

Below is a draft of mine from 2012 circulated and archived by email, containing ideas which eventually made it into 'An Epistemic Non-Consequentialism'.

See <http://www.kurtlsylvan.com/b-sides-demos-false-starts.html> for other old drafts.

The Foundations of Epistemic Kantianism

“Respect for the truth is an acquired taste” —Mark Van Doren

Preamble

According to epistemic Kantianism, what it is for truth (in belief) to be a properly epistemic value is for it to be involved in certain standards that we ought epistemically *to respect*. So far I've given no complete account of what respect for the truth involves: I've just relied on a more or less intuitive notion brought out by contrasts and examples. Yet I also claim that exhibiting respect for the truth will require complying with requirements of epistemic rationality—a claim I'll use to solve the problems of deontic significance for epistemic rationality I earlier uncovered, partly *via* analogies with the practical literature. One may complain that this solution cannot rest on some briefly glossed notion. Agreed. So I give an account in §1 and §2. Given the account, I note at the end of §2 why exhibiting respect for the truth would require complying with requirements of epistemic rationality. And I explain in §3 why giving an account of respect involves in other terms does not make the Kantian ideology superfluous. In doing so, I clarify what kind of theory epistemic Kantianism is.

1. Recognition Respect in General

1.0. *The target*

There are many uses of 'respect' talk. Many are irrelevant for me. So far my main dividing line has been one from Darwall. He distinguished between *appraisal respect*, which involves a kind of *esteem* most often directed at people, and *recognition respect*, which involves a *suitable constraining of choice and deliberation* that has a source of authority as its object—esp. the authority of rules and certain values. Only the second concerns me. Darwall glosses it thus:

There is a kind of respect which can have any of a number of different sorts of things as its object and which consists, most generally, in a disposition to weigh appropriately in one's deliberations some feature of the thing in question and to act accordingly. The law, someone's feelings, and social institutions with their positions and roles are examples of things which can be the object of this sort of respect. Since this kind of respect consists in giving appropriate consideration or recognition to some feature of its object in deliberating about what to do, I shall call it *recognition respect*.¹

I think there is a familiar phenomenon to which Darwall gestures, though his characterization is in some important ways misleading, as I will note soon enough. What matters off the bat is that the *phenomenon* is not some philosopher's invention, though Darwall's *term* for it is. When we speak of someone's

exhibiting no respect for some rule (e.g., a law),

for instance, we have this idea in mind. And we often talk this way: just keep your ears peeled! Moreover, it is clear enough prior to any philosophical theory-building that the problem with such a person is not a lack of “respect” in the sense at issue in claims like

1 Darwall (1977: 37)

“I have great respect for Philippa Foot as a philosopher”

When someone exhibits respect for the speed limit by deliberately slowing down, it is not as if she esteems the speed limit or anything like that. She may wish in her heart of hearts that a different limit were set and find this one dumb, or even feel nothing at all about it.

We can also distinguish the relevant use of 'respect' from a different use that sometimes arises when we talk about rules. Sometimes when we say “She is respecting the rules”, we just mean “She is conforming to the rules”. This sense of 'respect' usually arises when one doesn't put 'for' after 'respect'. This is not what concerns me. Indeed, I will claim that exhibiting recognition respect doesn't require conformity, even in the main. This should not be surprising: if exhibiting recognition respect for a rule is a matter of *appropriate deliberative constraint*, as Darwall says, then whether one counts as exhibiting respect for a rule will turn on the *quality of one's deliberative efforts* rather than on their *products*. If you try by all means available not to break the speed limit but your car is simply malfunctioning and so going too fast, *you* exhibit no lack of respect for the speed limit. Yet you still break the speed limit.

There is a relatively clear, familiar idea here. And Darwall's gloss seems partly apt. But it is worth saying more than Darwall does about its conceptual role, and what, in nonnormative terms, it involves. The second task is crucial. Note that Darwall explains our target *in normative terms*. He says that recognition respect for X involves “weigh[ing X] *appropriately* in one's deliberations...” This won't do here. I say that the response fundamentally called for by the epistemic value of truth is recognition respect. Yet this may seem a completely uninformative claim if we can only gloss recognition respect in normative terms. For if having recognition respect *just is* appropriately weighing certain considerations in one's deliberations, claims of the form “We ought to give recognition respect to X” may appear to be, in Parfit's words, *concealed tautologies*. For *any* X, we ought to appropriately weigh X in our deliberations in some sense, since “appropriately weighing X in our deliberations” could simply mean “weighing X in our deliberations *as we ought*”. That may even involve giving *little weight* to X in our thinking. We can avoid concealed tautologies, but it takes work.

Before turning to this, it is worth noting that Darwall partly misleads us about our target. Suppose that, in some particular case, I give a legal rule constraining authority over my practical deliberation in this strong way: I *deliberately comply* with this rule because I *realize* that it is a good rule, and *see* that the rule requires doing what I am about to do. We cannot complain about my behavior. I exhibit no lack of respect. Yet, crucially, my behavior may manifest no related disposition. It could be that I just “saw the light” for only a moment, and so made great efforts to block, however briefly, a disposition I have to do precisely the reverse. While it may be true in some *generic* sense that I “lack respect for this rule”, I exhibited no failure of respect in this case. Indeed, in this case, I exhibited recognition respect for this rule in an overwhelmingly clear way, by virtue of my witting, effortful compliance. Of course, this is not to say that witting, effortful compliance should serve as the paradigm from which theorizing should directly generalize. It is just to say that witting, deliberate compliance should be *one uncontroversial example* of recognition respect.

In what follows, I will only be interested in what it takes to exhibit recognition respect for a rule or value in a particular case. For me at least, this is the more fundamental phenomenon. Indeed, I think we could analyze the more generic sense of “having respect for a rule” in terms of this notion. We could see it as a disposition to exhibit recognition respect in particular cases. For this reason, it is misleading when Darwall speaks of recognition respect

as a “*disposition* to weigh appropriately in one's deliberations”. He could be *stipulating* that this is how he intends to use “recognition respect”. But if we are interested in the familiar phenomena, we would make a mistake to think that the disposition is the thing to discuss.

Indeed, there is a clear further reason why this would be a mistake. I am interested in a kind of respect that can serve as a *fitting response to certain things of value (of some kind, perhaps domain-relative)*. Darwall is too, particularly in more recent work. And I don't think only *we* should be interested. The topic is generally interesting, since it is plausible that there are *some* values to which a fitting response is recognition respect—whatever we say about *fundamentally* fitting responses, or all values. Yet if recognition respect were a *disposition*, this entire line of inquiry would make no sense. Having a certain disposition is not a *response* that a value can call for. At most, a value can call for a *manifestation* of a disposition. And even that is not what should interest us, given the example I just discussed earlier of witting, effortful compliance in a rare “moment of clarity” that blocks one's disposition to do otherwise. That is a clear case. Yet such compliance manifests no disposition (to try) to follow the rule.

