What We Talk About When We Talk About Epistemic Justification

Jack Lyons

University of Arkansas

Philosophers are reflective people, and it is good to sometimes turn our attention from the things we really care about—truth, beauty, goodness, and the rest—to our own habits, practices, and presuppositions in theorizing about such things. Stewart Cohen (this issue) thinks that epistemologists in particular have not been sufficiently reflective in this way and that it has adversely affected our theorizing. This is an important point, and Cohen's argument is an insightful and challenging one. It is especially challenging to epistemologists (like me) of a particular theoretical bent, because he thinks that epistemology would be better done if we'd stop talking about justification and just talk about rationality instead.

1. Justification and Rationality

'Epistemic justification', Cohen says, is a term of art—ask any competent English speaking non-philosopher which beliefs are epistemically justified, and she'll have little idea what's being asked. Because it's a technical term, however, its basic usage ought to be clearly specified before we go fighting about who has the right theory of what justification is. It hardly seems fruitful to engage in such disputes if we aren't even sure what it is we're arguing *about*. And of course, epistemologists are eager to offer substantive, contentious theories about the deep nature of epistemic justification, but pretty much silent about what we mean by 'epistemic justification'.

Cohen thinks that we have a much better pretheoretical grip on the term/concept 'rational', and this favors theorizing about rationality, rather than about justification.[1] "Justification theorists," however, (those epistemologists who frame their theories in terms of epistemic justification rather than, say, rationality) typically *deny* that what they're talking about is rationality, insisting sometimes that certain agents, like victims of an evil demon, for example, might be rational but unjustified. But then we have no clear idea what's meant by 'justification', and so we have no clear idea whether these theories of justification are at all adequate. And this must mean that we have no good reason to believe them.

Justification theorists are, of course, happy to tell us what they think justification is. But, Cohen argues, what they tell us is always something theoretically loaded, something controversial; it is a substantive epistemological position rather than a characterization of what that position is a position *on*. We are given a theory, rather than a specification of what it is that theory is trying to capture.

This puts the rationality theorist at a decided advantage, says Cohen. For even if the rationality theorist can't do any better at telling us in theory-neutral terms what rationality is, *we don't need to be told*. Questions about rationality are, unlike questions about epistemic justification, readily pretheoretically understood, without any special hand-holding or explanation from the epistemologist. So we should stop theorizing about epistemic justification and theorize about rationality instead. If it turns out that certain epistemic theories aren't plausible contenders for a theory of rationality (e.g., perhaps, reliabilism), then this is bad news for those theories, at least until they can help us understand what they're theories *of*. We already understand what the rationality theorist is talking about.

I tend to think that just the opposite is true. I suppose I'm one of these "justification theorists," because I write and talk and think about epistemic justification while avoiding writing

_____

[1] We should distinguish terms, concepts, and properties. Terms are linguistic entities: words or phrases. I will suppose for the purposes of this paper—though only for these purposes—that concepts are the (shared, public) meanings of terms: Fregean Senses or something similar. I will take terms to be individuated in such a way that a single term (an ambiguous one, for instance) can express more than one concept. Properties are the things picked out by the terms and concepts. So the terms 'water' and 'H_2O' express different concepts but pick out the same property. Because of the close connections among properties, terms, and concepts, my discussion will switch back and forth among them, as dictated by expository convenience, when the distinctions aren't pertinent to the current discussion.

and talking about rationality. The reason I do is that I feel like I have a much better handle on what's meant by 'epistemic justification' than on what's meant by 'rational'. Cohen claims (p. xxx) that justification theorists deny that justification and rationality are the same, but I for one don't understand 'rational' enough to endorse such a denial; this is one of the main reasons why I —and, I suspect, other justification theorists—opt to frame theories in terms of justification rather than rationality. It's not that we think that justification and rationality are different things, so much as that we see matters as exactly inverted from how Cohen sees them. *Justification* is the clear concept, while *rationality* is too squishy to theorize about in any serious way.

2. Naive Speakers and Rationality

Cohen's argument starts by noting that if we were to ask an ordinary, competent but not philosophically trained, English speaker whether some belief was epistemically justified, she would very likely have no idea what she was being asked. Asked if some action is morally justified, on the other hand, she would understand the question perfectly well, even if she didn't have a good answer. Questions about rationality, Cohen thinks, are more like questions about moral justification than like questions about epistemic justification. Ask whether a given belief is *rational*, and most people will understand the question just fine without any special training or indoctrination.

