitude towards it. I do not want to endorse this idea here; perhaps there are cases where one sincerely believes that one ought to do something without in any way favouring it. I am content with the vague albeit clear enough idea that when one believes that they ought to do something, they would normally (they are expected to, meant to, supposed to) favour it to some extent.

The plural of "mouse" ought to be "mouses". (Broome 2013, 9)

The first would normally be used to express what you expect to be the case — an "epistemic" reading of *ought*. It would not suggest that you somehow favour a cloudy sky. The second may be used not to express one's expectations about English nor one's recommendations for it, but to register instead a regularity.

The simplest view on Zeno's norms is that they are normative. Unfortunately, things are not so simple. For Zeno's norms may also be *hypothetical oughts*, which, if there are any, are neither of the straightforward normative type nor of the straightforward non-normative type. The idea is best illustrated with *have to*. Consider:

How can one get to the sarcophagus? — Well, it's not easy. You have to demolish the painted wall in the antechamber.

The dialogue may take place between two people to whom it is very clear that nobody ought in any sense to get to the sarcophagus. So the claim is not a straightforward normative *ought*. On the other hand, the claim has normative implications. For it clearly follows from what the second person says that *if one has to get to the sarcophagus*, then one has to demolish the wall. Thus the claim may be understood as a shorthand for the conditional form such as "if you want to get to the sarcophagus you have to demolish the wall". Similar phenomena may arise with *ought*. Call them *hypothetical oughts*.

Whether hypothetical *oughts* are normative or non-normative is moot. Conditionals of the form "If you want *A*, you ought to *B*" have at least in principle two readings, often labelled "narrow-scope" and "wide-scope".²¹ On the first reading, the claim is that if some condition obtains (you want *A*), some norm holds (you ought to *B*). On the second, the claim is that a norm holds, whose content is: (either you do not want *A* or you *B*). On the first reading, the claim is strictly speaking *not* normative, though its combination with additional claims may entail something normative. On the second reading, the claim *is* normative. It forbids a certain combination of attitude and action. The two readings would arise for *ought* claims that are implicitly hypothetical, if there are any.

²¹Broome (1999); see Kolodny (2005) and Broome (2013) for further discussion. Talk of scope should not be taken too literally. The two readings may be achieved by several linguistic mechanisms: for instance, one can get the "narrow-scope" reading by having a wide-scope *ought* whose domain of quantification is restricted by the *if* clause.

Now some philosophers would treat vast ranges of *ought* claims as hypothetical. Some would treat all *prudential ought* claims as hypothetical; some would treat all *pro tanto ought* as hypothetical. On a wide-scope account, some could even treat all *oughts* as hypothetical, that is, they could hold that correct *oughts* claims all bear on combinations of attitudes and actions.

I do not want to take a stake on such views. I want to leave open whether Zeno's claims are of the simple normative kind or of the hypothetical one. Since the later may turn out to be not strictly speaking normative, I leave open that Zeno's claims are not strictly normative. All that matters here is that they are no less normative than *e.g.* ordinary prudential *ought* claims are.

A distinct normative layer

Second, normative layers. Let us first illustrate the common distinction between "subjective" and "objective" *ought*. A doctor has a patient with a well-known disease. There are two treatments for it, the old and the new. The old has strong side-effects and is now almost entirely out of use. The doctor naturally prescribes the new. But the patient turns out to have an hitherto unknown allergy to it. The doctor then switches to the old and the patient is cured. Is the following true?

The doctor ought to have given the old treatment straight away.

We are pulled both ways. On the one hand, the right treatment for the patient was the old one. So the doctor ought to have given it straight away. If we had knew in advance of the patient's allergy, that is what we would have told the doctor, and it would have been correct for us to do so. On the other hand, the doctor did what she should have done. Since the new treatment is better, and there was nothing to indicate that the patient would react badly, she had to give the new treatment. Indeed it would have been inappropriate for her to give the patient the old one. So what ought the doctor to have done? A common answer to the puzzle is to distinguish two senses of *ought*, called "subjective" and "objective". What one *objectively ought* to do is what one ought to do in view of the facts. What one *subjectively ought* to do is what one ought to do in view of one's information or one's perspective on the facts. From the doctor's original perspective, the right action was to prescribe the new treatment. But in view of the facts, the right action was to prescribe the old treatment straight away.