In any case, regardless of whether we want to analyze the broader notion of simply *having respect* in terms of *exhibiting respect* or (somehow) the other way around, we can distinguish these notions. By stipulation, I will be discussing only the latter notion, not the former. Although I will typically use the explicit phrase “exhibit recognition respect”, I will sometimes drop “exhibit”, or use other formulations for just for variety's sake. All of these are intended, however, to mean the same thing meant by the more explicit phrase.

1.1. *The normative conceptual role of recognition respect*

While I will ultimately characterize what makes for recognition respect in nonnormative terms, showing how it links up with certain normative phenomena is not useless. Indeed, it helps to *avoid* the charge of concealing a tautology. The worry was that if respecting X *just is* appropriately weighing X in one's deliberations, then any claim that we ought to respect X would be trivial. But this is false if we understand the *kind* of appropriateness at issue. So I want first to note how recognition respect is necessarily linked with (albeit not definitionally) a distinctive kind of appropriateness, and stress that in *some* domains there are not always strong, non-optional, genuinely normative reasons to achieve this kind of appropriateness. Because there are not such reasons in some domains, it is very nontrivial to claim in the epistemic domain that what we always should “do” is respect truth.

The key connection is between failing to exhibit or exhibiting recognition respect for certain rules or values and *exhibiting carelessness* or *avoiding carelessness*. It can be expressed like so:

(CARELESSNESS LINK) Failing to exhibit recognition respect for X in a given case always makes one careless with respect to X's value or authority to some degree in that case, and to the extent that one is not at all careless with respect to X's value or authority in a given case, one cannot fail to exhibit recognition respect for X in that case.

“Carelessness” is to be understood here as a thick evaluative term. Moreover, although the suffix '-less' can make it look like the term should not be gradable, I take it that it is intuitively gradable: one person can be more or less careless than another. Furthermore, although we can talk about carefulness and carelessness as dispositions, we can also talk about particular examples of careless or careful behavior—examples that may be highly atypical for the people we are discussing. I have the particular-case notion in mind. Finally,

avoiding carelessness with respect to a rule does *not* require conforming to it. Indeed, it may, as I'll note in a moment, sometimes require precisely the reverse.

The CARELESSNESS LINK is, I think, a platitude. To see why, we can again consider the case of exhibiting recognition respect for legal rules. There are many cases where it would be natural to say that someone fails to exhibit respect for these rules, even if she *happens* to conform to them. In these cases, it will *also* be natural to say that she has been careless to some degree. Suppose the law requires driving no faster than 25 mph in the middle of Newark, DE. If one *isn't even trying* not to do what would constitute breaking this law, one is intuitively careless. Of course, unless one *actually breaks* the limit, one may not be legally punishable. But one could still be *fairly reproached* for not constraining one's decisions about how to drive *via* the rule—if, say, one drives 24 mph *only because* that was David Bowie's age when he wrote the song one listens to in Newark, DE. Conversely, insofar as one is careful with respect to this rule, one could not have failed to exhibit recognition respect.

This brings out why it is controversial to say that the response we ought always to have to certain values or rules is recognition respect. For one thing, we in effect saw in Chapter 3 and elsewhere that there are cases where we *really ought* to be careless in the relevant sense. This is, for instance, true when it ultimately appears—deeply misleadingly, so that the appearance is not truly good evidence—that A-ing would harm someone, when A-ing in fact is stably objectively likely to save that person's life and help her in other ways. The response more stringently demanded by the value of a life good for the person who lives it is *real protection*. If protecting some such life requires disrespecting the person's dignity, the con-reasons of disrespect are defeated. Indeed, the person could thank you for ignoring them—even silently cheer you on, if she was aware of the objective chances—and you could in retrospect be glad that you did it, and say plausibly that you “did the right thing”.

Moreover, recognition respect is not a response at all demanded as such by some values or rules. Perhaps this is true of speed limits. It hardly matters as such why one conforms. All that matters is that one does reliably conform. While passengers can still chide you when you explain your choice of 24 mph, this chiding has little deontic significance as such. Indeed, it may have no meaning at all if you have a stable disposition to listen only to Bowie songs in places whose speed limits match his age upon authoring them. No risk would then be imposed on your passengers or you. The chance that you'd conform if your disposition were unblocked by contingencies is 1. Still, you would in the relevant sense be clearly careless: *your thinking about how fast to drive wouldn't be governed at all by the relevant considerations*.

That is one of the reasons why I claimed earlier that I think that there is a substantive disanalogy between epistemic cases (as I see them) and some practical cases. Epistemic Kantianism implies that all that matters from a deontic point of view in the properly epistemic sense is respecting the truth. Producing true beliefs—while good *and epistemic*—does not matter *per se* in a *properly epistemic* sense. Something very different is true in many practical cases. If someone is at risk of death, what her life's value calls for is *saving*, if continued life would really be good for her. If you can only save her while also exhibiting substantial recognition *disrespect* for her dignity, you should still *just save her*. Plausibly, the value of continued good life calls more fundamentally for *real protection*. Some may disagree. But the point here is that this would be controversial—and, I would say, counterintuitive.

If so, it is *substantive* to say of any other value that what we always ought to do with respect to it is give it recognition respect. That may not be the response it fundamentally calls for.

If not, there will be cases in which we *really ought* to be careless with respect to it. So although I will want to characterize what recognition respect involves in nonnormative terms, it is not true that Darwall's characterization leads only to concealed tautologies. Everything turns on how we should understand “appropriately” in his gloss.

1.2. *What recognition respect cannot be*

Besides answering the charge of concealing a tautology, the CARELESSNESS LINK gives us a constraint on any account in nonnormative terms. Any such account must be compatible with the fact that failing to exhibit recognition respect makes for carelessness, and that fully avoiding carelessness makes for recognition respect. This rules out several theories that might have seemed natural, and helps us narrow in on the correct one.

1.2.1. *Naïve following theories*

One class of theories are theories that try to equate having recognition respect for rules with simply following these rules. Conformity with a rule in a given case seems necessary though insufficient for following it. You follow a rule only if you are in accord with it, and being in accord with a rule seems to demand conformity. If you are in fact driving 150 mph, you are plausibly not following the rule to drive no faster than 60 mph. You are not going to get away with a charge of failing to follow the rule by saying, “Look, I was following it. I was trying with all my might to drive no faster than 60. But my foot got paralyzed. I couldn't get it off the accelerator.” Even if your foot did get unexpectedly paralyzed and you would otherwise have conformed, you didn't follow the rule. Maybe there is a sense of “follow” that doesn't require conformity. But the dominant sense seems conformity-implying.

If the dominant sense of “follow” is conformity-implying, exhibiting recognition respect for a rule cannot imply following it. When ultimate appearances are sufficiently misleading, it may be that the only way to conform to the rule is to be careless. Indeed, avoiding carelessness may in some cases guarantee that one will *not* conform to the rule.

So while it *might* be true that following a rule implies exhibiting recognition respect for it, it cannot be true that exhibiting recognition respect for a rule implies following it. One can be immune from any charge of carelessness, and indeed be perfectly careful, and yet manage not to follow the rule. Indeed, if appearances are sufficiently misleading, one can systematically fail to follow a rule without being open to a charge of carelessness, and indeed while being praiseworthy for extreme carefulness.