It's hard to move immediately from data about words to conclusions about concepts, but suppose it's true that philosophically naive English speakers would report not understanding questions about epistemic justification (this is an empirical supposition that I lack systematic data for or against); suppose they're right, that they really don't understand the questions; and suppose further, for the moment (I don't think it's true), that the proper explanation for this fact is that *epistemic justification* is not part of our naive, pretheoretical conceptual repertoire (as opposed to our naive *linguistic* repertoire). What would that show? Only that the term/concept in unfamiliar. But that isn't something we really care about; what we care about is whether the term/concept is clear. This is especially true for the present purposes, where the issue is whether we genuinely understand what epistemologists are debating about. I mean for the difference

between familiarity and clarity to be roughly this: a familiar term is one that most ordinary speakers of the language can easily attach a concept to; a term is clear to the extent that it is free from vagueness, ambiguity, polysemy, and is completely grasped. I won't try to make this more exact, but it's obvious from the quick gloss that a term can be familiar without being clear, and vice versa. A lot of folk concepts/terms are highly familiar but highly unclear, while many mathematical concepts are quite clear but highly unfamiliar.

In light of this distinction, we might want to revisit the naive speaker test. A familiar term is one, such that a naive but competent speaker will generally *feel like* she understands the question, although if the term is sufficiently unclear, she may not really understand the question after all. The question might have been ambiguous, and she only understood one disambiguation. Or the term might have been so unclear that there wasn't really any genuine question to understand, yet familiar enough that she felt as if she were being asked a determinate but difficult question.

'Rational' is a more familiar term than 'justified', but that doesn't mean it's any more clear; in fact, I will argue later that 'justified' is clearer than Cohen thinks. Here and in the next section, I want to argue that 'rational' is less clear than he thinks. That 'rational' is less than fully clear is illustrated in part by the fact that even Cohen uses the term in two importantly different senses. He uses 'rational' and 'rationally justified' interchangeably, without distinction: "it makes no sense to distinguish between rational belief and justified belief. It is rational to believe just in case so believing is rationally justified" (p. xxx). But 'rational' is playing very different roles in these two phrases. In 'rationally justified', the term is used to specify a normative *domain* (the rational, as opposed to the legal, the moral, etc.); in 'rational' full stop, it picks out a normative *status* within one of these domains (epistemically rational, as opposed to epistemically irrational; or practically rational, as opposed to practically irrational, etc.).

To clarify, consider the familiar point that 'rational' has related but distinct uses in different contexts:

- That was the rational thing to believe, based on the evidence you had at the time.
- In the current market, it's rational to sell most of your stocks and invest in real estate

4

- She responded rationally to his attempts to draw her into a fight.

The first two examples here illustrate the contrast between what we normally call epistemic rationality and practical rationality (I don't know that the third has a name). 'Rational' and 'irrational' are polysemous (i.e., each has distinct but semantically related meanings), marking out different properties in different contextually determined domains of evaluation: epistemic rationality in one context, practical rationality in another. Suppose a friend smokes cigarettes and has been reading Pascal and James and chooses to believe that smoking doesn't cause cancer, in an effort to feel better about her habit. That belief is epistemically irrational (because of the evidence against it) and also practically irrational (because it causes her to keep smoking, the benefits of which are vastly outweighed by the costs and risks).[2] We've attributed to this friend two distinct properties, both normative and with negative valence, but different from each other.

This is the ordinary, status sense of 'rational' (/'irrational'): given a normative domain, the term picks out a status within it. Given a dimension of appraisal, the term locates things (beliefs in this case) in some region along that dimension; 'rational' locates them at one end, 'irrational' at the other end. But Cohen introduces a neologistic sense as well, a domain sense, whereby 'rational' distinguishes the epistemic domain from the practical, the legal, the moral, etc. Thus, some things are morally justified, some things are legally justified, and some things are "rationally justified." 'Rational' in the status sense has an antonym: 'irrational'; it has contrasts in the domain sense, but no antonym.

This domain sense comes out in the second stage of Cohen's argument. Not only (stage one:) do the philosophically naive not have a good grip on 'epistemic justification', but (stage two:) epistemologists have failed to rectify this deficit, since all their explanations of the target (in terms of truth, evidence, the difference between knowledge and unGettiered true belief, etc.) make substantive theoretical commitments and thus fail to offer a theory-neutral specification of what it is these theories are attempting to capture. What's at issue for Cohen here (Section I) is whether the justification theorist can specify what it is that distinguishes *epistemic* justification

---

[2] It doesn't matter exactly what is being assessed here for practical irrationality, the belief, the intention to believe, the trying to believe, etc.

from moral justification, legal justification, and the like. 'Justification' is clear enough: to say that a belief is justified is to say that it has some positive normative status, but we need to know which domain or dimension of appraisal it scores positively on. Because Cohen's opponent can't answer this without invoking controversial assumptions or unfamiliar technical terms, Cohen suggests we replace talk about epistemic justification with talk about rational justification, i.e., rational belief; 'rational' is familiar enough that he doesn't need to explain to us what it means—we already know.

