Objectivism and subjectivism in epistemology

Littlejohn, Clayton

Objectivism and Subjectivism in Epistemology

Clayton Littlejohn
King's College London

Draft of 1.8.2016

Abstract: There is a kind of objectivism in epistemology that involves the acceptance of objective epistemic norms. It is generally regarded as harmless. There is another kind of objectivism in epistemology that involves the acceptance of an objectivist account of justification, one that takes the justification of a belief to turn on its accuracy. It is generally regarded as hopeless. It is a strange and unfortunate sociological fact that these attitudes are so prevalent. Objectivism about norms and justification stand or fall together. Justification is simply a matter of conforming to norms. In this essay, I shall make the case for objectivism about justification.

1. Introduction

I will defend two objectivist views. The first is a view about epistemic norms. The objectivist about epistemic norms believes that when we specify the application conditions for these norms we'll sometimes refer to objective conditions (i.e., conditions that don't supervene upon our non-factive mental states or the subjective aspects of our mental lives). The second is a view about epistemic justification. The objectivist about epistemic justification believes that the justification of our beliefs depends upon certain objective conditions. The justification of our beliefs depends upon whether we violate any norms. Objectivism about norms and justification would seem to stand or fall together.

Norms identify the conditions under which someone should or should not believe, do, or feel something. If you shouldn't believe lottery propositions or should increase your confidence when you acquire new evidence for your belief, there should be some norm that says as much. The objectivist and subjectivist agree that some norms will have subjective application conditions. The fact that you are probabilistically coherent supervenes upon subjective conditions, so if there is a norm that requires coherence there is a norm that has a subjective application condition. Their disagreement is about whether there are any further norms, norms that require us to believe or refrain from believing when objective conditions obtain. They would disagree about norms such as these:

> You shouldn't believe *p* unless you know *p* (Only Knowledge).[1]
>
> You shouldn't believe *p* unless *p* is true (Only Truth).[2]

---

[1] For defenses of knowledge norms for assertion, belief, or reasoning, see Benton (2011), Bird (2007), Littlejohn (2013), Smithies (2013), Steglich-Petersen (2013), Sutton (2007), Turri (2011), and Williamson (2000).

[2] If Only Knowledge governs belief, so does Only Truth. For defenses of Only Truth that aren't intended to support Only Knowledge, see Boghossian (2008), Nolfi (2015),

> You should believe *p* if you are in a position to know *p* and it matters whether you believe *p* (Any Significant Knowledge).

While objectivists will disagree about which of these norms, if any, governs belief, subjectivists insist that all the real norms focus on relations between your beliefs and the subjective aspects of your mental life. The normative, they'll say, is concerned only with the subjective.

As I understand justification, an action or attitude is justified when it is proper, acceptable, permitted, right, or appropriate. If you have an attitude or perform an action and it's not the case that you shouldn't, the attitude or action counts as justified. Norms identify the conditions under which you should or should not have an attitude or perform an action, so we can characterize justification in terms of norm conformity and violation. On my simple view of things, objectivism about justification follows from objectivism about norms. While we'll see that some will object to this simple view, I would suggest that a belief is justified iff it violates no epistemic norms, which is to say that it is justified iff the thinker conforms to the epistemic norms. If we think of justification in this way, it would seem that objectivism about norms and justification should stand or fall together.[3]

For various reasons, most epistemologists reject the objectivist view of justification that I shall defend here. I shall make a case for objectivism about justification and norms in the next section and then consider a subjectivist response in the section after that. Convincing the reader that the view defended here is correct is a task that couldn't be done in the space of a single paper. I shall offer the arguments for objectivism that I take to be interesting and important. Some of the controversial assumptions of these arguments have been defended at greater length elsewhere. Space also won't permit an extensive discussion of the arguments for subjectivism, so I've decided to focus on a recent defense of subjectivism because it might help us see why so many epistemologists are objectivists about norms and subjectivists about justification. It is the only defense of subjectivism about justification that hasn't yet been discussed in the extant literature.

2. Objective or Subjective?

Let's start with a simple argument for objectivism about justification:

> *An Argument for Objectivism*
>
> O1. Your beliefs are justified iff you violate no epistemic norms in holding these beliefs.

Wedgwood (2002, 2013), and Whiting (2013). For criticism, see Brown (2008), Gerken (2011) and McKinnon (2013). For a helpful overview of the extant debate, see McGlynn (2014).

[3] While Littlejohn (2012), Sutton (2007), and Williamson (forthcoming) all defend views on which beliefs are justified only when they don't violate objective norms, most of the writers mentioned above won't defend objectivism about justification by appeal to some objectivist view about norms.

O2. You violate no epistemic norms in holding these beliefs iff these beliefs constitute knowledge. [4]

OC. Your beliefs are justified iff they constitute knowledge.

If the argument is sound, readers should take the 'factive turn' and embrace an objectivist account of justification. Naturally, I don't expect readers to take this argument to be sound at this stage of the discussion, but I shall offer my supporting arguments momentarily.

If the objectivist view of justification defended here is correct, it should be possible for pairs of thinkers to process things in just the same way, draw the very same conclusions using the very same patterns of reasoning after things appear precisely the same to them, and still form beliefs that differ in justificatory status. This suggests that a kind of rational perfection might fall short of ensuring that some subject's beliefs are justified. If you like to talk about luck, you could say that the objectivist view of justification that identifies justified belief with knowledge allows for a certain kind of luck that many epistemologists won't allow for and doesn't allow for a certain kind of luck that many epistemologists would allow for. I suspect that many epistemologists are opposed to the idea that a thinker might fail to acquire justified beliefs because they suffer the kind of bad luck that would prevent a rational thinker who reasons impeccably from acquiring knowledge. One reason that I'm opposed to the subjectivist view is precisely because they want to try to show that this kind of luck isn't possible.

---

[4] There is an interesting question about (O2) that we don't have space to discuss in detail. Is it possible to know *p* and violate some epistemic norm that requires you not to believe *p*? While it's an interesting question whether there could be cases in which someone shouldn't believe what they know, answering this question won't help us settle the debates between the objectivists and subjectivists. There are some epistemologists who think that it's just obvious that anything you know you justifiably believe because they think that justification is a part of knowledge (whatever that means). This is not my picture of justification. On my account of justification, the justification property is a property that a belief has iff it violates no epistemic norm (or, more carefully, does not do so without sufficient reason). Because justification is not part of the nature of knowledge, it isn't trivially true that anything known is thereby justifiably believed and thereby believed without violating any epistemic norms. On my view, knowledge is a relation between a thinker and a fact, a relation that holds between some thinkers and some facts where the thinker's beliefs could not have the justification property because these thinkers cannot be held accountable for their attitudes or actions. Our beloved dogs and cats presumably know things about their environments, but I don't think we'd want to say that their beliefs could be justified or unjustified. I think that only thinkers that can be held accountable can have justified or unjustified beliefs. On my picture, the relationship between justification and knowledge is analogous to the utilitarian attitude towards the relationship between right action and optimific actions. A dog's actions might be optimific without being right, but the actions of responsible agents could not be anything but right if optimific.

This knowledge account of justification implies controversially that a certain kind of luck that many epistemologists take to have no bearing on justification can rob you of justification. The lucky connection between a thinker and a fact that robs you of knowledge is one that robs you of justification on the present view. In the course of our discussion, I shall explain why I think it's important for our view of norms and justification to register that the kind of accidental connections between thinker and fact that we're familiar with from the Gettier literature prevents you from acquiring beliefs that have good standing.