Even if the dominant sense of “follow” were not conformity-implying, there is another respect in which the simple following account cannot be right: it is too weak. What following involves beyond conformity (if it involves that) is simply that one's behavior be explicable by one's having the rule or the satisfaction of its antecedent as a rationale for acting. This could involve suitable beliefs about the rule's applicability generating one's behavior. Or it could involve simply behaving in a certain way on the basis of occurrent appearances as of the satisfaction of the antecedent of the rule. The trouble is that, either way, this allows for carelessness. Careless rule following strikes me as a clear possibility.

Suppose that the rule calls for A-ing in one's circumstances, and that it occurrently appears that that B-ing would lead to A-ing in these circumstances. But also suppose that one is in a position to know that this occurrent appearance is misleading—say, by triggering a memory

one negligently fails to trigger, but that could be easily triggered with minimal reflection. One might believe that B-ing will lead to A-ing on this basis of the occurrent appearance. Knowing that the rules require A-ing, one might then intend to B, and B. Moreover, the unlikely can happen: B-ing does, against the chances, lead to A-ing.

If all this is true, it is plausible that one is following the rule. After all, it occurrently appears that the antecedent of the rule is satisfied, one acts for that reason, and conforms. Yet one follows it in a careless way, since one could *so easily see* that success was objectively unlikely. If so, following a rule cannot be sufficient for exhibiting recognition respect for it.

So following a rule is neither necessary nor sufficient for exhibiting recognition respect.

1.2.2. *Accepting and acting theories*

Another class of theories ruled out by the CARELESSNESS LINK is the class of theories that equate exhibiting recognition respect for a rule with accepting the rule and acting on one's acceptance. While these theories are better in not requiring conformity, they suffer from a similar problem to the second one just posed for the naïve following theory.

This is easy to see. One can accept a rule and act on one's acceptance while having fairly clearly but not obviously false beliefs about what it takes to conform. "Fairly clearly false" can mean "would be obviously false if certain stored appearances (e.g., stored apparent memories) were made occurrent." In good cases, this will amount to being in a position to know that one's beliefs are false but thinking otherwise, due to occurrent appearances to the contrary. Call this kind of acceptance "weaselly acceptance". When a person is guilty of weaselly acceptance—either by failing to take time to reflect, or by blocking certain appearances from becoming occurrent through subpersonal self-deception, or by negligence of another kind—she will be careless to some degree with respect to the rule, even in acting on her acceptance. This carelessness implies that she fails to exhibit recognition respect. So exhibiting recognition respect for a rule cannot be equated with accepting and acting on it.

Replies are imaginable. One could try to claim that truly accepting a rule is incompatible with failing to respond to the fairly clear falsity of one's beliefs about what it requires. While I do not think this is true, there is something behind this thought. A weaker thought is plausible. Suppose one *always* sticks to fairly clearly false beliefs about what a rule requires. Perhaps one reliably refuses to reflect on whether one's beliefs about what the rule requires in any given case, and dogmatically sticks to whatever is most convenient. If one had this tendency, it would be implausible to attribute genuine acceptance of the rule to one. One's "weaselly" dispositions seem incompatible with willingness to be governed by the rule. And if one really isn't willing to be governed by a rule, one plausibly doesn't *really* accept it.

Plausible enough. But this weaker idea cannot help the defender of the view that exhibiting recognition respect for a rule just is accepting and acting on that rule. For this weaker idea is compatible with one's failing to respond in *some* cases by sticking to a fairly clearly false belief through negligence. It is not as if subtly "weaseling" out of conformity while still accepting the rule is *impossible*—provided that the negligence that blocks the relevant dispositional appearances from becoming occurrent is not the result of a deliberate, knowing act (i.e., is not *recklessness*). Yet any case of acting on a "weaselly" acceptance suffices for some degree of carelessness. This is incompatible with one's exhibiting recognition respect. So the revised account of exhibiting recognition respect is too weak.

There seems to be a dilemma for accepting and acting theories. The defender of the view can either understand acceptance in a sense that makes weaselly acceptance impossible or in a sense that makes it possible. If the former, then the theory presupposes a false view of acceptance. While *systematic* weaseling may be incompatible with genuine acceptance, surely *some* weaselly acceptance is possible. If the latter, then the theory fails to predict that cases of weaselly acceptance are cases where someone does not fully exhibit recognition respect.

1.2.3. *Sophisticated following theories*

This point also undermines a less direct attempt to account for recognition respect in terms of following a rule. One might think that while exhibiting recognition respect for a rule does not require following it, it at least requires *revealing a disposition to follow*. “Revealing” must obviously mean something weaker than “manifesting”, lest the view be *even worse* than the simpler following view, by requiring *more* than it. We can make out what “revealing” means with examples of *trying*. Suppose you decided on a policy of eating a healthy breakfast every day, but that your just-prepared food gets knocked into the garbage by cats every time. Because you try, you *reveal* a disposition to follow the policy, though you don't *manifest* one in the strict sense of displaying the canonical manifestation (*viz.*, following it).²

While there may be plausibility in the idea that revealing a disposition to follow some rule is necessary for exhibiting recognition respect in a given case, I do not think it is sufficient. This is for the simple reason that what counts as a revelation of a disposition admits of enough flexibility to let in at least a bit of weaselly behavior in the sense introduced in the last section. At least it admits of enough flexibility to allow *a single case* of weaseling. It can't be that when one was careless because one failed to see certain facts by sticking to a certain belief about a rule's conditions of application by unnoticed dogmatism, one *didn't even try* to conform. But if one really tried, that counts as revealing a disposition to conform to the rule—particularly if carelessness of this kind is *made rare* but *not ruled out* by the disposition.

Imagine an otherwise scrupulous vegan who on this unusual (say, distracting) occasion didn't check the ingredients of something because she “was confident that it didn't contain animal products”. Suppose she did ask herself: “Is this vegan?” She then hastily answered: “Yes. I recall that it is,” where the “I recall” turned out to signal a product of subtle, unintentional self-deception that could have been overturned by a bit of further reflection or a look at the ingredients. She was, let's stipulate, in a position to know otherwise. Her confidence that she remembered was a side-effect of two factors: the food tastes like similar vegan foods, and also, well, tastes *really good*. Did this person fail to reveal a disposition to follow vegan rules? I doubt it. She really has such a disposition. She consulted her memory, and then tried to follow her rule *because of* her disposition. It's just that her “memory” only delivered some hogwash that could have been overturned without difficulty. Yet she was to some degree careless. So she failed to exhibit full recognition respect for her principles. Hence, revealing a disposition to follow cannot be sufficient for exhibiting recognition respect.