There are three problems here. First, it's now unclear why Cohen calls his opponents "justification theorists," given that the part of 'epistemic justification' he's puzzled by isn't the justification part but the epistemic part. Second, Cohen seeks to solve the justification theorist's problem by invoking a term/concept that's already part of our pretheoretical repertoire. But 'rational' *in the domain sense* doesn't fit the bill. This is a neologism; I've never heard a non-philosopher describe a belief (or anything else) as "rationally justified." It's only in the status sense that 'rational' is an old friend. This puts Cohen on an equal footing with his opponent; he's using an old word for a new purpose; we're using a new word for the same purpose. Either way, we all have to find some way of indicating what the relevant normative domain is and how it differs from the legal, moral, etc. Third, this isn't a *serious* problem for either Cohen or his opponent, because we all understand well enough (though surely imperfectly) the difference between the moral, legal, and rational-cum-epistemic domains. I know well enough what Cohen means when he says "rationally justified"; it's what I mean when I say "epistemically justified." And I think he's right that efforts to explicate the domain in theory-neutral terms are bound to fail. And he's right that this does not pose any crippling methodological problems, given that we have a decent enough pre-theoretical grasp on the matter. The *word* 'epistemic' is unfamiliar to most competent English speakers (most speakers probably don't even know it's a word; some—especially in English departments, for some odd reason—know it's a word but are wrong about what it means). But that doesn't mean the *concept* is unfamiliar, which is what really matters. Even if 'rationally justified' had been part of the folk vernacular, this would have solved the justification theorist's problems for her: epistemic justification is no different from rational justification; if we understand one term we understand the other.

Because the concept of the epistemic/rational domain is both familiar and clear, it's little problem that the terms are new or neologistic. If someone needs to know which domain we're talking about, it's easy to key them in. One way is to use Cohen's helpful suggestion that the domain epistemology is concerned with is "the kind of rationality [/justification/normativity] that applies exclusively to believing" (p. xxx). Another way, however, would be to say that it's the domain that's most closely connected to evidence, or the domain one of the statuses within which is knowledge. It's not controversial that 'knowledge' is a term of epistemic (as opposed to legal, moral, etc.) appraisal, or that evidence is more closely connected to the epistemic domain than to the others; these aren't things that are debated in the journals. Of course, none of these gestures —including Cohen's—would generate the concept of the epistemic/rational domain in someone who didn't already have that concept. But that's not our problem; our problem is getting a normal adult (all of whom do already have this concept) to understand our talk.

Cohen insists that there's no difference between saying a belief is rationally justified and saying it's rational. 'Rational', therefore, is, when used in the domain sense, supposed to substitute for 'epistemic' in 'epistemically justified'; and in the status sense, it's supposed to substitute for 'justified'. Renaming "the epistemic" as "the rational," however, hurts him more than it helps. One virtue of the standard terminology is that neither 'justified' nor 'epistemic' admits of the domain/status ambiguity that 'rational' does. A minor annoyance that Cohen's introduced ambiguity creates is that if 'rational' distinguishes the epistemic from the moral, and 'irrational' distinguishes the justified from the unjustified, then if he wants to distinguish the things my smoker friend has done wrong, he has to say that she's both practically irrational and rationally irrational. I dislike the latter phrase for obvious reasons, and greatly prefer 'epistemically irrational'. (He can of course avoid the awkward phrase, but not in a few syllables; he has to say something like, "irrational in the sense that applies exclusively to believing.") A more serious problem, as discussed above, is that it leads him to think he's taking a term from our pretheoretical repertoire when in fact, he's coining a new one.

More serious yet is that this ambiguity invites us to think that there's no difference between a rationally justified belief and a rational belief, and to think that this is trivially true, as if the former expression were just a pleonastic corruption of the latter. But once we see that

'rationally justified' just means 'epistemically justified' we see that this is a substantive, indeed controversial, claim—one that justification theorists doubt or deny. The new domain sense allows 'rational' to behave like 'moral' (which also has a domain sense and a status sense). Cohen is right that something is morally justified iff it's moral, and he thinks, similarly, that something is rationally justified iff it's rational.[3] But whether this analogy is apt or not is an open question, and one that's settled not by giving 'rational' a domain sense, but by investigating whether our commonsense commitments regarding epistemic justification are at all different from our commonsense commitments regarding (epistemic) rationality. I take this up in the next section.

Cohen thinks that when we doubt or deny that a epistemically justified belief is the same as a rational belief, we must be thinking about a different domain of normative appraisal, distinct from the moral, the legal, the rational, etc. (p. xxx). This would indeed be puzzling. But we're emphatically *not* thinking about a different *domain*; we're thinking there might be different kinds of *status* within a single domain—the one that he calls the rational and we call the epistemic.

We can refuse to equate the rationally/epistemically justified with the rationally/ epistemically rational, while still equating the rationally justified with the epistemically justified. Even if 'rational' and 'epistemic' are interchangeable, this doesn't mean that 'rational' and 'justified' are. It's only the latter equivalence that the justification theorist resists, not the former. Thus, if Cohen can make sense of 'rational' in the domain sense, we can make sense of 'epistemic'. Whatever he says to make sense of it, we justification theorists can subsume as our own. We will prefer the word 'epistemic' to 'rational' for this purpose, because even though the former is less familiar, it's unambiguous and less prone to generate confusion or the sorts of problems we've just seen.