There are two misconceptions to avoid here. The first is to think that subjective *ought* claims are not normative. For instance, one may think that "S

subjectively-ought to F" is roughly equivalent to "S believes that they objectivelyought to F". The fact that it expresses S's belief about what they in fact ought to do would explain why we expect S to act accordingly. But the fact that it merely expresses S's belief about what they ought to do would mean that it is not normative. But that picture of the relation of the two oughts is wrong. To see this, it is best to consider a case where the two come apart and the subject knows that they do. Regan's Mine Shafts story (1980, 265n1, see also Parfit 2011, 159) is one such case. Ten miners are trapped either in shaft A or shaft B, but we do not know which. The water is rising, and we have three options : open gate A, open gate B, or open both. If we open only the gate of the shaft where they are, they will all die; if we open the gate of the other shaft, they will all be save. If we open both gates, one of them, but only one, will die, no matter what shaft they are in. In that case we know that what we objectively ought to do is either to open gate A or to open gate B. It is not to open both gates. But arguably what we subjectively ought to do is to open both gates. Since we do not know which shaft the miners are in, we must minimize risk and avoid the death of all. The cases illustrates several points about the relations between 'subjective' and 'objective' ought claims. First, it shows that we subjectively ought to do is not what we believe we objectively ought to do : for we do know that closing both shaft is not what we ought to do in view of the facts. Second, it shows that there is something genuinely normative about 'subjective' oughts: there is a clear normative sense in which we ought to close both shafts. It is neither the expression of some non-normative standard nor a mere appearance or illusion. Third, it shows that there is something genuinely normative about 'objective' oughts: if we learned that we objectively ought to do close shaft A, then that would become what we subjectively ought to do as well. There is a (hard to specify) sense in which 'objective' oughts prevail over 'subjective' ones wherever possible. So both oughts are genuinely normative. Now once we have said what S objectively ought to do and what they subjectively ought to do, it is tempting to react as follows: "granted, what S ought to do in view of the fact is this, and what they ought to do in view of their information is that, but what ought they to do in the end? What ought they to do, simpliciter?". But the question makes no sense; the two oughts both hold, they are both normative, and they do not conflict.

The second misconception is that the 'subjective' and 'objective' oughts express different sources of normativity. We we are confronted with sources of normativity, there is conflict: what we owe to the state vs. what we owe to our family, what we ought to do for ourselves vs. what we owe to do to others, what we ought to do for the task at hand vs. what we ought to do for

our long-term goals, and so on. (Of courses two sources of normativity may prescribe the same thing; but at least conflict may in principle arise.) Conflict is solved by compromise, prevalence of one norm, or even not resolved at all. But it always involves some considerations in favour of doing something and some considerations against that are balanced against each other. Nothing such arises with 'objective' and 'subjective' readings of ought claims. First, 'objective' ought is not one category of ought alongside moral, legal, prudential and so on. There is no situation where what we "objectively" ought to do is *F* but what we "prudentially" ought to do is not *F*. Rather, for each of the categories moral, legal and so on, there are objective oughts. It may be, for instance, that in view of the facts the *legal* thing to do would be F but, still in view of the facts, the *prudent* thing to do would not be F. The same holds for 'subjective' ought: it is not one category alongside moral, prudential and so on. One may be tempted to think so, if one calls it the *ought* of rationality; one could think that in some cases we have a conflict between what is rational to do and, say, what is morally right to do. But I think this is a confusion. For each normative source such as the legal, the moral and so on, there is a rational way to pursue it; to each of these correspond distinct 'subjective' ought claims. The kind of cases where we seem to pit morality against rationality are in effect cases where we pit what is morally required against what is prudentially required — for instance, what we subjectively ought to do, morally speaking and what we subjectively ought to do in view of our interests alone. Second, 'objective' ought and 'subjective' ought are not such a conflict with each other. Suppose we observe somebody caught in a dilemma between two moral duties. We will see the two duties in opposition; we will often look for a compromise; if one duty prevail, we will still feel the force of the considerations brought by the other. We may say, for instance: "On the one hand she must be true to her mother; on the other hand she should not break the promise made to her sister to keep their secret; the only way for her to do both is to avoid talking to her mother at all; but if it came to that, she would have to betray her sister." Contrast when we observe a "conflict" of 'objective' ought and 'subjective' ought, as in the Mine Shafts case. We do not say: "On the one hand the miners are in A and she objectively ought to open gate B; on the other hand she subjectively ought to open both gates". We do not try to find a compromise between the two oughts. If we are about to offer advice, the 'objective' ought alone will matter and considerations of what the agent subjectively ought to do at that time will have no force whatsoever. If we are discussing whether the agent acted in a stupid or evil manner, the 'subjective' ought alone will have weight. Each corresponds to a distinctive layer of normative claims. They are not different sources in conflict within a single layer.