For their part, the subjectivist is quite likely to think that I've just made three significant errors. It's a mistake to think that truth is required for justification (and possibly a mistake to think that its required for conforming to epistemic norms), so naturally it is a mistake to think that a non-accidental connection to the truth is required for justification. It's a mistake to think that the kind of bad luck we suffer from when we're in hostile environments robs of justification by leading us into error precisely because these environments don't have any direct bearing on how our beliefs fit with appearances. This matters, they might think, because all epistemic assessment is concerned with the kind of processing that takes us from appearance to judgment or belief. Finally, they might also question the suggestion that rational perfection might fall short when it comes to satisfying a genuine normative requirement. The real norms, they might say, are only concerned with the relations between appearances and belief precisely because the norms are concerned only with how we reason, how we process what's given to us, and how we exercise our rational capacities.

Perhaps this ideology explains why many prominent subjectivists about justification would find Huemer's recent argument for subjectivism compelling:

> *An Argument for Subjectivism*
> S1. If *p* and *q* seem the same in all epistemically relevant respects to a perfectly rational person, she would accord the same credence to *p* and *q*.[5]

---

[5] In the context of the paper, it is clear that Huemer intends the antecedent to be read as saying that *p* and *q* seem the same in all epistemically relevant respects *to the particular thinker* whose attitudes we're assessing. He is not interested in whether *p* and *q* seem the same in such respects to an external observer who knows, say, that *p* is true and *q* is false. It should be noted that (S1) is *not* just a restatement of the idea that justification supervenes upon appearances. It is also not a trivial consequence of that supervenience thesis. Someone could reject (S1) and still defend that supervenience thesis. Suppose there's some rational requirement that says, in effect, that Agnes should suspend on *p* and not suspend on *q* when things appear a certain way to her. Things could appear the relevant way to Agnes, Agnes could be required to suspend on whether *p* but permitted to believe *q*, but it *also* might seem to Agnes that *p* and *q* are the same in all epistemically relevant respects. (Someone could say that this means that Agnes is somehow less than fully rational, but I don't see that. Provided that Agnes doesn't violate these rational requirements, I don't see her as rationally simply because she doesn't appreciate the rational difference between the propositions that she entertains but doesn't believe.) For

S2. If a perfectly rational person would accord the same credence to *p* and *q*, then *p* and *q* have the same degree of justification.
S3. The propositions *p* and *q* seem the same to our thinker in all epistemically relevant respects.
SC. Therefore, *p* and *q* have the same degree of justification for our thinker.[6]

We should read (SC) as implying that the thinker is in a position to justifiably believe *p* iff she is in a position to justifiably believe *q*. Because it is possible for the thinker's beliefs in *p* and *q* to differ in truth-value when *p* and *q* seem to them the same in all relevant respects, this argument, if sound, would force the objectivist to reject (O2) and possibly (O1).

In choosing my arguments for (O1) and (O2), I chose arguments that would show why I think it's important for us to resist arguments like this argument for subjectivism. It's clear that Huemer thinks that facts about how things seem or appear have a great deal of normative significance, but this argument doesn't tell us why we should think that. Some would ask us to consider error cases to try to elicit intuitions about rationality to shore up support for the key premises.[7] Some would instead suggest that appearances play the role of evidence or ground and look to defend this argument by appeal to some norms that tell us that justification is entirely a matter of how our beliefs fit the evidence or grounds available.[8] The arguments I've chosen should help readers see why I don't think that these ancillary considerations do much to support the subjectivist cause.

Does this argument need support from ancillary considerations? I think that it does. I have to confess that there was a time when I probably found this kind of argument quite troubling, but I don't now. The argument runs from some suggestions about rational credence to a conclusion about the justification of full belief. Since I'm interested in the justification of full belief, it's worth asking whether there's any good reason to think that the justificatory status of full belief in *p* and *q* would be the same if the subject rationally invested the same credence in these propositions. Lottery cases cause trouble for this proposal. When it comes to beliefs in lottery propositions and beliefs based on what you read in the paper, it seems plausible that you should have higher credence when it comes to lottery propositions but it also seems that you cannot justifiably believe such propositions even though you can justifiably believe what you read in the paper. Part of this is down to the fact that what you read in the paper isn't generally a known unknown (to use Sutton's terminology), something you know you couldn't know. Beliefs in lottery propositions are known unknowns. This explains, I'd

---

criticism of (S1) and its use in Huemer's argument for phenomental conservatism, see Littlejohn (2011).

[6] Huemer (forthcoming).

[7] See Cohen (1984) and Wedgwood (2002).

[8] See Dougherty (2011) for a discussion of combining Conee and Feldman's (2004) evidentialist view with phenomenal conservatism.

argue, why the respective beliefs can differ in justificatory status and why we should be skeptical of (S2).

Of course, it could be said that this is just a distraction because the objectivist and subjectivist can both agree that you shouldn't believe lottery propositions and can both agree that lottery propositions show that (S2) is mistaken. I agree. Once we strip away the questionable assumptions about the relationship between rational degrees of credence and the justification of full belief, it looks like we're left with nothing but the bare claim that if things appear the same with respect to two prospective beliefs, they have to agree in justificatory status. Why should we accept this? It's at just this point that the subjectivist will have to offer some fresh argument for thinking that the norms only care about relations between the subjective aspects of our mental life and our beliefs or some fresh argument to show that these appearances are the only possible rational basis we have to go on and use some implicit evidentialist assumption to try to justify (SC). My arguments for objectivism are supposed to speak to both issues. The first argument is an attempt to show that it's important to recognize that there are norms that are concerned with more than just relations between beliefs and appearances. The second is an attempt to show that it's a mistake to try to show that the appearances wholly determine what evidence we have and how we ought to be guided by it.

2.1 The Transcendental Argument

My first argument for objectivism about norms and justification is a transcendental argument. The argument starts from an assumption about practical requirements. There are practical requirements to conform to objectivist norms. It then proceeds to show that these norms are binding only if there are objectivist norms that govern belief and made justification turn on whether we conform to Only Truth. This argument will not vindicate Only Knowledge, but it establishes an interesting form of objectivism and moves the discussion to a discussion about the best version of the objectivist view.

The starting assumption is that there are practical requirements to conform to objectivist norms, norms that have external application conditions:

> You shouldn't sentence the innocent for crimes they didn't commit (Only the Guilty).
>
> You should sentence those you know to be guilty (Any Known Offender).
>
> You should repair your past (objective) wrongs (Repair).[9]

The fact that Agnes wasn't involved doesn't supervene upon the subjective aspects of your mental life. We can imagine a world in which you know that she was involved and you punish her in light of what you know. This might be permissible. We can hold the subjective states constant and then imagine another world where Agnes was framed. In this world you'd violate Only the Guilty if you punished. You would run the risk of violating Repair if we didn't make appropriate reparations.