3. The Rational and the Irrational

---

[3] Pushing the parallel further, he claims that a justification theorist ought to endorse the absurd claim that a belief is epistemically justified iff it's epistemic (p. xxx). This is unintelligible, of course, but only because 'epistemic' doesn't have a status sense in the way that 'moral' and 'legal' do. This fact doesn't in any way impugn the justification theorist's view.

Let us turn, then, to 'rational' in the status sense and its relation to justification. Once again—and this is, unfortunately, mostly just autobiography—I don't feel like I understand the term well enough to deny that the epistemically justified and the epistemically rational are the same. This is why I prefer to theorize in terms of justification; even though it's a relatively unfamiliar technical term, I'm not worried that I'm going to be accidentally conveying ideas that I didn't mean to, in a way that I might well be by using 'rational'.

Although I don't feel like I have much of a handle on the term 'rational', I think the situation is interestingly different concerning the term 'irrational'. Maybe it's just my own idiolect, but I think that 'irrational' is quite a bit clearer than 'rational'. To give a rough approximation, I tend to call something—a mental state or act—irrational just in case the agent should have known better. This applies to irrationality of all sorts, not just epistemic irrationality. I'm not sure how precise this can be made, but something like the following schema seems about right:

It is irrational for S to A iff S knows something that is contrary to A,

where contrariness is understood differently for the different kinds of state that A might be. If A is a belief, then x is contrary to A iff x is (obviously) incompatible with A. If A is a fear, then x is contrary to A iff x (obviously) entails that the thing feared is harmless, or so remote a possibility as to not be worth worrying about. If A is a desire, then x is contrary to A iff x (obviously) entails that the thing desired will bring substantially more unhappiness than happiness. And so on. The unifying idea is that in all the various cases, S is Aing even though S knows something that should have led her not to A.[4]

The schema is stated here in terms of knowledge, and that's probably too strong; 'knows' should be replaced by some other epistemic status, like 'is ex ante justified in believing' or something similar. I won't worry here about the details; we can get a family of distinct but related theories of irrationality by varying how we unpack contrariness and which epistemic status we offer in place of knowledge.

---

[4] This proposal resembles one Cohen discusses in note 11, if I understand it properly.

Even without settling all the details, a few important features are apparent: (i) it unifies all or many of the different types of irrationality, (ii) it imposes epistemic requirements on these various irrationalities, (iii) it's not a reductive account of irrationality, (iv) it seems to make the irrational out to be something narrower than the complement of the rational; relatedly and most importantly, (v) even the weaker forms of the schema (those that replace 'knows' with something less stringent) make epistemic irrationality a pretty serious failing—more serious, in particular, than regular old unjustified belief.

Taking (i) and (ii) first, it is clear that this is an understanding of irrationality that covers a lot more than just epistemic irrationality, though it sees all irrationality as having an epistemic component. Notice, however, that not all irrationality involves epistemic irrationality. You might be irrational in acting in a way that you know to be disutile, without this irrational action proceeding from any irrational beliefs. And we can still allow that believing *p* might be quite rational a thing to *do*, in some Pascalian way, even if it's an irrational thing to *believe*.

Partly because 'irrational' is an established term in a living, natural language, I doubt that we'll be able to get very precise about it, while still remaining faithful to the folk concept, simply because the folk concept isn't itself precise. And as for (iii), I haven't tried to spell out irrationality in nonepistemic terms. I haven't even tried to figure out which epistemic terms we have the best pretheoretic grip on, using just these as a basis for making sense of irrationality.

Regarding (iv) and (v), irrationality, on this account, is more than just a matter of not having a good reason for a belief (fear, etc.); it's a matter of having some positive reason *against* that belief. This leaves room for epistemic failings that fall short of irrationality. We could stipulate that *A* is rational iff *A* is not irrational, but this doesn't seem to be how 'rational' is standardly used, for it would allow a rational belief to be one that doesn't have anything positive going for it, epistemically speaking, so long as it doesn't embody the particular flaw marked out by 'irrational'.

I don't think that this is just a consequence of my own, potentially idiosyncratic, understanding of irrationality. Intuitively, it's not necessarily irrational to believe (so long as I"m not too confident) that the next toss of the coin will be heads—I don't, after all, have any evidence to the contrary. But it's probably not *rational* to believe it either. I'm certainly not

justified in believing it.[5] Thus, to say that a belief is irrational is to level a stronger condemnation than to claim that it's unjustified.[6]

By tacit consent among epistemologists, a belief is either justified or unjustified (even though these, of course, come in degrees). 'Rational' and 'irrational', however, seem to be different; there seems to be a middle space between the two of them, for beliefs that are neither rational nor irrational. This would at least allow for the possibility that the rational is coextensive with the justified, although it would also allow them to be distinct. Perhaps, for example, the rational is that (probably strict) subset of the justified, for which the agent possesses *reasons* in some fairly demanding sense. This would mean that all rational beliefs are justified, though not vice versa; and all irrational beliefs are unjustified, though not vice versa.