The best way to understand Zeno's norms is that they intend to capture a distinctive *layer* of normativity. The 'subjective' *ought* is supposed to correspond to what is best in view of one's 'information' or from one's 'perspective'. But there are various notions of 'information' or 'perspective' to consider. In our original example, giving the new treatment was what the doctor ought to do in view of what they knew, but also what they believed. But the two can come apart. In some cases, the doctor irrationally believes that the patient will respond well to the new treatment. Doing so would then be what she ought to do in view of what she believes but not in view of what she knows or rationally believes. In some cases, what the doctor ought to do in view of what they know may differ from what they ought to do in view of what they rationally believe. That may be so, for instance, if all the doctor as ever heard about the new treatment is in fact a fabrication; she rationally believes it, but there is nothing to it. If so in view of all that she knows — namely, that the *old* treatment works — what she ought to do is to give the old treatment, even though in view of what she rationally believes it is to give the new one. Now one may argue that these difference correspond to distinct but genuine normative layers. If someone does what she believes to be best, without rationally believing that it is best, then there is a sense in which they do what they ought to be doing and a sense in which they do not. If someone does what she rationally believes to be best, without it being best in view of what they know, then, it is argued, here as well there is a sense in which they do what they ought to be doing and a sense in which they do not.

So the most charitable way to assume Zeno's norms is to grant that there is a distinctive *layer* of normativity about which they hold. There is a sense of *ought* in which one ought to believe only what one knows and one ought to act only what one knows. It is distinct from some 'subjective' *oughts*, such as (some notions of) rationality: it is sometimes rational to act on something we do not know, for instance when everything misleadingly indicates that we know it.²² It is not an 'objective' *ought*: in the Mine Shaft case, it matches the 'subjective' one instead.

²²Philosophers sometimes distinguish "procedural rationality", which is merely a matter of having coherent attitudes and "substantial rationality", which requires more, *e.g.* having beliefs that fit the evidence and desire things that are worth desiring. These would also correspond to distinct layers of rationality; the one is clearly distinguished from a knowledge-based normative layer; the second less clearly so.

No distinctive source

We should not read Zeno's norms as expressing a distinctive *source* of normativity. That is most defensible for the Knowledge–Action Principle. We may imagine a case where *in view of what you know*, you morally ought to help someone but it is not in your best interest; while *in fact*, you morally ought not to help them but it is your best interest. That may happen for instance if an eccentric rich man pretends to be in dire poverty and need your help. In view of what you know, the moral thing to do is to help them, though it is not in your interest. In view of all relevant facts, there is no moral requirement to help him, though doing so will happen to bring you a hefty reward. In such a case we do not have a knowledge-based *ought* that enters in conflict with a moral one and a prudential one. Rather, morality and prudence each generate their knowledge-based ought alongside their 'objective' one. The requirement to act on what you know is a distinctive normative *layer*, not a normative source.

With the Knowledge – Belief principle, the claim is more debatable. Consider standard cases of believing for practical reasons. An athlete knows that they have no chance to win the race, but they also know that believing that they will win will improve their time. One may feel a conflict analogous to a conflict of prudence and morality here. What the athlete *epistemically* ought to do is to believe that they will not win; what they *prudentially* ought to do is to convince herself that she will win. The two *oughts* conflict; and both are at the layer of knowledge-based *oughts*.²³ So one may be tempted to count the Knowledge – Belief principle to reflect a particular normative source.