I don't have any complicated arguments for Only the Guilty or Repair. They both strike me as intuitively plausible.[10] Consider the case in which you punish Agnes

---

[9] An objective wrong is a wrong that involves violating an objectivist norm, such as Only the Guilty.

for something that she did not do. You discover that she was innocent and that the evidence that led you to punish was misleading. Bad things happen all the time and you don't always have any special connection to them. It seems that this bad thing is something that you have a special obligation to do something about. You should assume responsibility for this harm and for trying to make this right even though it's not true that you should assume responsibility for similar harms, much less similar harms that you cause.[11] This suggests that there's a reparative duty, a duty to right something that wasn't simply some harm you caused Agnes to suffer, but some harm that you wrongfully brought about. The duty arises because of a morally significant relation between you and Agnes that doesn't hold between Agnes and other parties not involved in the punishment who could take steps to make her better off.[12] The fact that you weren't aware of the morally significant relationship you'd stand in by virtue of punishing her for something she didn't do doesn't erase the debt that you uniquely owe to her.[13]

I think it's important to resist the urge to rewrite these norms to fit the subjectivist framework. The subjectivist is right that we need to think about the subjective aspects of an agent's life when assessing the agent and her conduct. The subjective matters when we're interested in things like credit, blame, criticism, and the like, but this is only an aspect of morality and it isn't somehow more important than other aspects of morality. There is another part of morality that isn't directly concerned

---

[10] I should also add that if there were a group of people who acted as if these norms were binding, it would be hard to see what could be objectionable about it. We can imagine alternative practices, practices where people acted as if only subjective correlates of these norms would be binding and this wouldn't seem to be a terrible practice. It would, however, be a practice that, to my mind, left out some nice features of the objective practice (e.g., a set of duties that required compensation for acts that that objective view takes to be wrongful) and wouldn't appear to have many particularly important additional virtues that would offset the cost. In short, it would be nice if morality incorporated norms like Only the Guilty, Any Known Offender, and Repair and it wouldn't be nicer if morality dropped these in favor of subjective correlates. If true, might this be some evidence that morality involves objective elements? Perhaps. See Enoch (2009), Preston-Roedder (2014), and Sayre-McCord (MS) for discussion of whether it might.

[11] For an extended discussion on this point, see Littlejohn (2012).

[12] For further defense of this description of the case, see Littlejohn (2012). Zimmerman (2008) defends a kind of subjectivist view, one that insists that what we should do is always determined by a kind of expected value. He thinks (rightly, in my view) that subjectivists have to deny Repair and deny Only the Guilty. We shall discuss this further in §3.

[13] The intuition that suggests that it's possible for conscientious agents to violate moral norms owing to factual ignorance or mistake is reasonably widespread, even among those who defend subjectivist theories. See Enoch (2012), Herman (1993), and MacFarlane (2016) for interesting attempts to vindicate the intuitions that seem to support objectivist norms.

with such things and it's the part of morality that gives us guidelines that help to determine when someone has the right to proceed without interference, when someone should be protected against those who would harm them or their interests, and when some party need to take steps to apologize and make reparations. It's this part of morality that the objectivist gets right. Getting this right requires identifying standards that should guide action even when not seen perfectly by an agent. I fear that subjectivism, if left unchecked, leads us to draw mistaken conclusions about the kind of responsibility we should assume in the wake of violating some (putative) objectivist norm because it tries to deny the possibility that such guidelines have normative force when some agent has imperfect access to them.

The second assumption in the argument is *unificationism*, the view that says that there is a principled connection between the justificatory status of beliefs about what's required of us in the situations we're in and the justificatory status of the responses we believe to be required. Unificationism says that if someone justifiably judges that they are required to Φshe couldn't be required to do other than Φ[14] The *segregationist* denies this. The segregationist says that we can be required to do other than Φin situations where we nevertheless justifiably judge that we are required to Φ The disagreement between the unificationist and segregationist is about a kind of detachment. Everyone agrees that if someone believes that they should Φ nothing interesting follows about whether they really should. The interesting question is about the case in which they believe they should Φand this belief is normatively appropriate. If this belief is sanctioned, could it then be wrong for the subject to Φ The unificationist thinks that it couldn't be wrong for the subject to Φin line with her normative judgment.[15] The segregationist thinks that it can be.

We can use the Enkratic Requirement to defend unificationism:

> Enkratic Requirement: You ought to see to it that: if you judge that you yourself ought to Φyou Φ

The Enkratic Requirement is a wide-scope requirement, one that rules out certain combinations or patterns. The unificationist doesn't defend a wide-scope requirement that rules out combinations or patterns, but instead defends a conditional that says, in effect, that whenever you shouldn't Φyou shouldn't judge or believe that you yourself should ΦSome think that it is quite intuitive that we shouldn't be akratic, but how does this intuitive requirement about acceptable combinations give us a conditional that tells

---

[14] The relevant requirements are those that have a special kind of rational authority such that we count as irrational or unreasonable if we give these requirements no weight. Requirements of the law or etiquette

[15] For defenses of unificationism, see Foley (2001), Gibbons (2010), Littlejohn (2012), and Way and Whiting (forthcoming). Greco (2014), Ichikawa and Jenkins (2013), Kiesewetter (forthcoming), Titelbaum (2015), and Smithies (2011) defend limited versions of unificationism that apply just in the epistemic realm. For defenses of segregationism, see Coates (2012), Fantl and McGrath (2009), Feldman (2008), and Lasonen-Aarnio (MS).

us that the justification of our attitudes about Φing is linked to the justificatory status of Φing?

Titelbaum (2015) shows us that there is a very simple argument from the Enkratic Requirement to unificationism. Suppose that you ought not ΦIf so, could you nevertheless justifiably judge that you yourself ought to Φ No, not if the Enkratic Requirement is a genuine requirement. We can rewrite the requirement as follows: you ought to see to it either (you do not believe that you ought to Φor you Φ. Because we can assume that you ought not Φwe can derive that you ought to see to it that you do not believe that you ought to ΦThus, *if* it is not the case that you ought not believe that you ought to Φit is not the case that you ought not Φ In the case where you justifiably judge that you ought to Φyou couldn't be in a case where you ought to do other Φ

Using the Enkratic Requirement, we can establish that false propositions about what you yourself are required to do are not things that you can justifiably believe. Let's consider two lessons we can take from this.

The first lesson is negative. It looks like we can argue from the impossibility of a certain kind of false, justified normative belief to the denial of these orthodox accounts of justification:

> Apparent Truth: If it seems to a subject that *p* and she has no available defeaters, her belief in *p* would be justified if based on this seeming.[16]
>
> Sufficient Strength: If a subject who doesn't believe *p* has sufficiently strong evidence for believing *p*, her belief in *p* would be justified if based on this evidence.[17]
>
> Mentalism: The justificatory status of a belief supervenes upon the subject's non-factive mental states and the subjective aspects of her mental events.[18]
>
> Reliabilism: If there is a sufficiently reliable process that produces the subject's belief in *p*, that belief is justified.[19]

On these views, justification is not generally factive. By focusing on propositions about what you yourself should do (where the thing done requires responding to things in your environment), we can generate counterexamples to these views.[20] We can imagine pairs of subjects with the same evidence, with the same mental states, with the same equally reliable processes coming to believe that she is required to Φwhere one gets it right and the other gets it wrong because she is prohibited from Φing and thereby show that justification involves something that we don't get with sufficient strength,

---

[16] This is an increasingly popular view. See, for example, Brogaard (2016), Huemer (2007), and Tucker (2010).

[17] See Conee and Feldman (2004).

[18] See Cohen (1984) and Conee and Feldman (2004).

[19] See Goldman (1986).

[20] Proponents of these views might say that their views are *not* subject to these counterexamples, but see the next section for a response to a line of response I think isn't very promising.

something that doesn't supervene upon the mental, and something that the most reliable processes and methods will get wrong.[21]

The second lesson points us in a positive direction. We can now offer our first argument for Only Truth:

A Transcendental Argument

T1. There are objective practical norms that sometimes say that we should not Φeven when we're not aware that we should not Φand not aware of the conditions that constitute the relevant norm's application conditions.