I have been arguing that 'rational', though quite familiar, is relatively unclear. It is less clear, I will argue, than the technical concept of justification. But what exactly is clear here and what isn't? Where does the unclarity of 'rational' arise?

First, off the term is polysemous; it has a domain sense and a status sense. The domain sense is fairly clear.[7] We have a pretty good pretheoretic grasp on the distinction between epistemic goodness/rationality/justification and practical goodness/rationality/justification, for example. But the status sense is quite a bit more problematic. Although we probably have a fair handle on what counts as irrational, this doesn't go very far toward giving us an understanding of the rational. The rational might just be the not-irrational, in which case a belief might be rational without yet passing epistemic muster. Or it may be that a belief is rational only if it does pass

---

[5] Bayesians often claim that any belief that violates the probability calculus is therefore irrational; this belief would qualify, on the assumption that I believe it's a fair coin, etc. But this isn't an objection to what I'm claiming here, since Bayesians are pretty clearly not relying on the standard, pretheoretical use of 'irrational'. Very subtle inconsistencies among a very large set of beliefs might render me susceptible to a Dutch book, but no one other than a Bayesian would count that as an *intuitive* instance of irrationality.

[6] This is why I claim (Lyons 2013), as Cohen notes, that we might say that victims of an evil demon are unjustified, without imputing to them any irrationality. Some of their beliefs are unjustified, but not in such a way that they ought to know better. These beliefs would then be unjustified but not irrational. This, for reasons already mentioned, is not to say that the beliefs are rational although unjustified.

[7] That is, it's fairly clear *once* we understand that Cohen's new usage makes it a synonym for 'epistemic'; the neologism derives its clarity from the clarity of concept *epistemic*.

epistemic muster, if it's epistemically proper, if it's justified. Alternatively, it might be that only some justified beliefs—e.g., those that are based on *reasons*—count as rational.

It is for these reasons that I think that 'justification', although admittedly a technical term, is a better vehicle for epistemic theorizing than 'rational'. 'Epistemically justified' does not admit of a domain/status distinction and so is not subject to any attendant confusions. It is obvious that a belief is epistemically justified just in case it passes epistemic muster (whatever exactly that means) and is unjustified otherwise. Partly because 'epistemic justification' is a technical term, there are fewer opportunities for confusion than there are with 'rational'. Of course, the fact that it's a technical term also means that we have to find some way of communicating what it is we are talking about when we use the term, but I think in the end, the result is a better understanding than we would get by using a term that's highly familiar but also highly unclear. What we get with 'rational' is an immediate feeling of understanding, but a feeling that is highly illusory. With 'epistemically justified', we start out cold but with a much better potential for genuine, rather than merely apparent, understanding.

4. Getting a Grip on Justification

Let us suppose, then, that the concept *epistemic justification* is entirely unfamiliar and that it only gets any foothold at all in the context of substantive epistemological theorizing. I think this is not true; I think that we have a fairly good antecedent understanding of the distinction between the epistemic and other domains (even if the English *word* 'epistemic' is novel), and of what would count as justified (i.e., appropriate, passing muster, etc.) within that domain. But let's suppose it anyway and see whether substantive theorizing could shed any light on the concept for anyone who needed any light shed.

One way to try to elucidate a novel term is by reference to some that are already familiar. Cohen thinks that 'rational' is an antecedently understood term that avoids most of the problems facing 'justified'. I've pointed out what I think are problems for this view. 'Evidence' and 'knowledge', however, are better candidates. I'm not an evidentialist, and I don't endorse a knowledge-first approach to epistemology, but I think that these terms can give us a good entry

into what's meant by 'epistemic justification'. Not only are they quite familiar; they are, unlike 'rational', fairly clear. So, when introducing my epistemology students to the term 'justification' for the first time, I usually do so in terms of knowledge and evidence. Justification is the salient, normative feature that distinguishes knowledge from a lucky guess. It is the normatively positive property that beliefs based on good evidence have, even if they happen to be false.

Now, Cohen argues that introducing justification in terms of knowledge and evidence in this sort of way is illegitimate, in part because it involves substantive theoretical assumptions. It assumes, for example, that a belief must be justified to be known, and this is something that some epistemologists deny. I doubt that many of these epistemologists really do deny this; they deny that an analysis of knowledge requires a separate clause for justification, perhaps, but not that safety or what have you might entail justification; or they might deny that knowledge requires justification in some highly intellectualist, reason-giving sense of 'justification' perhaps, but this doesn't mean that they deny that knowledge entails justification in the relevant sense. But setting this aside, it doesn't follow from the fact that some philosophers deny that *p*, that those who assert that *p* are making a substantive and post-theoretical assumption. The supposition that p might still be the theory-neutral position, even if there are theories that involve denying *p*. Insisting otherwise—that no position is theory-neutral if some philosopher has ever denied it—would make it far too hard for any position to count as theory-neutral. In part, this is because a position doesn't *cease* to be theory-neutral just when some philosopher comes along and tries to get tenure by staking out a contrary position. If the existence of a controversy shows that a position is non-neutral, it does so by revealing something that was always there. Thus, it's not only positions that have actually been denied that would, by this criterion, be non-neutral; it's also true of any that could be, with any reasonable degree of credibility, denied, given a sufficiently ingenious critique.