If we did so we would get Meno's thesis fairly straightforwardly. An attractive hypothesis about normative sources is that they all reflect values. What one ought morally to do derives from what is morally good, what one ought prudentially to do derives from what is good for one, and so on. More precisely, what one 'objectively' X-ly ought to do is what is X-ly good; what one X-ly ought to do in view of what one beliefs is what has higher expected X-ly value, where the expectations are given by one's beliefs; what one X-ly ought to do in view of what one beliefs is what has higher expected X-ly value, where the expectations are given by one's knowledge; and so on. Now if the hypothesis holds, one 'objectively' ought to believe only what one knows

²³To parallel the foregoing case, one can imagine a variant where *in view of what the athlete knows*, they have no chance to win but believing that they do is likely to improve their time, while *in fact* they are likely to win but the belief would lower their chances. We seem to have an epistemic *vs.* prudential conflict both at the 'objective' and knowledge-based layer, and the two layers are distinct.

if and only if it is bad to believe what one does not know. So the norm would entail (part of) Meno's Thesis. ²⁴ The resulting set of claims is virtually indistinguishable form the non-consequentialist derivation of the Knowledge–Belief principle we examined earlier. While we would have derived Meno's Thesis, it would not be clear, however, that we would have explained it. For one may think that the hypothesis holds because values ground norms; if so, we are in effect explaining the Knowledge–Belief Principle by Meno's Thesis and not the opposite.

Be that as it may, I will focus on another construal of the Knowledge–Belief Principle. It is best seen by rewriting the principle thus:

Knowledge–Belief Principle (rewritten) One ought to believe only what is true in view of what one knows.

We start with the idea that having true beliefs and not having false beliefs is a normative source; that is, something that a source of oughts that may in principle conflict with moral, prudential, legal considerations and so on. The source will generate various layers of oughts. Roughly: that one 'objectively' ought to believe what is in fact true; that one purely subjectively ought to believe what is true in view of what one believes; and that one 'knowledgeably' ought to believe what is true in view of what one knows. The latter is the Knowledge-Belief Principle. It reflects both a normative source, in its requirement of believing the truth, and a normative layer, in its focus on what one ought to believe in view of what one knows. To take the dimension of layer apart, we may focus on a more general principle:

Generalized Knowledge–Belief Principle One ought to form one's belief in view of what one knows.

The principle generates epistemic oughts, when combined with the idea that one ought to form true beliefs, but also prudential oughts, when combined with the idea that one ought to form useful beliefs, for instance.

Back to the derivation

The Knowledge-Action principle and the Generalized Knowledge-Belief principle delineate a significant normative role for knowledge. Can we derive

 $^{^{24}}$ To entail Meno's Thesis we should add the positive principle that one should believe what one is in position to know.

from them the claim that knowledge is better than belief that is not knowledge? Engel (2009) takes the answer to be negative, but does not discuss why. Prima facie, there seems to be a way. First, we assume that believing the truth is good and that acting in the light of true propositions is good. These are assumptions about value, but distinct from the straightforward assumption that knowledge itself is good. Second, it follows from Zeno's norms that in order to do these good things, we are in some sense required to have knowledge—and not merely true belief. (The sense in which we are required to do so is the one that corresponds to the sense *ought* has in Zeno's norms.) From this, it seems, we can conclude that knowledge is better than true belief.

The derivation has some appeal. It seems plausible that (normative) conditions for doing good things are themselves good. If knowledge (normatively) allows you to form beliefs and act on them, then insofar as these beliefs and acts are good, knowledge would seem to be a good thing.

As it stands, the derivation fails. It is not in general true that (normative) conditions for doing good things are themselves good. Apologizing for one's faults is good; one ought to apologize for one's mistakes only if one made mistakes; but making mistakes is not good. That being said, there are undeniably many cases where conditions for doing good things seem good, and seem good precisely in virtue of being such. One would need a restricted version of the principle. If the restricted version applies to the case of knowledge, the derivation would succeed.