T2. If there are such norms, there are situations where we should not judge that we are required to Φeven when we're not aware that we shouldn't Φand are not aware of the conditions that constitute the relevant norm's application conditions.

T3. The best explanation of this is that Only Truth governs belief.

TC. So, Only Truth governs belief.

This is an abductive argument, so much rests on (T3). While (T1) and (T2) only rule out the possibility of a kind of false, justified belief (i.e., one that is about what we're required to do), there doesn't seem to be anything special about such beliefs from the epistemic point of view. It seems that only norms like Only Truth and Only Knowledge provide any explanation here as to why justification requires truth. So, using the Enkratic Requirement, we can show that objectivism about epistemic norms is required for objectivism about practical norms.

2.2 Normative Beliefs and Normative Mistakes

Let's consider two potential lines of objection. The first is concerned with the objectivist assumption about practical norms in the transcendental argument. The second is concerned with the assumption that we should invoke norms like Only Truth to explain why certain false, normative beliefs couldn't be justified.

Some readers would urge us to drop the assumption that there are objectivist practical norms. If we drop that assumption, we lose the argument sketched in the

---

[21] It is worth noting that the argument for unificationism could go through even if we dropped our objectivist assumptions about practical norms. Even without these norms, we could cause trouble for these orthodox accounts of justification for it seems that each of these accounts should allow for the possibility of false but justified beliefs about what we ought to believe, feel, or do even if we assume that facts about these oughts supervene upon a subject's subjective mental life. Once we have established that there cannot be false, justified normative beliefs, it will then be difficult to resist an argument for Only Truth because we couldn't appeal to these (apparently) falsified accounts of justification to resist the arguments for Only Truth and (as we'll see below) because once we have the argument for the claim that normative beliefs have to be true to be justified we can then argue that the set of beliefs that would have to be true if justified includes non-normative beliefs.

previous section. It's interesting that we don't save the orthodox views of justification that imply that there can be false, justified belief about what a subject is required to do. It still looks as if we might be able to run a revised abductive argument for Only Truth even if we drop (T1).

Using the Enkratic Requirement, we've established that when a subject is required not to Φ this subject cannot justifiably judge that she's required to Φ. This means that the subject cannot justifiably judge that she's required to required to Φ when it appears to her that she is so required, that she cannot justifiably judge that she's required to Φ when she has sufficiently strong evidence for this belief, and cannot justifiably judge that she's required to required to Φ when this belief is produced by a fallible but highly reliable process. Because of this, the question still remains: if belief isn't governed by Only Truth, how could it be that it's impossible to justifiably judge that you should Φ in circumstances where you're required not to Φ?

To see that there's a lingering issue here, let's suppose we reject objectivism about practical requirements and adopt a view on which a subject is required to Φ when the expected value of Φing is greater than that of any alternatives to Φing. It isn't clear how this subjectivist view of practical requirements helps the proponents of the orthodox accounts of justification. A subject who is ignorant about the relevant values or misled about the values could, it seems, come to well-founded or reliable but mistaken judgments about which option maximizes expected value. She could presumably come to the judgment that she's required to Φ when in fact there's an alternative to Φing that has greater expected value. Imagine pairs of thinkers with very similar beliefs but different evaluative evidence when it comes to the moral status of fetal harms. Because of this difference in evaluative evidence, we'd expect that there would be resultant differences in which well-founded beliefs they could have about the options that would or would not maximize expected value. It isn't hard then to imagine situations where these thinkers are both required to refrain from Φing but differ in that one has the well-founded belief that Φing maximizes expected value when the other has the well-founded belief that it does not. A unificationist might say about this case that the subject who has the mistaken but well-founded belief about which options maximize expected value couldn't have a justified belief about what maximized expected value if such a belief could be justified only if justifiably acted upon.[22]

In formulating the transcendental argument, I helped myself to objectivist assumptions about practical norms because I wanted to fill out the details of a philosophical view that I think is true. This assumption can be relaxed and we can still argue from the impossibility of a certain kind of false, justified normative belief to the denial of orthodox accounts of justification that allow for justified, false beliefs in cases where the beliefs are supported by good evidence or produced by a reliable process by focusing on cases of evaluative ignorance and mistake instead of factual ignorance and mistake.

---

[22] There are interesting disagreements about whether evaluative ignorance, uncertainty, or mistake has any effect on what an agent should do. For arguments that it can, see Zimmerman (2008). For arguments that it cannot, see and Littlejohn (2014).

There is a second objection to the transcendental argument to consider. The unificationist view can be understood in two very different ways. If it's understood as a *bottom-up* view, the transcendental argument should go through. If, however, it is understood as a *top-down* view, the argument might seem less plausible. Remember that unificationism tells us that a specific kind of false belief cannot be justifiably held. If you're in a case in which you shouldn't Φ it couldn't be that your belief that you yourself should Φ is both justified and mistaken. The disagreement between the bottom-up and top-down theorists is a disagreement about why this conditional is true. If we understand this view as a bottom-up view, we would say that there are some facts about what a subject should or should not believe, feel, or do. They are independent targets that our beliefs are supposed to faithfully represent. They don't shift, sway, or move when we form beliefs about them, not even if the beliefs are reasonable, backed by good evidence, or produced by reliable processes. Facts about what should or should not be done, felt, or believed are explanatorily prior to facts about what should or should not be believed, so these normative facts place normative constraints on our normative beliefs.

We could understand unificationism as a top-down view.[23] On the top-down version of the view, we'd say that the reason that the belief that you ought to Φ could not be both justified and mistaken is that these normative facts are malleable or shifty. Part of what determines whether you should or shouldn't Φ are your normative beliefs and the things that figure in their justification. On the top-down view we would say that if the processes that produce our normative beliefs are reliable, the evidence that supports them is sufficiently strong, or they fit the appearances, they count as justified and *then* it couldn't be that you shouldn't Φ if you've judged that you should.

The abductive argument for Only Truth might appeal to bottom-up theorists but top-down theorists can object as follows. If, they might say, you had shown that normative beliefs couldn't be justified *because* they are false, that would provide support for Only Truth. The argument for unificationism, however, only establishes that there are some beliefs that cannot be justified *and* false. The top-down view gives us that, but it doesn't give us any reason to think that the relevant beliefs couldn't be justified by virtue of being false. Thus, there's no reason to think that we need to appeal to Only Truth to explain why normative beliefs are never justified if false.

This is a significant theoretical worry, but it shouldn't be difficult to put to rest. The top-down view faces serious problems that don't arise for the bottom-up view. We can see this if we think about inferences of this form:

I must Φ.
I cannot Φ unless I Ψ.
So, I must Ψ.

And think about inferences of this form:

I must give this to A.
A is B.
So, I must give this to B.

---

[23] This seems to be the kind of view that Foley (2001) and Gibbons (2010) prefer.