2 Revealing a disposition is not, I'll stress, just a matter of displaying a property that *indicates* possession of the disposition. The source of the revelation must be the very same categorical property that grounds the disposition. The idea is that when a thing or person has a disposition to M, the categorical property that grounds that disposition will also dispose that thing or person to behave in certain distinctive ways in bad conditions—either bad with respect to the stimulus, or with respect to the background conditions. These ways of behaving must be “related” to the canonical manifestation in some rough and ready way, but will not amount to a strict manifestation of the disposition. Note that *mere indication* of a disposition could amount to manifesting a disposition conferred by some *different* categorical property that is, as a point of regular fact, instantiated in things that instantiate the categorical property that grounds the disposition.

There is an even simpler route to this conclusion. What we *in effect* saw in §1.2.2 was that all it is to accept and act on one's acceptance of a rule is to reveal a disposition to follow it. If so, then the fact that accepting and acting on a rule lets in some cases of weaselly behavior shows that any account of exhibiting recognition respect in terms of revealing such a disposition is bound to fail. So we can more easily transfer the final objection from §1.2.2.

The other problem with these theories is that they are too strong. I agree that there is some intuitive pull to the thought that exhibiting recognition respect for a rule in a given case requires revealing a disposition to follow the rule. But this pull is misleading. Take a case where someone, for the first time ever, realizes painfully that she has a disposition *not* to follow some rule. With a great exhibit of effort, she manages to prevent this disposition from manifesting, and manages to follow the rule deliberately. Was this person relevantly careless? No. Indeed, she's praiseworthy here for the care taken to comply. Yet the sophisticated theory implies otherwise, since she revealed no *disposition* to comply.

Any attempt to cash out exhibits of recognition respect in similar dispositional terms cannot work for this reason. Although not all cases of exhibiting recognition respect must involve deliberate, witting compliance with the kind of perspective we just considered, surely *some* of these cases exhibit recognition respect. Even if one is disposed to violate a rule, one can block one's disposition in certain moments of clarity. Examples like this can be paradigms of praiseworthiness and, indeed, carefulness. In these cases, it is not true that one acquires a better disposition. Clarity may fade immediately after one's success, and one may return to violation. That does not imply that one never succeeded—i.e., that one was not careful, and did not exhibit recognition respect at all. For this reason, all dispositional accounts that I can imagine must fail by ruling out some clear examples.

As I in effect flagged much earlier, there is an explanation of why such accounts are tempting. There is a difference between the generic idea of

having recognition respect for a rule,

and the local, particular-case idea of

exhibiting recognition respect for a rule in acting or forming attitudes in some way

The first has a dispositional flavor. The second can *sound* dispositional, especially given 'exhibits', which sounds like 'manifests'. Yet we saw at the outset that this cannot be right, by considering the case of witting, deliberate compliance in a way dispositionally atypical for the agent: behaving because sees in a moment of clarity that a correct rule requires it, and with great effort to overcome, however momentarily, a disposition to the contrary. Even a person who is generally disposed to do the reverse of what the rule requires might have these rare moments of clarity where she sees the good, and knowingly complies with the right rules. Here the person exhibits recognition respect. Yet this is not backed by any disposition to follow. If so, requiring that the person reveal a disposition to follow seems fundamentally misguided: it excludes rare moments of witting, deliberate compliance.

1.2.4. *Plus-reflection views*

Cases of dispositionally atypical, witting and effortful compliance rule out other views. A problem with every view considered so far is predicting that some cases of carelessness are

actually cases of recognition respect. The cases have been similar: in each one considered earlier, occurrent appearances suggest one thing and the person settles hastily for what they suggest, but reflection on available facts would easily defeat these suggestions. One might have wanted to handle these cases by adding some reflective constraint. For example:

AA+R: You exhibit recognition respect for R by A-ing iff you accept R and A after seeing that it is apparent upon reflection that A-ing allow you to conform to R, where *that A-ing allow you to conform to R* is also your rationale.

Yet AA+R is false, since such reflection isn't necessary if one really knows *in the first instance* that A-ing in C would constitute conformity with R. This is true in some cases of deliberate, sighted compliance. I bet it is true in lots of other cases of genuine recognition respect. Reflection is unnecessary, since one can often know “off the bat” that one is doing the right thing. Or one can at least know off the bat what would be just as serviceable—say, that one's act would be precisely of a type that (as a *side-comment*) facilitates conformity to R. These cases are more realistic, since people may rarely explicitly represent the rules followed.

AA+R overintellectualizes in other bad ways. It requires one to have the capacity to think the thought that A-ing in C would constitute conformity to R. In many cases, this may not be unreasonable. How else is one going to comply with various legal rules, if not by conceptualizing them? But in other cases of ultimate interest, it won't do. Someone can exhibit sufficient recognition respect for various logical rules without conceptualizing them as such. This isn't to say that having recognition respect for certain rules is not going to require *any* concept application. To engage in any cogent inference whatever, one plausibly needs the concept expressed by 'so', where that picks out some consequence relation. Still, one can have that concept without conceptualizing the rules in any detail.

This is a problem for two dimensions of AA+R. Acceptance plausibly involves concept application. One accepts X or that P only if one has a concept of X or of the constituent particulars and properties of the fact that P. Overintellectualization in AA+R would then stem from the acceptance bit of AA+R. But it also stems from the “seeing that” bit. While I do think responses are possible here, it would be nice to directly avoid having to give them.

1.3. *Towards a better view*

The foregoing subsections suggest five desiderata on an improved theory:

- (1) It should rule out all cases of carelessness, including “weaselly” cases.
- (2) It should not require manifestation or revelation of a disposition to follow.
- (3) It should rule in cases of dispositionally atypical, effortful witting compliance.
- (4) It should not require conformity.
- (5) It should avoid overintellectualization.

The view in the last section did make one stab in a better direction—just way too hard a stab! In the cases of carelessness considered before, the problem certainly appeared to have *something* to do with the *insensitivity* of the agents to certain relevant facts or appearances that

could have been uncovered on reflection. Yet the failure of AA+R shows that this sensitivity cannot generally *require* any sophisticated and time-consuming reflection, though such reflection could in *some* cases actually be operative. So the crucial question to ask is:

(Q) What does this sensitivity require, if not such reflection?

In answering (Q), it is tempting to violate desideratum (2). One thinks: if an agent just had a competence to A only on the basis of appearances that *in fact* would not be outweighed by further considerations reflection would shore up, she could exhibit respect for R without needing actually to reflect. Yet we cannot require this much in general. If an agent has a “moment of clarity” and manages, just for now, to block some *incompetence* of hers that inclines her to contrary rule-violating behavior, she can manage to exhibit recognition respect without manifesting or revealing any disposition to (try to) follow the relevant rule.

We need a theory that lets both examples count—cases of manifesting a sensitivity, and cases of dispositionally rare, witting and effortful compliance in moments of painful “moral” clarity. Thankfully, there is a theory of this kind that secures the other desiderata. To state it neatly, a couple of pieces of terminology are worth introducing. Let a

stable occurrent appearance for S be an occurrent appearance of S's whose content would appear not to be made significantly less objectively likely by contents of any other occurrent appearances, stored appearances (e.g., apparent memories), or dispositional appearances (e.g., intuitions that would strike S upon addressing a question) in S.³

While attributing such appearances to someone may require her to possess the concept of objective probability, I do not think this is overintellectualization. It is hard to imagine beings capable of much theoretical or practical reasoning who lack this concept. This is not to say that they must have the capacity to use the *phrase* 'objective probability', or to *articulate* their thoughts about it—just that they can *regard* things as more or less chancy.