The route from Norms to Values is more promising. Whether it ultimately succeeds depends on whether we can find a plausible motivation of the idea that knowledge is good because it is normatively required to do good things. It is worth stressing that, if it is ultimately successful, it would yield a picture of the value of knowledge that is at odds with much of the current literature. Much current literature tries to show that knowledge is valuable by showing that it is a worthwhile thing to aim at, either as an end or as means to some end. On the present perspective, knowledge would be good because it is normatively required to aim at anything.

²⁵Engel (2009, sec. 5) calls the idea that knowledge is relevant to what we ought to do "pragmatic relevance" and writes: "Of course if the phenomenon of pragmatic encroachment reduces, as I have claimed, to that of pragmatic relevance, question (2) [the question whether pragmatic encroachment explains why knowledge is valuable] has to be answered in the negative". He does not discuss the idea further, but assumes that it is only if pragmatic encroachment is *more* that something like the Knowledge–Action principle that it could explain the value of knowledge.

4. Doing Without Value

We started with a pair of ideas: that knowledge has value, and that it enters certain norms. We have tried to derive one idea from the other. I have argued that the norms cannot be derived from the value and I have expressed doubts as to whether the value can be derived from the norms. Now I want to stress another perspective: namely, that once we assume the knowledge-involving norms, it is unclear whether there is any way to motivate the idea that knowledge has value. I have made the suggestion elsewhere (Dutant, 2012; forthcoming). Here I will discuss a couple of objections to it.

Why should we think that knowledge has more value than true belief? Many philosophers treat it as a (at least *prima facie*) platitude in need of no defence. To them, it is (at least *prima facie*) obvious that there is something commendable about believing something when you know it that there is not about believing it when you do not know it. But note that that is explainable on the basis of the Knowledge – Belief principle alone. If that principle holds, then there is a sense in which you do what you ought to do when you believe p while knowing p and you do not when you believe p while not knowing p. That is enough to explain that there is something commendable about knowledge that mere true belief lacks. That does not require or entail that there is something genuinely *better* about it. What we ought to do and what is good may come apart. If you face a choice between A and B, if A is bad but everything indicates that it is good and better than B, then there is a sense in which you ought to choose A, but we may deny that choosing A is genuinely good. ²⁶

Some will think that doing what one ought to do has value in itself.²⁷ They would endorse the straightforward derivation of Meno's thesis from the Knoweldege-Belief principle that we sketched earlier. But there is no need to accept that view. In particular, normative *sources* and *levels* may stand differently in their relation to value. There is some plausibility that normative *sources* reflect different values: the morally required derives from the morally good, the prudentially required derives from the personally good, and perhaps the epistemically required derives from the value of believing the truth.

²⁶We use adjectives like "good" fairly liberally. There are very natural uses on which "you should do that" and "doing that would be good" are virtually interchangeable. With these uses we may well say that choosing *A* was the "good" choice. But in laying out a theory of value and norms one may need use "good" and "value" more strictly. Strictly speaking, choosing *A* was not good; it just *seemed* good. That it seemed good made it the choice you ought to make; but that need not itself make it genuinely good.

²⁷Piller (2009) pursues this line to explain the value of knowledge.

But normative *layers* do not require additional values. Rather, each normative layer correspond to a different way to derive an *ought* from a value. Given moral goods, there is when we objectively ought to do for the moral good; what we ought to do for the moral good in view of what we believe; what we ought to do for the moral good in view of what we know; and so on. We have argued that the Knowledge–Action principle and the Generalized Belief–Action principle express a normative layer and no particular normative source. If that is so, they need not be associated with any value at all.

Do we have any other reason to think that knowledge is valuable? Knowledge is definitely something that matters. But as we have seen, its normative role is enough to explain that it does.

The best reason to think that knowledge has value is, I think, the idea that knowledge is worth aiming at. We want knowledge; we strive for it and we are ready to make sacrifices for it. We are not foolproof; but assuming that we are right in this, knowledge is something we ought to aim at. Conversely, a theory that denies that knowledge is something we ought to aim at has to claim that we are misguided in that respect. If knowledge is good, we have a straightforward explanation of why we ought to aim at it. For in general, we ought to aim at good things. Conversely, if knowledge has no value, it becomes doubtful whether we should aim at it.