These inferences cause trouble for the top-down view. Let's suppose that the first premise in each inference is known and the second premise is false but (if this is possible) justifiably believed. Under these conditions, the conclusions should be justifiably believed if competently deduced. Depending upon how the details are filled in, we could have cases where the subject cannot both Φand Ψor cannot give the item to A and to B. Given the unificationist assumptions, we get the result that while the subject ought to do something (i.e., Φor give something to A) they may do something that would prevent them from doing it (i.e., Ψor give something to B). Thus the top-down view loses the seemingly sensible principle that if you may do something, you can do it without thereby preventing yourself from fulfilling an obligation.[24]

There are further difficulties for the top-down approach that have to do with blame, responsibility, and fetishism. We would generally want to blame people for failing to do things that they know they ought to do and shouldn't blame people for actions that do not manifest *de re* unresponsiveness.[25] It's difficult to see how the top-down view could respect both of these constraints. The top-down explanation as to why there couldn't be false, justified normative beliefs doesn't appeal to Only Truth, but some mechanism by which facts about what a subject should do are 'shifted' to fit normative beliefs. Consider a case of shifting. If a subject fails to Φwhen she fails to act on the justified belief that she must Φin a shifting case (i.e., a case where the factors that explain the belief and account for its justification help to move the normative facts into

---

[24] Way and Whiting (forthcoming) rightly note that previous arguments against the possibility of false, justified belief found in Littlejohn (2012) only support the thesis that if you justifiably judge that you ought to Φthat you may Φand seek to defend something stronger using argument strategies that seem essentially the same. They try to show that if you justifiably judge that you ought to Φthen you ought to ΦIt isn't clear whether they prefer the top-down approach to the bottom-up approach. They also don't have any settled view about whether non-normative beliefs could be both justified and false. The worries about 'ought' and 'can' raise much more serious difficulties for their stronger principle if combined with the top-down approach. Facts about what an agent ought to do depend upon facts about what the agent has the opportunity and ability to do. Consider the first inference. We might imagine that in the actual world Agnes knows the premises and knowingly infers the conclusion. We then might imagine Agnes knowing the first premise, justifiably but falsely believing the second, and inferring the conclusion when the act isn't one that she can perform. The stronger principle that Way and Whiting endorse commits us to saying that Agnes ought to do what she cannot do. This problem doesn't arise on a view that incorporates Only Truth or on the bottom-up views, but unificationists who don't adopt these positions either have to deny that 'ought' implies 'can' or adopt a highly skeptical view. See Littlejohn (2012) for discussion. Of course there are some authors who do not accept 'ought' implies 'can', but those who do not sometimes think that there are certain 'crazy actions' (e.g., those involving time-travel) that would be good that we cannot be obligated to do, so they might endorse something like 'ought' implies either 'can' or 'not a crazy action'. The weaker principle will steal wreak havoc.

[25] See Arpaly (2002) for discussion.

position to make the belief true), the failure to Φ would *not* manifest *de re* unresponsiveness.

Finally, closure principles for justification cause difficulty for the top-down view. The top-down view is supposed to vindicate the idea that certain kinds of mistaken normative beliefs cannot be justified and false without appeal to Only Truth. Consider, then, inferences of the following form:

I ought to Φ
If I ought to Φ, I can Φ.
So, I can Φ.

The top-down theorist wants to say that the first premise might be true in virtue of the subject's evidence, how things appear to the subject, or because the subject's beliefs are produced by reliable processes. Facts about what the subject can or cannot do are not among the sorts of things that constrain what evidence she has, how things appear, or whether her beliefs are produced by reliable processes, so difficulty will arise for the view when (a) the subject's strong epistemic position with respect to the normative proposition inclines us to accept that the subject justifiably believes the first premise and (b) the conclusion is false. If the conclusion is false but 'ought' implies 'can', it seems that the top-down theorist would have to say that in spite of the strong evidence, the appearances, the reliable processes, or whatever it is that she thinks would generally ensure that a belief is justified being in place, these things could not justify belief in the first premise if the conclusion of this inference is mistaken.[26] Thus, the top-down theorist would either have to say that *nothing* justifies belief in the first premise, nothing could justify belief in the first premise if the agent's ability to Φ depends upon how things are in the environment, or adopt a very bizarre objectivist view of the justification of normative beliefs according to which their truth isn't a constraint on whether they can be justifiably held but facts about what we can do are among the facts that constrain whether the belief can be justifiably held.

2.3 Follow (Just) the Evidence

In recent discussions of the nature of evidence and its theoretical roles, these principles concerning evidence and belief-revision enjoy widespread acceptance:

> Revise: If you have rational conditional beliefs and you acquire a new piece of evidence p then you ought to adjust your credence in q so that $P_{new}q = P_{old}(q|p)$.
>
> Remain: If you have rational conditional beliefs you should not update on p unless you acquire p as a new piece of evidence.

---

[26] I appreciate that 'ought' implies 'can' is controversial. My impression is that its critics tend to agree that certain kinds of 'crazy' actions aren't actions that we ought to perform even if such actions would bring about great amounts of good (e.g., travelling through time, changing the past from the comfort of an armchair in the present). If 'ought' implies 'isn't a crazy action', we can get a version of this argument off the ground. Thanks to xyz for discussion.

In combination, Revise and Remain tell us that we ought to update and revise our beliefs on the evidence and nothing but the evidence.

We can use these principles to generate an argument for objectivism about norms and justification. To see how, consider two questions about evidence, a question about the role of justified belief in reasoning and a question about the ontology of evidence. Suppose that you come to justifiably believe *p*. You see that *q* would be true if *p*. What attitude should you take towards *q*? If we accept closure principles for justification, it would be proper or acceptable for you to believe *q*. That suggests that if you justifiably believe *p* and infer *q*, you wouldn't have violated Remain. That suggests, in turn, that whenever you justifiably believe *p* and see that *q* is true if *p* that there is something or other in your evidence that is the kind of thing you should serve as the basis for belief-revision and updates. A natural explanation for this is the following:

> Justified Evidence: If you justifiably believe *p*, your evidence includes *p*.[27]

Now let's think about the ontology of evidence. We're interested in *propositionally* specified evidence or reasons, the things that are (potentially) the subject's reasons for believing, feeling, or doing things. The subject's reasons are the things that figure in her reasoning and stand in logical relations. Ascriptions that specify the agent's propositionally specified reasons all entail corresponding propositional knowledge claims.[28] If we say, 'Agnes' reason for ϕ-ing is that *p*', what we say is true only if, 'Agnes knows that *p*' is true. Thus, we have some linguistic evidence for half of Williamson's equation:

> E=K: Your evidence includes *p* iff you know *p*.[29]

The argument I shall offer only requires that all propositional evidence is an object of propositional knowledge, but we get a quick argument for identifying all objects of propositional knowledge with evidence if we accept Justified Evidence and take knowledge to be sufficient for justified belief:

> *An Evidentialist Argument for Objectivism*
>
> E1. If you justifiably believe *p*, *p* is part of your evidence.
>
> E2. If *p* is part of your evidence, you know *p*.
>
> EC. If you justifiably believe *p*, you know *p*.

This gives us the result we want. A belief is justified only if it is knowledge. Thus, there is a norm that requires us to believe only what we know, Only Knowledge. If we violate it, our beliefs aren't justified.

Like unificationism, Justified Evidence can be understood in different ways. On one way of interpreting it we're supposed to start with some independent notion of what justification is, think about whether beliefs that have the justification property can

---

[27] For defenses of Justified Evidence, see Comesana (2016), Fantl and McGrath (2009), and Littlejohn (2012).

[28] See Hyman (1999), Unger (1975), and Williamson (2000).