Against this, it is more plausible to hold that if credibly denying *p* requires a significant enough amount of ingenuity, then the existence of such an alternate view does nothing to detract from the theory-neutrality of *p*. This is especially plausible given that the notion of theory-neutrality we are working with here is merely the kind required for *understanding* what a debate is about. For example, you could bring me to know what you're referring to with your use of

13

'justification' by telling me that you mean that normative property entailed by knowledge, even if it happens to be controversial whether there is any normative property not identical with knowledge but entailed by knowledge—provided that I'm unaware of this controversy. In fact, I think I can get my students to understand what I mean by 'justification' by telling them that it is what needs to be added to true belief to yield knowledge. Justification is, of course, *not* the only thing that needs to be added to true belief to get knowledge, but as long as my students aren't aware of the Gettier problem, the oversimplification quite plausibly allows them to latch onto the property of justification and to link it up with my use of 'justification'. It may then be that this inaccurate description, although instrumental in securing reference to the relevant property in the course of learning the word, becomes unnecessary for further, later, acts of reference, in such a way that they are later able to think directly about the property, without relying on this description. The tricky part comes later, in getting them to understand the Gettier problem without losing their grip on justification, even though they remember how they originally got this grip. Even this seems not to be a problem, if the students are no longer relying on the original description for their access to the concept of justification. In this way, concept introduction might still be theory-neutral in the relevant sense, even if the concept is introduced by way of controversial assumptions, or indeed even by way of simplifying falsehoods!

Suppose we introduce the concept of justification in evidential terms but then proceed to explicitly deny evidentialism. Is it possible to then *retain* our grip on the concept of justification? This is, to me, a pressing question, because I think that evidence is probably the pretheoretically clearest and most familiar of the relevant concepts, and yet, I reject evidentialism. Nevertheless, I think that we can get a reasonably good handle on justification in the following, or similar, terms: *that positive epistemic property that a belief has when it is well supported by good evidence, even if that belief happens to be false*. Such a belief is epistemically appropriate; it is reasonable, rational, in some important sense proper, even though unfortunately false. We then notice that certain other beliefs, which Sellars called "IPM beliefs"—introspection, perception, and memory beliefs—have this same epistemic status, that they are also proper in that same sense, even if false. However, we might further feel that terms like 'reasonable', 'rational', and 'evidence' seem a bit out of place in connection with IPM beliefs, perhaps because these terms

14

suggest a kind of explicit ratiocination that is clearly absent in the case of IPM beliefs. This wouldn't prevent us from understanding justification by appeal to the paradigm case of beliefs based on cogent (though possibly false or misleading) evidence, even if we denied that all justification involved evidence.[8]

Now of course, many epistemologists will respond to this by insisting that IPM beliefs are in fact rational, reasonable, and based on evidence. But that's not the point. The point is that even if we denied this, we could still retain a perfectly good understanding of what positive epistemic property all these beliefs shared.[9] So even if we understand 'justification' initially in terms of evidence, this doesn't mean we're locked in to an evidentialist understanding of justification.

5. Tightening the Grip

There's a fairly serious worry that arises in connection with this move, and that is that threatens to result in an absurd proliferation of concepts. The justification theorist—at least in the nonevidentialist mode considered here—claims that many beliefs share some positive epistemic property with rational beliefs that are supported by evidence, even though these beliefs are neither rational nor supported by evidence. But surely if they share one positive epistemic property, they will share many. Put differently, if there's one good thing to be said on their behalf, there must be many. So once again the justification theorist is at a disadvantage compared to, say, the rationality theorist or the evidentialist: we already know which property the latter two are talking about, but we only have some vague idea what the former is talking about; it's some epistemically positive feature, but we're not sure which one.

_____

[8] It might work best to invoke evidence and knowledge in tandem. IPM beliefs don't have *all* the same properties as the results of cogent ratiocination. Which ones are we interested in? The ones relevant to knowledge.

[9] For what it's worth, my own linguistic intuitions balk at using 'rational' or 'reasonable' in connection with IPM beliefs, except insofar as these terms are meant to indicate absence of irrationality, as described earlier. But we're ascribing a stronger epistemic status to them by calling them "justified". Claiming that IPM beliefs are supported by evidence doesn't violate any of my linguistic intuitions, but I think there are good theoretical reasons for denying it (Lyons 2009, forthcoming), even though not the simplistic ones glossed here for the sake of illustration.