The idea raises a difficulty for the view that knowledge plays a important normative role but has no value. Zeno's norms alone do not entail that we should *acquire* knowledge. One may comply with both norms by avoiding belief and action altogether. Less radically, one who complies with the norms at one point may comply with them onwards by not acquiring any new belief and acting only on what they already know. In reply, I would point out that the norms entails requirements to acquire knowledge *in conjunction* with other norms or values. For instance, if we assume that it is better to act on more relevant facts, we may derive that one ought to acquire more knowledge of relevant facts. Similarly, we may assume that it is good to have true beliefs on a range of topics; if so, we may derive that one ought to acquire more knowledge. It is not trivial to work the reply out properly. For instance, it is not obvious that it is *always* better to act on more relevant facts. For our purposes, it is sufficient to show that there are some ways to ground requirements to acquire knowledge that do not assume that knowledge is itself valuable.

Another difficulty for the view that grants knowledge a normative role but no value is worth discussing. The view assumes Zeno's norms, or something like them. The assumption is not trivial. As we have argued, the best way to understand them is to postulate a normative layer distinct from the tradi-

tional 'objective' and 'subjective' ones. One may fell uneasy about the very idea of layers of normativity, and about the idea that knowledge-based ought claims delineate a distinctive layer. In reply, I will make a couple of points. First, as I have argued, it is difficult to make sense of cases such as Regan's Mine Shafts case without accepting that there are at least two distinct layers of normativity; each genuinely normative but not conflicting. Second, once we have admitted the idea of layers, it is easy to see how they multiply. Within a belief-based layer alone, we can often distinguish opposite ought claims derived from various natural subsets of one's beliefs. For instance, one may have a set of salient beliefs in view of which one ought to do A; but one may at the same time have some deeply buried beliefs, such that in view of them and the salient ones, one ought not to do A. In such cases, an onlooker may feel both the pull of "they ought to have done A" and the pull of "they ought not to have done A''. Or again, one may impeccably infer a conclusion from a set of crazy premises. In such a case we may feel both the pull of "they ought to have inferred that conclusion" and of the opposite; one is what one ought to do in view of one's premise beliefs, the other what they ought to do in view of some broader background. The divisions may be multiplied: there may be cases where one's conclusion is correct in view of the premises, insane in view of the premises of the premises, correct again in view of the premises of the premisses of the premisses, and so on indefinitely. I see no reason to reduce the profusion of these normative layers to one or two. As long as each oughtquestion we care about manages to pick up a specific enough layer, we can leave with many oughts. On the backdrop of these many normative layers, the knowledge-based one is not a cost.

Summing up, once we grant that knowledge plays a central normative role along of the lines of Zeno's norms, it is not clear that there is anything left to motivate the idea that knowledge is of distinctive value. So Zeno's norms can be used to explain *away* Meno's thesis. In that perspective, the normative role of knowledge allows us to explain why it matters without having value.

5. Conclusion

I have highlighted two possible answers to the question why knowledge matters. One is that it has value. Another is that it plays a significant normative role. I have granted that if knowledge *had* value, or if it *did* play the alleged normative role, then it would matter. For most of the discussion I have remained neutral on whether knowledge has value or does play that normative

role. My central question has been instead whether we can derive one idea from the other. That is, whether assuming the idea that knowledge has value — and some defensible general hypotheses about norms and values —, we could derive the claim that it plays the alleged normative role. Or whether, assuming that knowledge does play that role — and some defensible general hypotheses —, we could derive the claim that it has value. I have found the route from Value to Norms unsuccessful. The main problem here is that the idea that knowledge has value does not seem enough to derive the idea that one should act on what one knows. I have found the route from Norms to Value more promising, though a complete path is missing. The main idea here is that knowledge is good because it is normatively required to do good things, such as believing the truth and acting in view of true propositions. But since not all normative condition for doing something good is itself good, we still lacked an explanation of why knowledge would be so. Finally, I have suggested an alternative perspective, on which we would not try to derive the idea that knowledge has value from its normative role, but rather use its normative role to explain away the idea that it has value. The general idea is that if knowledge does play the normative role in question, then the fact that it does explains while knowledge seems to be something that has value. But there is no need to think that it has; all that matters about knowledge could be explained by its normative role.