[29] For defenses of E=K, see Williamson (2000). Hyman (1999) defends a similar view about potential motivating reasons. For criticism that doesn't focus on the truth-requirement on reasons or evidence, see Hughes (2014) and Locke (2015).

properly figure in practical reasoning, and conclude that because they can it must be that they contribute evidence to that reasoning. On another reading, the reading that I prefer, the idea isn't that a proposition gets turned into evidence by virtue of being the content of a belief that is well-supported or produced by a good process. Instead, the idea is that the justification of a belief turns on whether its contribution to rational processes (e.g., theoretical reasoning or practical reasoning) is a good one. If not, the belief doesn't meet the crucial standard that we use to assess beliefs. The proper order of explanation is like this. If a belief cannot supply what it should to the rational processes that involve it (i.e., a potential motivating reason), it couldn't be justified because it couldn't meet the crucial standard. If, however, your belief in *p* can provide what it's supposed to and put the thinker in a position to Φfor the reason that *p*, it meets the standard that we use to assess belief and couldn't fail to be justified.

These different readings of Justified Evidence point to very different explanations as to why something counts as evidence. On the first, there's something that counts as evidence or a potential motivating reason because the subject has some *assets* that make a belief that has the proposition as its content justified and thereby turns the belief's object into a piece of evidence or potential motivating reason. On the second, Justified Evidence is true *not* because we have assets that help turn propositions into evidence by virtue of making our attitudes towards them good in some way. Justified Evidence is true because of the liability we take on whenever we take on the commitment that is distinctive of belief. If a belief cannot deliver what it's supposed to (i.e., a potential motivating reason), it couldn't possibly have some other features that make it appropriate to take on. I think that the second reading is preferable because it supports a view on which facts about the kinds of things that can be reasons or evidence place constraints on the kinds of beliefs that can be justified.[30]

However we understand Justified Evidence, we can use norms like Remain to get an argument for (O1) and (O2). If a putatively justified belief didn't provide us with evidence, it cannot properly figure in theoretical reasoning because it would violate Remain. If you violate Remain, however, it's hard to see how the belief could be justified. It's a belief that you shouldn't have. You only conform to Remain if *p* is part of your evidence. So, your belief in *p* is justified only if *p* is part of your evidence. If E=K is true and *p* is part of your evidence, you know *p*. Thus, we get our argument for objectivism about norms and justification.

To resist this argument the subjectivists about justification must challenge Justified Evidence or E=K. Dialectically, it doesn't make much sense for them to challenge Justified Evidence. The subjectivist about norms would not want to say that satisfying the norms that tell us when it's appropriate to update on a piece of evidence require *more* than justified belief. If they did, there would be pairs of cases where a subject justifiably believes a proposition in both but would only be permitted to update on that proposition in one. The requirement to refrain from updating in the one case would have to be sufficiently obscure to the subject so that they didn't threaten the

---

[30] When we try to run things the other direction, we'll run into difficulties that we'll discuss below.

justificatory status of her beliefs but nevertheless sufficiently 'internal' to be binding so that the subject shouldn't draw the conclusions she knows would have to be true given the things that she justifiably believed. It's hard to see how this view could possibly be attractive.

The subjectivist will say that internal duplicates will do equally well conforming to norms like Revise and Remain. They wouldn't want to say, then, that owing to differences in their knowledge about the external world it is possible for pairs of internally indiscernible subjects to draw the same inferences and reason in just the same ways to differ in that only one of them satisfies Remain. It wouldn't be wise for the subjectivist to insist that all of our evidence consists of propositions about our subjective mental life. If they did this and didn't also challenge E=K, their view would lead to an unpalatable form of skepticism. They shouldn't deny that whatever we know belongs to our evidence. If they challenged this, they'd have a hard time accommodating Justified Evidence. Their best strategy would be to challenge the idea that propositional knowledge is required for propositional evidence. Specifically, their best strategy is to challenge the idea that all propositional evidence consists of facts or true propositions.

This is precisely what some of Williamson's critics have done. Comesana thinks that views that incorporate Justified Evidence and E=K run into trouble in cases of ampliative inference because this view rules out the possibility of false evidence:

> According to Williamson, a subject is justified in believing a proposition only if she knows it. Suppose now that a subject knows a proposition *p* because she infers it from another proposition *q* (which doesn't entail *p*). In that case, according to Williamson the subject's total evidence entails *p*, for the simple reason that it includes *p*. But let us suppose that the evidence on the basis of which she believes *p* (namely, *q*) does not entail *p*: the inference from *q* to *p* is ampliative. In that case, it is possible for another subject to also know *q* and to similarly come to believe *p* on its basis even when *p* is false. This subject, of course, does not know that *p*, for the simple reason that *p* is false in her situation. And yet her evidence for believing that *p* is exactly the same as the evidence that the first subject, who knows that *p*, has [or had prior to learning *p*]. Williamson's position has the unfortunate consequence that the first subject is justified in believing that *p* whereas the second is not (Comesana 2015: 86).

Is this consequence unfortunate? It doesn't seem so to me. To show that it's unfortunate, it would be helpful to have some rationale for thinking that the inferences yield beliefs with the same justificatory status and/or the same evidence. It isn't clear why we should accept this crucial claim about the justificatory status of the inferential beliefs. An objectivist like Williamson can insist that the beliefs in the bad case aren't justified either because they violate Only Knowledge or because that's jut an implication of

Justified Evidence and an independently motivated account of what our evidence consists of.

It might seem that Williamson is on weak ground because this response would force him to say that there can be justificatory differences without some difference in the reasons that bear on whether to hold the relevant beliefs. Before the subject draws the relevant inferences, the subject's evidence for *p* could have been the same, so shouldn't they be in the same position to justifiably draw the same inferences? To this, Williamson has an answer. It's important to distinguish cases where subjects have the same available evidence for an inferential belief and cases where the same reasons bear on whether to draw the relevant inference. On Williamson's view, there *is* a difference in the reasons that bear on whether to believe *p*. In the bad case, the subject would violate Only Knowledge. Only Knowledge tells us that a subject in such a case shouldn't believe *p*. Because 'should' and 'ought' imply 'reason', Only Knowledge tells us that there's a reason for the subject in the bad case to refrain from believing *p*. There is no such reason in the good case because that is a case where the subject wouldn't violate Only Knowledge by coming to believe *p*. In light of this difference, the difference in justificatory standing is just what we'd expect.[31]

Instead of focusing directly on the justificatory status of the inferential beliefs produced in the good and bad cases, Comesana could try to bracket the question about justificatory status and focus just on the question about whether the same inferences would provide the subject with the same evidence. It's going to be difficult to establish that the thinkers acquire the same evidence in these cases. One way to test proposals about when subjects possess reasons is to think about cases where the relevant propositions figure in reasoning. Suppose Agnes comes to believe *p* via ampliative inference but forgets her original grounds. She retains her belief in *p* and sees that *r* is true if *p* is. She infers *r*.

If *p* were among Agnes' reasons in the good and bad case, we'd expect that this reason could be her reason for thinking, feeling, and doing things. Consider three possible views about Agnes' evidence or reasons:

> Same: In the bad case, Agnes' reason for believing *r* is that *p*. It is the same reason as in the bad case.
>
> Different: In the bad case, Agnes' reason for believing *r* is not that *p*. There was something that was Agnes' reason, but it isn't that *p*.[32]
>
> Neither: There was nothing that was Agnes' reason for believing *r* in the bad case. There are reasons why Agnes

---

[31] Some people find it intuitive that the beliefs produced in the right error cases are justified. This intuition has been addressed elsewhere. See Williamson (forthcoming). If that intuition carries the day, further arguments like the one we're considering aren't needed. If, however, that intuition doesn't undermine the view that justification requires knowledge, the present argument doesn't seem to create any serious difficulties for Williamson's views about evidence, justification, and knowledge.