This objection, as just formulated, surely overstates the sharpness of our naive concepts of rationality and of evidence. More importantly, it neglects the possibility that our conception of justification is something that might be shaped and honed by epistemic theorizing. Even if we don't know *exactly* what's being theorized about from the very beginning, this is something that can be, and obviously often is, clarified as the theory is developed, articulated, and defended.

We might, in a Quinean spirit, go so far as to claim that the concept is straightforwardly determined by the theory and that there is no difference between deciding which theory is correct and deciding which concepts are most fruitful for understanding our world. We don't need to go so far for the present purposes, since all we're worried about is *understanding* what concept some theorist is employing. But let's push the Quinean point a bit. To take a well worn example, deciding whether whales are fish is not an entirely empirical matter, nor a matter of deciding whether whales satisfy our pretheoretical concept of *fish*. Part of what we are doing in deciding whether whales are fish is deciding which ways of carving up the world are going to be theoretically fruitful. Of course, this Quinean account couldn't begin to work if the only virtue a marine biology might have is fidelity to our folk biology. But it's clearly not the only virtue; in this case, it's probably not a virtue at all.

But it's not the only virtue in the epistemological case, either. Meshing with our pretheoretical beliefs and concepts counts for something. But other things count, too. Clarity, precision, integration with other of our best theories (both normative and empirical), demystification of otherwise puzzling primitive terms, providing a satisfying answer to skepticism, resolving or avoiding various paradoxes, explaining why various epistemic statuses are valuable, etc. leap to mind. This is just for starters. Suppose it turns out (I don't think it does) that, say, the process reliabilist and the phenomenal dogmatist are employing different concepts of justification and each is right about the nature of the concept she is analyzing; it wouldn't follow that neither theory is superior to the other, any more than it would follow that the whales-are-fish theory is just as good as the standard one.

Let's set the whales aside and consider a case closer to home. Psychology is probably more beholden to folk psychology to establish references for its theoretical terms than biology is to folk biology. For one thing, we can simply kill, preserve, and tag a type specimen in the

biological case and define the relevant taxon in reference to it; in the psychological case, there isn't any obvious way to pick out the relevant tokens, let alone the relevant types, without some kind of appeal to pretheoretical concepts. Thus, many of the terms of scientific psychology might well get their initial grip by way of their already understood role in our commonsense understanding of the mind. But this doesn't mean that the term/concept can't evolve with the further articulation of psychological theory. 'Perception', for example, is a term that is quite familiar and relatively clear outside of the works of scientific psychology. Pretheoretical perceptual concepts, however (i.e., *perception*, as well as *seeing, hearing, tasting*, etc.), tend to have both phenomenological and factive aspects that we have rightly wanted to drop from serious psychological theorizing. That is, we have discovered that there is a psychological phenomenon that is similar to the paradigm (conscious, successful) cases of perception *in all of the most theoretically important respects*, even though it is unconscious; we have discovered another that is again similar to the paradigm in all the most theoretically important respects, but is false and thus not actually seeing in the ordinary folk sense. That is, although our pre-theoretical perceptual concepts (*seeing*, *hearing*, etc.) entail both consciousness and factivity, this is no bar to developing post-theoretical versions of these concepts that omit both the factivity and the consciousness requirement.

In like manner, we might well start epistemological theorizing with pretheoretical concepts, like evidence and knowledge, that have consciousness and/or factivity requirements; but we may leave those requirements behind if we see that beliefs that are neither known nor supported by evidence still resemble the paradigms in all (or sufficiently many) of the most theoretically important respects. Whether or not this resemblance claim is true, of course, is up for debate; an evidentialist or a knowledge-firster might deny this. But I see no insurmountable obstacle to claiming that we at least understand what that debate is *about*.

Of course, philosophy isn't science, and epistemological theorizing is not the same as theorizing in scientific psychology. But if we're doing anything *real* in epistemological theorizing, it's similar enough to scientific theorizing to make my point. Psychology embraces the more inclusive concept not because of any straightforward empirical findings that showed the folk to be mistaken, but because the broader concept does better with respect to unification,

explanation, simplicity, and the like. I should hope that epistemologists value these kinds of theoretical virtues in addition to that of saving the phenomena.

6. Conceptual Pluralism

One reason why these issues matter—and a reason why it is important that we engage in this very discussion—is that epistemology (i.e., the discipline, as we know it) is concerned with a distinct if not autonomous domain of normative theorizing about proper belief formation. For this reason in addition to the ones already mentioned, our pretheoretical concepts and intuitions about cases can only be part of the whole story.