It is easiest to get an intuitive feel of the notion of a stable occurrent appearance by thinking about *good cases*, where pathways to knowledge work. In them, there are many things we are *in a position to know*: we have *the capacity to exploit various pathways to knowledge* available to us with no further changes in the course of our experience. We could consider our memories more carefully, consider the consequences of other things known noninferentially, scrutinize the (veridical) experiences we are having, and so on. In good cases, a stable occurrent appearance is simply an occurrent appearance with a content that is not made significantly less likely by anything else we are in a position to know, and knowably so. In bad cases, things *feel* exactly the same. But the pathways don't actually help us get knowledge.

So we can think of a stable occurrent appearance as an occurrent appearance whose content would not be made significantly less probable by anything that we are *as if* or *virtually* in a position to know—i.e., that a nonfactive mental duplicate of ours *in a good case* is in a position to know. Indeed, I suspect that we could do without the complicated appeal to such notions as dispositional and stored appearances with the kind of Williamsonian move made by Bird (2007). His view was in effect that doxastic justification is virtual (“as if”) knowledge, and that you're propositionally justified in P iff one of your non-factive twins in a good case is in a position to know P. While I won't pursue them here, I'll indicate that I now am very much inclined to make similar moves.

3 The 'would appear' is simply meant to indicate a further dispositional appearance in S.

Given this piece of terminology, we can consider the following *Improved Theory*:

Where a case C is *A-encouraging* iff A-ing is objectively likely to approximate compliance with R in C,

S exhibits recognition respect for R by A-ing iff

(i) S As *because* it stably occurrently appears to S, of some possible case that would be A-encouraging, that such a case obtains,

and

(ii) S As in order to do precisely what approximates compliance with R,

where

(iii) this is the only motive needed to explain S's A-ing.⁴

The idea behind this theory is that recognition respect is a matter of attempted compliance *via* reliance on a stable occurrent appearance as of a good opportunity for compliance.

Now, two points of clarification are in order to forestall unnecessary concerns.

Firstly, the 'because' in the Improved Theory is the (factive) 'because' of sense-making explanation or, as it is more often called, "normative explanation".⁵ The idea is that we can use precisely the existence of a stable occurrent appearance as of some A-encouraging case to make the best sense of S's decision. It is a familiar fact that this 'because' is inherently open-ended about precisely what causal mechanisms are at work. It depends on the case. There is no more general story to be had. Maybe the person realizes after reflecting that A-ing in this case would approximate compliance with some correct rule, and she manages to A *via* great effort with this as an explicit rationale. This could be a case of witting, effortful compliance, where the agent manages to block a disposition she has to reliably violate the rule out of a brief period of clear-headedness. In other cases, this person may simply have some automatic competence, and it may be revealed by her A-ing in response to the stable occurrent appearance, and only appearances of that kind. This competence could be innate, or it could be acquired in a subtle way, without the person's seeing it as such.

It is the inherent open-endedness of the 'because' of sense-making explanation with respect to modes of causal implementation that enables the Improved Theory to capture desiderata (2) and (3). And ignoring this concept and its inherent open-endedness will lead to failure. It will lead to failure by creating an artificially forced choice between (a) requiring that the person does what would appear to constitute following the rule by virtue of some unreflective competence or disposition, and (b) requiring that the person does what would appear to constitute following the rule by virtue of sighted choice that may not be backed by any stable competence. The fact is that both of these examples are good examples of exhibiting recognition respect. They can be brought under one umbrella only with the (factive) 'because' of sense-making explanation in this way.

Now for the second point of clarification, which may already be obvious. In clause (i), I am *not* assuming that the appearance or S herself must explicitly represent the situation *as an opportunity for facilitating compliance with R by A-ing*. The situation is just represented to be one

⁴ I leave room for other motives in S. It is just that they are not needed: S would still A without them.

⁵ *Nota bene*: in agreeing that this 'because' is factive, I am *not* agreeing that locutions like "S's reason for A-ing was that P" are factive. I still think that those locutions are non-factive. Indeed, for that reason I deny that if S's reason for A-ing is that P, S can A *because* P. These are simply different phenomena—one is weaker.

precisely of a type *where S could approximate compliance with R by A-ing*, with that used as a side-comment. Now, S *does* need to be reacting to precisely this *property* of the apparent case. But this can happen in different ways—e.g., *via* clear knowledge that the case is A-encouraging, or simply *via* some competent sensitivity to the apparent feature of being a case of the type.

Suppose for illustration that S is driving 70 mph and it now occurrently appears to S that she is driving that fast, given a look at her speedometer and outside. The rules call for driving between 40 and 60 mph, inclusive. This is a case where slowing down by at least 10 mph approximates (indeed, guarantees) compliance. Suppose S then decides to slow down by 20 mph, given the look. This gets her close to exhibiting recognition respect for the speed limit, since the appearance here is presumably stable—S's speedometer isn't malfunctioning, and S isn't in a position to know anything that would make this likely. Yet S may never think to herself: “This is a case where in order to approximate conformity with the speed limit I must slow down.” Clause (i) is deliberately phrased to accommodate this kind of example.

More is required for a display of recognition respect, and that is the point of clauses (ii) and (iii). If S is slowing down mainly in order to reach a speed that coincides with Bowie's age upon writing his last good song, S is not exhibiting recognition respect for the speed limit. S must be slowing down *in order to reach the speed limit*, and that must be all we *need* to explain S's behavior—though she may have the coincidence with Bowie's age in mind, and rejoice at it.

Now we hit a crucial extension of the second clarification. To explain someone's act as being done *in order to do what approximates compliance with a rule*, it is *not* necessary to represent that person even as having the concept of that rule, or as representing it as such. In this case, it is plausible that the person has the relevant concepts. But if we considered other cases—cases where S changed her wording *in order to avoid violating some syntactical rule*—this would be unnecessary. I think this is a common way of using 'in order to'. The Improved Theory does not presuppose otherwise. This is one way it avoids overintellectualization.

We can now see how the Improved Theory accommodates the other desiderata. Given the correct background views about the 'because' of sense-making explanation and what is meant by 'in order to', it avoids overintellectualization and captures desideratum (5). It does not require one to recognize the stability of the appearance by engaging in any reflection. If one *simply knew* the content of the appearance, there would be no need to engage in such reflection. The Improved Theory does not rule this out as part of a legitimate way of exhibiting recognition respect for a rule. Moreover, it clearly does not require conformity: even stable appearances can be misleading. So it captures (4). Finally, because the cases of carelessness we considered were ones where someone acted on an unstable appearance, and our theory requires acting on a stable appearance, it seems to capture desideratum (1).