[32] See Dutant (forthcoming), Hornsby (2007), and Lord (forthcoming).

believed what she did, but these reasons weren't Agnes' reasons.[33]

If someone is going to argue that the relevant falsehood (i.e., p) has to be part of Agnes' evidence in the bad case, it would be helpful to show that Same does better than Different and Neither in identifying Agnes' reasons in these cases. Unfortunately, Same hasn't been carefully defended in the literature and all the linguistic evidence suggests that Different and Neither better account for the apparent entailments between ascriptions of reasons and knowledge.[34] Since someone who accepts Different or Neither wouldn't have any good reason to think that Agnes' evidence would have to include the falsehood, I don't yet see why we should part company with Williamson in insisting that the relevant falsehood that Agnes comes to believe via ampliative inference never constitutes a piece of evidence.

There is a further difficulty that arises for the kind of view that Comesana has proposed. On this view (the 'target view'), evidence or potential motivating reasons (i.e., the things that could be the subject's reason for thinking, feeling, or doing something) are propositions that can be true or false that get to be part of the subject's evidence by virtue of standing in some epistemically significant relation that we can bear to truths and falsehoods. Justification or justified belief, as ordinarily understood, is supposed to play the role of the relevant possession condition. Difficulties arise for the target view when we think about the properties of processes like conditionalization and the properties of 'ought'.

Let's suppose, in keeping with the target view's understanding of Justified Evidence, that a proposition will belong to Agnes' evidence if rationally or justifiably believed and that this means that falsehoods will sometimes constitute part of Agnes' evidence (e.g., when they appear to be true, when they are backed by strong evidence, or when they are the content of a belief produced by a fallible but sufficiently reliable process). Imagine that Agnes completes *Here I Am*, a carefully researched if somewhat self-indulgent autobiography. Because it is so carefully researched, Agnes has a well-founded belief that corresponds to each claim in the body of the book. Thus, it would seem that she should have a justified belief that corresponds to each claim in the body of the book. After completing the body, her fact-checker tells her that she found precisely one error in the manuscript. Alas, her fact-checker dies suddenly and unexpectedly before she can reveal the error's location. After searching in vain for the error, Agnes fails to find the error. She notes in the book's preface that the book contains an error and it is sent to press.

As I understand the story, Agnes satisfies the conditions that orthodox accounts of justification impose, so we should say that she has a justified belief that corresponds to each claim in the body of her book (i.e., p_1-$p_{1,000,000}$) *and* in the preface. It would seem that the target view would allow that (P1)-(P3) could all be true:

---

[33] See Alvarez (2010).

[34] For discussion of the linguistic evidence, see Hawthorne and Magidor (forthcoming), Littlejohn (2012), and Unger (1975). For defenses of a truth-requirement on evidence or reasons, see Alvarez (2010), Mantel (2013), and Mitova (forthcoming).

P1. Agnes justifiably believes p_1, Agnes justifiably believes p_2, ..., and Agnes rationally believes $p_{1,000,000}$.
P2. Agnes justifiably believes $\sim(p_1 \,\&\, p_2 \,\&\, \ldots \,\&\, p_{1,000,000})$.
P3. Whatever Agnes justifiably believes belongs to her evidence.
P4. Agnes ought to update on p_1, Agnes ought to update on p_2, ..., and Agnes ought to update on $p_{1,000,000}$. [(P1), (P3), and Revise]
P5. Agnes ought to update on $\sim(p_1 \,\&\, p_2 \,\&\, \ldots \,\&\, p_{1,000,000})$. [(P2), (P3), and Revise]
P6. Agnes ought to (update on p_1 & update on p_2 & ... & update on $p_{1,000,000}$). [(P4), Agglomeration]
P7. If Agnes ought to update on $\sim(p_1 \,\&\, p_2 \,\&\, \ldots \,\&\, p_{1,000,000})$, Agnes ought not update on $(p_1 \,\&\, p_2 \,\&\, \ldots \,\&\, p_{1,000,000})$.
P8. Agnes ought not update on $(p_1 \,\&\, p_2 \,\&\, \ldots \,\&\, p_{1,000,000})$. [(P5), (P7)]
P9. If Agnes ought to: (update on p_1, update on p_2, ..., update on $p_{1,000,000}$), it is not the case that Agnes ought not update on $(p_1 \,\&\, p_2 \,\&\, \ldots \,\&\, p_{1,000,000})$.
P10. It is not the case that Agnes ought not update on $(p_1 \,\&\, p_2 \,\&\, \ldots \,\&\, p_{1,000,000})$. [(P6), (P9)]

Because (P10) contradicts (P8), it seems that we either have to abandon the target view or reject some other premise in the reasoning that leads to this contradiction.

If you accept the kind of objectivist view that I've defended (i.e., that incorporates Justified Evidence and E=K), you know that it couldn't be that (P1) and (P2) are correct, but I don't see how a proponent of the target view could plausibly reject both claims. [35] Justified Evidence is part of the target view, so its defenders cannot reject (P3). Revise generates a number of 'ought' claims about updating and they support (P6) because of agglomeration. As for (P7), this is the plausible claim that if you ought to update on p, you should not (also) update on its negation. As for (P9), this is plausible given that a process like conditionalization is cumulative. The result of conditionalizing on the conjuncts is the same as the result of conditionalizing on the conjunctions. It is hard to see how there could be a requirement to update on p, say, a requirement to update on q, and *also* a requirement to refrain from updating on the

[35] Some writers would argue that we shouldn't accept (P1) and (P2). Some would argue this because they think that there's some kind of defeater operative that prevents them from both being true (e.g., Ryan (1991)). For arguments that we shouldn't deny that it's possible for (P1) and (P2) to be true, see Easwaran and Fitelson (2015) and Worsnip (2015). Some would try to show that we don't have the right evidence or grounds for (P2) to be true (e.g., Smith (2016)). Using testimony, it's hard to see how denying (P2) could be a viable strategy.

conjunction when we know in advance that the result of the one-step update and two-step update has to be the same.

The trouble with the target view is that when it drops the truth-requirement on evidence, it does so because it offers some non-factive epistemic relation to propositions as sufficient for the possession of evidence. This, in turn, prevents the proponents of the target view from allowing for inconsistent bodies of evidence. Once we have that, we should not then let the evidence possessed determine what we 'ought' or 'ought not' do. If we did that, we couldn't vindicate the agglomeration of 'ought'. If evidence is the kind of thing that we ought to be guided by and ought to compel us to update, it would have to consist of consistent sets of propositions. The best way to achieve that, theoretically, is to adopt a factive account of evidence. Once we have that, Justified Evidence vindicates the objectivist account of justification and norms defended here. The reason that justified beliefs have to be true, remember, is not that they need some maximal independent support, but because only beliefs that contribute truths to rational processes meet the standard that distinguishes the beliefs that can do what they're supposed to do from those that cannot.

3. Subjectivism about Justification

Recall the argument for objectivism about justification:

> *An Argument for Objectivism*
>
> O1. Your beliefs are justified iff you violate no epistemic norms in holding these beliefs.
>
> O2. You violate no epistemic norms in holding these beliefs iff these beliefs constitute knowledge.
>
> OC. Your beliefs are justified iff they constitute knowledge.