6. References

Selim Berker. Epistemic teleology and the separateness of propositions. *Philosophical Review*, 122(3):337–393, 2013a.

Selim Berker. The rejection of epistemic consequentialism. *Philosophical Issues*, 23:363–87, 2013b.

John Broome. Normative requirements. Ratio, 12:398–419, 1999.

John Broome. Rationality through Reasoning. Wiley-Blackwell, 2013.

Edward Craig. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford University Press, 1990.

Julien Dutant. The value and expected value of knowledge. *Dialogue*, 51(01): 141–162, 2012.

Julien Dutant. In defence of swamping. Thought, forthcoming.

- Pascal Engel. Pragmatic encroachment and epistemic value. In Adrian Haddock, Alan Millar, and Duncan Pritchard, editors, *Epistemic Value*. Oxford University Press, 2009.
- Roderick Firth. Are epistemic concepts reducible to ethical concepts. In Troyer, editor, *In Defense of Radical Empiricism: Essays and Lectures by Roderick Firth*, pages 237–249. Rowman & Littlefield, Lanham, Md., 1998a.
- Roderick Firth. Chisholm and the ethics of belief. In John Troyer, editor, *In Defense of Radical Empiricism: Essays and Lectures by Roderick Firth*, pages 143–155. Rowman & Littlefield, Lanham, Md., 1998b.
- Roderick Firth. Epistemic merit, instrinsic and instrumental. In John Troyer, editor, *In Defense of Radical Empiricism: Essays and Lectures by Roderick Firth*, pages 259–71. Rowman & Littlefield, Lanham, Md., 1998c.
- Michael Frede. Stoics and sceptics on clear and distinct impressions. In *Essays on Ancient Philosophy*, pages 151–176. Clarendon Press, 1983/1987.
- Michael Frede. Stoic epistemology. In K. Algra, J. Barnes, Mansefld J., and Schofield M., editors, *The Cambridge History of Hellenistic Philosophy*, pages 295–322. Cambridge University Press, 1999.
- H.J. Hankinson. Stoic epistemology. In B. Inwood, editor, *Cambridge Companion to the Stoics*. Cambridge University Press, 2003.
- Mark Kaplan. It's not what you know that counts. *Journal of Philosophy*, 82: 350–363, 1985.
- Thomas Kelly. Epistemic rationality as instrumental rationality: A critique. *Philosophy and Phenomenological Research*, 66:612–640, 2003.
- Niko Kolodny. Why be rational? *Mind*, 114(455):509–563, 2005.
- Angelika Kratzer. The notional category of modality. In *Collected Papers on Modals and Conditionals*. Oxford University Press, 2010.
- Maria Lasonen-Aarnio. Unreasonable knowledge. *Philosophical Perspectives*, 24(1):1–21, 2010.
- Anthony A. Long and David Sedley, editors. *The Hellenistic Philosophers*, volume 1. Cambridge University Press, 1987.
- Derek Parfit. On What Matters, volume 1. Oxford University Press, 2011.
- Charles S. Peirce. *The Philosophy of Peirce: Selected Writings*. Harcourt Brace, 1950.

- Philip Percival. Epistemic consequentialism: Philip percival. *Supplement to the Proceedings of The Aristotelian Society*, 76(1):121–151, July 2002.
- Christian Piller. Valuing knowledge: a deontological approach. *Ethical Theory and Moral Practice*, 12:413–28, 2009.
- Donald Regan. Utilitarianism and Cooperation. Oxford University Press, 1980.
- Robert Stalnaker. Epistemic consequentialism: Robert stalnaker. *Supplement to the Proceedings of The Aristotelian Society*, 76(1):153–168, July 2002.
- Timothy Williamson. Knowledge and its Limits. Oxford University Press, 2000.