Indeed, it rules out a further kind of carelessness. Now, explaining someone's behavior by adverting to the fact that it would constitute compliance with a rule doesn't require seeing that person as someone engaged in “rule worship”. Often we care about rules because these rules are independently good, and worth imposing on our behavior or attitude-formation. It may be this goodness at which we *ultimately* aim. Nevertheless, one kind of carelessness we do want to rule out is doing what would constitute compliance with a rule primarily for some completely different, irrelevant reason. We saw this when we discussed the case of driving just below the speed limit just because that was David Bowie's age at some time that interests you. Our Improved Theory rules out this kind of carelessness with clause (ii).

2. Respect for the Truth in Particular

Now that we have a better account of what recognition respect for rules involves in general, it is time to use this broader idea to explain “respect for the truth” as it figures in epistemic Kantianism. Since I gave an account of recognition respect in general for *rules*, I will want to reduce respect for the truth to recognition respect for certain rules. These rules may, I’ll stress, not always or often be our explicit “decision guides” in theoretical reasoning. But they are ones whose correctness is needed to explain the serviceability of any explicit decision guides we use. And virtuous epistemic subjects will use whatever rules they use because, in some deeper sense, they apparently approximate compliance with the fundamental ones. Figuring out just what that sense of ‘because’ involves would get us into hard questions about the nature of belief and the sense in which the rules I’ll discuss are *constitutive*. I set aside those questions here, though I’d hope to address them at some point.

2.1. *Fundamental rules for which recognition respect is demanded*

The only rule I’ve discussed much before is the following fundamental one:

(T-Rule) Believe P only if P.

There are many states that would constitute compliance with the T-Rule: (i) believing P truly, (ii) suspending on P when P is false, (iii) suspending on P when P is true, (iv) disbelieving when P is false. We certainly need to add another rule:

(F-Rule) Disbelieve P only if \sim P.

We need to add this because it is not always equally permissible to try to comply with the T-Rule by either suspending or disbelieving. Yet the mere claim that we ought to try to comply with the T-Rule doesn’t have this implication. It only has the implication that we should *not believe* P if \sim P. And there are two ways to not believe P—suspending and disbelieving. The T-Rule by itself says nothing about how to choose between these ways.

Do we need more rules yet?

So far I’ve simply been talking about the familiar trichotomy of doxastic attitudes of traditional epistemology. If there were a clear reductive story about how credences relate to these states, we could try to subsume these rules under more fine-grained rules. If, for instance, flat-out belief were credence 1 and flat-out disbelief were credence 0, we could subsume the two rules above under some relative of the Principal Principle such as:

(Coordination Rule) Have credence X in P only if the chance that P is X (and you lack “inadmissible” information).

Since I do not know how to reduce belief and disbelief to credence, I will stick with the familiar trichotomy, though I will note later how useful it would be to have a reductive story.

But what of the third state beyond belief and disbelief? Do we need some rule for it?

Yes, but some observations are in order first. While I just used the term ‘suspension’ above, it is crucial to distinguish at least three different phenomena that one might have in mind in

using this term. One phenomenon is the simple absence of belief and disbelief. Another phenomenon is withholding belief and disbelief for a reason. A third state is having withholding all doxastic attitudes whatever for a reason, including credal attitudes. Really only the latter two states are properly called “suspension”. And there is a difference between them. Suppose you are asked whether the coin will land heads. You can assign a definite credence of .5 to this proposition if you know it is a fair coin. For that reason, you will refrain from believing and disbelieving. Suppose instead that I ask you a completely random question about the hottest star in the Andromeda Galaxy. You will have no clue whatsoever, and it would be a mistake for you to assign *any* definite credence to any answer. For this reason, you will refrain from believing and disbelieving—and a lot more.

These are very different states. Call the first state “weak suspension” and the second state “strong suspension”. There are different rules for these types of suspension. We can use:

(WS-Rule) Weakly suspend on P only if P and \sim P have equal chances.

And also:

(SS-Rule) Strongly suspend on P if you lack any information about whether P.

What about the other state, which I claimed not to be properly called “suspension”?

Crucially, this state—which I will call “the mere absence”—is *not an alternative to belief and disbelief*. It is not an attitude at all, or something that can be had for reasons, but rather just what the name indicates—a mere absence of them. We can *cause ourselves* to suffer this mere absence for reasons. But we cannot *be in* the mere absence for reasons. We can only strongly and weakly suspend for reasons. There can, then, be no epistemic rules for the mere absence. There are *practical rules* for the *distinct act of causing ourselves to suffer mere absence*. But no rules for the mere absence, or of a kind that could be of interest to epistemologists.

This point is of great importance. Look back at the T-Rule and the F-Rule, and notice that both say “only if” and not “if and only if”. You might have thought if these rules were the only fundamental rules for belief and disbelief, we would have to believe and disbelieve *very little*. After all, we know that our faculties are fallible. So even if we are presented with a *stable* occurrent appearance as of P, it might be that \sim P. We must admit it. This means that there is almost always going to be a *risk* of violating the first two rules if we try to comply by having positive doxastic attitudes like belief and disbelief. Yet, one might say, there is a sure-fire way to comply with them: neither believe nor disbelieve! Doesn't that mean that we should *generally neither believe nor disbelieve*, since that would always guarantee compliance, and there is always a risk of violation if we either believe or disbelieve? And wouldn't that imply radical skepticism, and show that we made some huge mistake?

No. You cannot *comply* with those rules *via* mere absence, since that state simply isn't a state that can be “had” for epistemic reasons. Of course, there is a *distinct act* one could perform—causing oneself to get into that state. You can *exit* the game of epistemology with respect to a given proposition by performing such an act. But there are only *practical reasons* for that act (which may, of course, concern values *that are epistemic*), not properly epistemic reasons. So, if you want to *comply* with those rules without believing or disbelieving, your options are strong and weak suspension, or credence that falls short of belief or disbelief. Yet we now have rules for strong and weak suspension, and rules for credences not hard to generate.

The rules for the former states are stringent: only strongly suspend if you lack information about whether P, and only weakly suspend if the chances are even. So the mere risk of violating the T-Rule and F-Rule will *not* license us to strongly or weakly suspend: often we would be *certain* to violate the rules for those suspended attitudes.

That highlights the need for credal rules, which I have been setting aside. Obviously, it would be bad if avoiding violation of the WS-Rule and SS-Rule forced us to believe or disbelieve almost every proposition we consider, or perform some practical act that enabled us to “exit the game of epistemology”. If we stuck with only four rules and pretended that the trichotomy captured everything, we would be forced to say that, and that would be crazy. But there is a huge number of alternatives—any real-valued credence between 0 and 1 that doesn't, say, ground some disposition to affirm. This is an embarrassment of riches. If we want the full story about what else we can “do”, we are going to have to look at these.

That, however, is not the most important point for present purposes. The most important point is this: far from encouraging Pyrrhonian silence, our rules, if left unsupplemented by rules for credences, are going to recommend us to either believe or disbelieve tons of propositions *if* we are already in the business of attempting to take a stance on these propositions. The further reason why this matters is that it shows that there is *no pressure from concerns about avoiding skepticism to replace the T-Rule and the F-Rule with biconditional rules*.