Epistemology is not even the only discipline to be concerned with epistemic normativity. The law, for example, makes frequent appeal to the concept of reasonable belief, by which is surely meant *epistemically* reasonable (or proper, or appropriate) belief. But we can't assume that the law is or ought to be employing the same concept of *the epistemically proper* as epistemologists are. It's usually legal to kill someone if you "reasonably believe" that that person poses an immediate threat to your own life. This belief doesn't have to be true, of course, but it has to be somehow epistemically proper. But the law requires a conception of the epistemically proper that is constrained in ways that the epistemologists' conception isn't. The law needs a conception that jurors can readily understand and apply. Thus it's no surprise that a standard legal gloss on 'reasonable belief' is that it is what an ordinary, careful person would believe in those circumstances. Whatever its other merits or shortcomings, it's something that an untrained juror, who has never seriously deliberated about epistemological issues before, can use to determine with a reasonable amount of confidence and consistency whether a litigant was believing reasonably. But the conceptual framework required by legal institutions isn't necessarily the conceptual framework most suited to epistemological theorizing. Even if most ordinary, careful people commit the conjunction fallacy, that wouldn't make it the right inference to make; it wouldn't thereby be epistemically proper in the sense epistemologists are interested in. Epistemology and the law *might* converge on a single normative epistemic property. But they

18

might not, and their failure to do so wouldn't necessarily indicate that either discipline had made any kind of mistake.

Does this kind of conceptual pluralism threaten to trivialize apparent disagreement among epistemologists? It's sometimes suggested that internalists and externalists are simply employing distinct epistemological concepts and thus there's no genuine disagreement.[10]

No. First, as argued above, even if internalists and externalists are working with different concepts, they might still disagree about whose concept/theory complex does a better job satisfying a common set of desiderata (clarity, integration, explaining the value of knowledge, etc.). The contrast between internalism and externalism is quite different from the contrast between epistemology and the law. Internalism and externalism are "playing the same game": there's more or less agreed on set of work that internalists and externalists want their concept(s) of epistemic normativity to do. The law isn't playing this same game at all; it has a different job description for its concepts, and it's trying to satisfy a different set of desiderata.

Second, if we claim that any disagreement implies the use of different concepts, then we trivialize *all* first-order disagreement (e.g., we couldn't disagree about what justification is, because we don't use the same concept; but we could have a second-order disagreement about what which concept better satisfies some desiderata), as per Fodor and Lepore (1992), not just in epistemology. Better to say that the fact that we have the same set of desiderata indicates that we are employing a shared concept and merely disagree about the nature of justification. And it's clear that at least some of us do share these desiderata, despite our disagreements. This isn't necessarily to claim that all epistemologists are working with the same desiderata and the same concepts. Philosophers of science and Bayesian epistemologists might be operating with a different set of desiderata and a different normative epistemological concept than are post-Gettier epistemologists steeped in the 'S knows that p' tradition.

Third, not only do we (some of us) tend to agree about what the standards are for an epistemological theory, we also agree about a great many of the particular intuitions, even when those are problems for our theories. For example, I think—contra internalists—that victims of a

---

[10] Or maybe the opposition should be between evidentialists and nonevidentialists, or responsibilists vs. nonresponsibilists, etc.

Cartesian demon (Cohen 1984) are unjustified in many of their beliefs (Lyons 2013). But I don't think this is obvious, and I don't think it's sufficient defense of this to point out that demonworlders aren't forming beliefs reliably. I think that responding to the internalist here requires making a number of moves that undercut the intuitive support for the idea that anyone who has the same experiences I do is justified in believing all the same things I am, a principle I too find intuitively appealing but which I think there are sound and non-question-begging, though nonobvious, theoretical reasons to deny. Similarly, it's not obvious that IPM beliefs are supported by evidence, but evidentialists have sophisticated and nonobvious defenses of the claim that they are. If epistemologists of different theoretical persuasion were, simply in virtue of these theoretical differences, trying to analyze distinct concepts, epistemology wouldn't be so difficult. We wouldn't feel the pull of many of these objections, and we wouldn't have to complicate and further articulate our theories to deal with them.

Thus, while it's likely that some epistemologists really are focused on distinct epistemic properties, and while this would account for at least some of the divergence in their theorizing without anybody being wrong, it is also quite likely that many epistemologists have genuine disagreements about the nature of a particular epistemic property and about which beliefs have it. So a pluralism about epistemological concepts doesn't threaten the intelligibility or reality of apparent disagreement among epistemologists.

In the end, I think that Cohen is right that 'epistemic justification' is a term of art. And he's right that this highlights a need to be clear about what it is that our epistemological theories are theories *of*. I'm less optimistic than he is about the clarity and theoretical utility of our pretheoretical concept of rationality. On the other hand, I'm more optimistic about our ability to extract a theoretically useful concept of epistemic justification from other epistemic concepts that we're more familiar with. Whoever is right here, it surely helps the larger project to at least sometimes pause our first-order investigations of the relevant normative properties and make sure our methodology is on reasonable footing.

References:

Cohen, S. (1984). Justification and truth. *Philosophical Studies* 46, 279-96.

Fodor, J. and E. Lepore (1992). *Holism: A shopper's guide*. Oxford: Blackwell.

Lyons, J. (2009). *Perception and basic beliefs*. New York: Oxford University Press.

——-. (2013). Should reliabilists be worried about demon worlds? *Philosophy and Phenomenological Research* 86, 1-40.

——-. (forthcoming). Experiential evidence? *Philosophical Studies*.