Often, when people discuss rules of this particularly fundamental kind, they start out with biconditional rules like:

Believe P iff P

They feel pressured to use these rules, precisely since they worry that otherwise the least risky way to comply with the T-Rule is to have no beliefs at all. Yet these biconditional rules are crazy. They are far worse than my rules, left unsupplemented by credal rules. While we cannot ever in any genuine sense epistemically rationally transition from belief or disbelief to simple absence, we can transition epistemically rationally from simple absence to belief, disbelief or some degree of credence. If we bought biconditional rules, there would be a lot of pressure to do this with any proposition whatsoever, even dreaded propositions about the number of blades of grass in the lawn. The correct response to this really is to replace the biconditional rule with rules like the T-Rule and the F-Rule. It is not to tack on some clause like “for propositions of interest”. Indeed, then there would be no story to be told about the epistemic irrationality of believing uninteresting propositions—and yet the mere fact that propositions are uninteresting does *not* mean we can't believe them irrationally. And when we choose to replace the biconditional rules with the T-Rule and the F-Rule, we are not courting skepticism. The fact that there are rules for strong and weak suspension, and that simple absence *simply isn't an option* in the relevant sense, allows us to avoid that worry.

Once we've appreciated these points, we can simply stick with these fundamental rules, and analyze respect for the truth in terms of recognition respect for these rules. A completely satisfying story would require rules for credences, and a story about how credences relate to flat-out belief and disbelief. But I will set this aside.

2.2. *What respect for the truth involves*

Given the Improved Theory of recognition respect for rules in general, our work is not

hard. The main complication stems from the fact that we are dealing now with several rules, not just one. We can start off with the following reduction of respect for the truth to respect for fundamental epistemic rules:

(Thesis 1) S respects the truth in forming some doxastic attitude D iff (i) forming D exhibits recognition respect for one of the fundamental epistemic rules, where (ii) the case also stably occurrently appears to be one of a type where D-ing would not constitute a (likely) violation of any other fundamental epistemic rule.

To see why we need two clauses, consider a normal case and a fair “rational reconstruction” if it. You ask yourself whether P, and are struck with a genuine memory as of a fact that clearly entails P—a memory that wouldn't be rendered unclear by anything else you are in a position to know. Since P for these reasons seems true, you believe it. Here you believe P *because* your case stably appears to be one of a type where believing P would be believing a truth—and that is a compliance state of the T-Rule. Notice that this explanation does not involve your believing P because it would comply with other rules. This seems typical. Nevertheless, often complying with one rule could risk violation of another. But one needn't form an attitude because it would comply with all the rules for this reason. One need only form the attitude so that it complies with one of them, where it also stably appears that no other rule would be violated. Put briefly: respecting truth requires recognition respect for at least one fundamental rule, and no *disrespect* for others.

Now clause (ii) need not be unpacked further here, given our previous remarks about the notion of a stable occurrent appearance, and our working list of fundamental rules. But we should unpack clause (i), using one of the rules as our example. We can use the T-Rule:

(Thesis 2: Example with the T-Rule) Where a D-encouraging case is a case where having a doxastic attitude D would be objectively likely to approximate compliance with the T-Rule,

S exhibits recognition respect for the T-Rule by D-ing iff

(i) S forms or maintains D *because* it stably appears to S, of some possible case that would be D-encouraging, that such a case obtains.

and

(ii) S forms or maintains D in order to do what would approximate compliance with the T-Rule,

where

(iii) this is the only ultimate epistemic motive needed to explain S's D-ing.

Notably, simply believing for stably apparent sufficient epistemic reasons plausibly *suffices* for exhibiting recognition respect of the T-Rule. If you have an apparent sufficient epistemic reason for believing P, it will appear sufficiently objectively likely that P is true. If you believe P just for that reason, you will, in a clear sense, be forming an attitude *in order to what would approximate compliance with the T-Rule*. After all, you cannot *believe for* a sufficient epistemic reason in the relevant sense unless your believing *submits to* the T-Rule in a deeper—e.g., functionalist—sense. Exactly what that sense involves would require us to get into questions about the nature of belief, and into how it is that truth is its standard of correctness. The point for now is simply that such stories have been told, and it isn't irresponsible to simply defer to them for further explanation of the precise way in which believing for only good epistemic reasons counts as submitting to the T-Rule.

Believing for stably apparent sufficient epistemic reasons will also plausibly be *necessary* for exhibiting recognition respect of the T-Rule by believing. If you believe but lack stably apparent sufficient epistemic reasons for believing P, clause (i) will not be satisfied. So, for any S who plans to believe P, S's respecting the truth will require S to believe P for stably apparent sufficient epistemic reasons. And given the earlier point, believing for them will suffice for S to respect truth. That establishes strong connection between responding to apparent epistemic reasons and exhibiting recognition respect for the truth—though the connection is a substantive one that must be derived, not one true by stipulation.

2.3. *The payoff*

That is precisely why I believe that epistemic Kantianism is in a position to give a deeper explanation of the deontic significance of requirements of epistemic rationality. Having respect for the truth is going to require complying with these requirements, since these requirements centrally demand responding to apparent epistemic reasons. So if we really ought to have respect for the truth—if that, as I've claimed before, *just is* what it is for truth to be epistemically valuable—then we really ought to comply with these requirements. That gives us a deeper explanation, and enables us to solve the problems of deontic significance for these requirements—at least *modulo* a further defense of epistemic Kantianism as compared with other comparably fundamental accounts of the source of deontic significance, such as epistemic teleology in any of the various forms I've considered.

3. How Can Respect Be Central? Remarks on the Explanatory Order

I have now given a theory of recognition respect, and sketched how it and epistemic Kantianism will jointly require us to believe P when and only when we would believe P for stably apparent sufficient epistemic reasons. It may now become tempting to ask why I bothered. Why introduce the ideology of respect for the truth at all if it can simply be explained in the terms made available by the Improved Theory, and Theses 1 and 2? Indeed, once we have seen how these notions could be explained away, why doesn't epistemic Kantianism simply collapse into some variation on Huemer/Pryor-style internalism?

3.1. *Two levels of theorizing*

The second question admits of a direct response. It confuses two importantly different levels of theorizing. Phenomenal conservatism *a la* Huemer is, in effect, a first-order normative theory: it attempts to give us a general story about when, *in nonnormative terms*, our beliefs are justified. Epistemic Kantianism is a second-order, metanormative theory that seeks to explain in more fundamental normative terms what it is for truth (in belief) to be an epistemic value, and what grounds the standard of quality for epistemic reasons, given such an account of truth (in belief)'s epistemic value. Its claim is:

The pro-response whose correctness is fundamentally linked with the properly epistemic value of truth (in belief) is recognition respect, not production. So the fundamental epistemic norm is not “produce true beliefs”, but rather “respect the truth”, assuming the fundamental epistemic value is truth (in belief).

This kind of theory is intended to serve as a unifying underpinning for a first-order theories in the ballpark of Huemer view. For we can argue from the veritist version of epistemic Kantianism and claims about the nature of respect to the truth of theories of this kind.

Such an underpinning is needed. For it is easy to argue, as I did in Chapters 2 and 3, that internalist theories *a la* Huemer conflate epistemic rationality with epistemic justification, and also fail to tell us *anything* about how it is that the factors that make us epistemically rational also give us genuinely *good* epistemic reasons. Given the strong analogies with the practical literature, and the pressure that exists in that literature to claim that rationality has no deontic significance as such, this is a striking failure on behalf of these theorists. Indeed, on their own, internalist theories like these are, I believe, dead ends. Given that almost everyone thinks that truth (in belief) *is* fundamentally an epistemic goal, and that this belief makes it compelling that the standard of quality for epistemic reasons should be an objective teleological one—so that good epistemic reasons must be such that complying with them increases the chance of promoting the “truth-goal”—these theories will remain indefensible until that fundamental belief is undermined.

This is why epistemic Kantianism is not superfluous, and does not collapse into Huemer-style internalism. It is a theory of a different, more fundamental type that serves to underpin theories like Huemer-style internalism. It purports to solve what I take to be the deepest concern about all such theories. If one operates on standard teleological assumptions, distinguishes between epistemic rationality and justification, and sees the link between the latter and truly good epistemic reasons, such theories will appear not only false but confused. Yet if one rejects the standard teleological assumptions and instead embraces epistemic Kantianism, then even though epistemic rationality and justification remain *notionally* distinct, they will turn out *extensionally coincident*, as a point of substantive fact.

3.2. *Unanalyzability vs. fundamentality of theoretical role*

This does not, however, fully address the first question. There is a question about why epistemic Kantianism as I paint it should rely on the ideology of respect, given that this ideology appears to be analyzable in other terms. While I will ultimately deny in the next subsection that I have given an *analysis* of the concept of recognition respect—just an *elucidatory biconditional*—it is worth pausing on the assumption behind this question.

The assumption appears to be that insofar as some notion (or property/relation) admits of an analysis, it cannot play a fundamental theoretical role. This assumption is, I believe, false. Consider the following argument:

Williamson's Double Wrongness

1. Williamson is wrong on one count: knowledge is analyzable (say, as safe-apt belief).
2. If knowledge is analyzable, it cannot play any fundamental theoretical role.
3. So Williamson is doubly wrong: knowledge cannot play such a role.

Premise (2) is suspect. Even if knowledge consists in safe-apt belief, it could be that safe-apt belief is shown to be *important* only by being that in which knowledge consists. There simply is this difference between *conceptual* and *metaphysical* fundamentality. When it comes to the *importance* of a notion, something like a principle of organic unity often seems to hold. To take another example, pleasure is clearly not a metaphysically fundamental phenomenon. Indeed, we clearly do not need the concept of pleasure as one of the basic ingredients for “writing the book of the world”. Nevertheless, whatever turns out to constitute pleasure will be shown to matter for us only because it constitutes pleasure. It is not as if pleasure will be shown not really to matter because it reduces to certain patterns of neurons firing, or

whatever. Nor will it turn out that we can even helpfully illuminate our ordinary thinking about what matters by substituting in different concepts for the concept of pleasure.

If Williamson's Double Wrongness is unsound for these reasons, I don't see the complaint for the epistemic Kantian who wants to make recognition respect for the truth central.

3.3. *The concept of recognition respect may be irreducibly normative*

There is a second answer that we can give to the first question. While I have given an *account* of what *makes for* recognition respect for rules in general in nonnormative terms, it does not follow that I have shown how we can eliminate the concept of recognition respect *even in principle* in favor of nonnormative concepts. For this concept is, I suspect, an irreducibly normative concept. To see what I have in mind, an analogy is worth considering.

Consider the concept of a competence, as virtue epistemologists and ethicists (should) understand it. This is the concept of a disposition to *succeed*. Success is a broadly normative idea. Of course, it is easy to see how, for any given competence, we can provide an *account* of what *makes for* this competence in nonnormative terms. Success *conditions* are easily stated in nonnormative terms. On one kind of virtue epistemology, the success condition simply is true belief, and a competence will involve a disposition to believe truly (and not believe falsely). The fact that such an account can be given does not, of course, show that virtue epistemology of a suitably robust kind is impossible. A suitably robust virtue epistemology ought to claim that a fundamental normative concept needed to understand epistemic normativity in general is the concept of an epistemic virtue. This is a metanormative claim, about the fundamentality of one normative concept with respect to others. The fact that we can give a precise account of the nonnormative facts on which the facts picked out by the concept of an epistemic virtue supervene does nothing to harm this claim.

Even outside the context of such a robust virtue epistemology, I think it is doubtful that we can reduce the *concept* EPISTEMIC COMPETENCE to any nonnormative concepts. This is not out of any broad skepticism about conceptual analysis. Rather, it is because fully grasping this concept *requires* grasping concepts like SUCCESS, whereas fully grasping the concept DISPOSITION TO BELIEVE TRULY or any similar complex nonnormative concept requires nothing similar. Although concepts cannot simply be *identified* with their grasping conditions, it is plausible that if the (full) grasping conditions of two concepts are distinct, those concepts are distinct. So the concept of an epistemic competence would appear to be distinct from any complex (or simple) nonnormative concept.

I suspect that something similar might be true of the concept of recognition respect. Remember how we started out by noting that this concept seems closely tied as an *a priori* matter to the concept of carelessness, as expressed in the CARELESNESS LINK. Indeed, to fully grasp the notion of recognition respect for some rule, it seems like we need to have some grasp of the notion of respectful behavior with respect to that rule. Yet respectful behavior really just is behavior that is not, in the relevant sense, at all careless.

None of this is to suggest that I embrace some nonnaturalist metaphysics. I am simply making claims about the irreducibly normative character of certain *concepts*, not of the properties picked out by those concepts. Indeed, I assume that the relation that one always stands in when one has recognition respect for a rule *just is* the relation on the right-hand side of the Improved Theory. Or at least something nearby, since I'm not certain that this

theory is true—just that it is an improvement over all the other theories we considered.

K. L. Sylvan

References

- Bird, A. 2007. "Justified Judging." *Philosophy and Phenomenological Research* 74.1: 81–110.
- Darwall, S. 1977. "Two Kinds of Respect." *Ethics* 88: 36–49.
- _____. 2006. The Second-Person Standpoint. Cambridge MA: Harvard University Press.
- Huemer, M. 2001. Skepticism and the Veil of Perception. Lanham MD: Rowman and Littlefield.
- _____. 2007. "Compassionate Phenomenal Conservatism." *Philosophy and Phenomenological Research* 74: 30–55.
- Parfit, D. 2011. On What Matters. Oxford: Oxford University Press.
- Pryor, J. 2000. "The Skeptic and the Dogmatist." *Nous* 34: 517–49.
- Williamson, T. 2000. Knowledge and its Limits. Oxford: Oxford University Press.