In the previous sections, I've offered a transcendental argument that Only Truth governs belief and an evidentialist argument that Only Knowledge governs belief. The transcendental argument, if sound, shows that the standard objections to (O2) miss their mark. The evidentialist argument, if sound, shows that you cannot believe what you should unless you believe only what you know to be true. The upshot is that the false, justified belief is a myth.

Most epistemologists reject (OC), but it isn't always clear where they think my argument for objectivism goes wrong. Huemer (forthcoming) suggests that the problem with arguments like my argument for (OC) is not so much that they assume that there are objectivist norms, but with the further assumption that these norms might tell us something about justification. He agrees that there's a sense in which we 'ought' to conform to epistemic norms like Only Truth and could not justifiably believe what we 'ought' not believe, but he thinks that this lends no support to objectivism about justification.

Huemer thinks that some of our normative vocabulary is ambiguous. Terms like 'ought', 'should', and 'norm' can be read objectively (so that external facts that the subject might be ignorant of can matter to the truth of claims about what a subject 'ought' to do, think, or feel) and also can be read subjectively (so that they do not matter

to the truth of claims about what a subject 'ought' to do, think, or feel). When it comes to justification, justification is entirely a matter of conforming to subjective norms.

What should we take from this? When it comes to debates between the objectivist and subjectivist about norms, he thinks we should see that these debates aren't substantive:

> Now suppose someone asks, "*Does what a person should do* depend upon external facts of which that person may be unaware, or does it only depend upon information available from the subject's own point of view?" This would be a misguided question. There is no deep issue there; there is only the perfectly shallow, semantic question of whether you want to use the internal sense or the external sense of the word "should". Both senses, as far as I can tell, are established in ordinary English; neither is more correct than the other (Huemer: forthcoming: 7).

Because these debates rest on the mistaken presupposition that there is a substantive debate to be had, they should end. The objectivist and subjectivist about norms should see that their opponent's views get something important right. Live and let live.

When it comes to debates about justification, however, things are quite different. He thinks these debates shouldn't continue because (P3) is true only on its subjective reading:

> The truth of this thesis [i.e., that justification is entirely a matter of conforming to subjective norms], by the way, strikes me as just an obvious semantic point. Unlike the word "should", I think the word "justified" is not ambiguous; rather, it has only the internal meaning in standard English (Huemer forthcoming: 8).

Even if the objectivist is right about epistemic norms, there is nothing that the objectivist about justification gets right. When it comes to justification, it only matters whether we conform to subjectivist norms. Only the subjectivist about justification has a sensible view and that's why this debate should end. Live and let die.

In discussions of the subjective 'ought' and the (alleged) ambiguity of our normative language, philosophers tend to exhibit one of two tendencies. They tend to either be dividers or debaters.[36] The divider thinks that pointing to the ambiguity will help us see that there is no need to debate a certain issue because both parties to the debate are right about something. The parties to the debate are only mistaken in that they fail to see their opponent's insights. The debater thinks that introducing the subjective 'ought' helps us see that the objectivist is wrong about some normative notion or notions. As I mentioned earlier, it's long seemed strange to me that so many epistemologists would agree that there are norms like Only Truth or Only Knowledge that governs belief when these same epistemologists would insist that false beliefs can be perfectly justified. Maybe Huemer has his finger on something important here. Should we be dividers about norms and debaters about justification? Specifically, should we

[36] See Sepielli (forthcoming).

adopt his subjectivism about justification and accept that debates between objectivists and subjectivists about norms are non-substantive?

I think not. We've seen above why Huemer is so confident that the subjectivists are right about justification. I don't here want to survey the possible arguments for subjectivism about justification because most have already been addressed in the literature and I've gestured at some reasons why I don't think that Huemer's most recent argument will advance the subjectivist cause. Instead, I'd like to focus on the proposal that we should see these debates between the objectivist and subjectivist about norms as non-substantive. When we see why there is room for substantive disagreement here, we should also see why it isn't obvious that the subjectivist is right about justification. When these norms come into conflict, there is a substantive question about which norms to follow, a question that is clearly about whether we could act or believe with justification if we conformed to one set of norms while violating another.

If the dividers can show that debates between objectivists and subjectivists about the formulation of norms is misguided or not substantive, this would be because the subjectivist can vindicate everything the objectivist says so that there is no remaining substantive disagreement. Can it be done? I have my doubts. Subjective and objective norms can come into conflict in such a way that it looks as if subjectivists will simply have to deny that objectivist norms are genuine.

Consider the kind of disagreement that the divider would take to be non-substantive. One such disagreement would be about the objectivist view of punishment (e.g., a view that incorporates Only the Guilty and Any Known Offender) and one that replaces these norms with some subjectivist alternative:

> You should sentence the accused iff their guilt is beyond all reasonable doubt (Beyond Reasonable Doubt).

According to Only the Guilty and All Known Offenders, Agnes would have done nothing wrong if, say, she sentenced 1,000,000 persons standing trial provided that she knew of the accused to be guilty. Assuming, as we will, that Agnes believed the accused to be guilty beyond any reasonable doubt if she knew them to be guilty, the subjectivists who accept Beyond Reasonable Doubt would agree that Agnes did no wrong by sentencing these 1,000,000 persons for the crimes they committed.

The interesting issues don't arise when everyone agrees that Agnes got it all right. They arise in the cases where it seems to someone that Agnes got something wrong. Let's imagine again that Agnes serves on 1,000,000 juries and that the evidence in each case leaves no room for reasonable doubt. Thus, we get:

> 1. Agnes subjectively ought to convict in case 1.
> 2. Agnes subjectively ought to convict in case 2.
> 3. Agnes subjectively ought to convict in case 3.
> ...
> 1,000,000. Agnes subjectively ought to convict in case 1,000,000.

Suppose that after her 1,000,000th case is brought to a close Agnes learns that in one of her cases the accused was framed. She thus learns:

> 1,000,001. Agnes objectively ought not: (convict in case 1, convict in case 2, ... convict in case 1,000,000).

If Agnes knows that she objectively ought not do something, it seems plausible that she subjectively ought not do that thing:

> 1,000,002. Agnes subjectively ought not: (convict in case 1, convict in case 2, ... convict in case 1,000,000).

If the subjective 'ought' acts anything like the 'ought' of standard deontic logic, we can use agglomeration and (1)-(1,000,000) to give us this:

> 1,000,003. Agnes subjectively ought to: (convict in case 1, convict in case 2, ... convict in case 1,000,000).

And if the subjective 'ought' acts anything like the 'ought' of standard deontic logic (1,000,002) gives us:

> 1,000,004. It is not the case that Agnes subjectively ought to: (convict in case 1, convict in case 2, ... convict in case 1,000,000).

Since (1,000,003) and (1,000,004) are incompatible, something has to give. It seems we have four options:

> Option 1: Deny at least one of (1)-(1,000,000) and deny Beyond Reasonable Doubt.
>
> Option 2: Deny (1,000,001) and deny Only the Guilty.
>
> Option 3: Deny that the subjective 'ought' acts anything like the 'ought' of standard deontic logic.
>
> Option 4: Deny the bridge principle that states that you subjectively should not do what you know you objectively should not do.

Dividers want to show that the disagreement between the subjectivists and objectivists isn't substantive, so they cannot avoid the contradiction by means of Option 1 or 2. If they deny Only the Guilty or Reasonable Doubt because of some (apparent) tension between them, they would enter into just the debate that they wish we'd all stay out of.

The third option is problematic for dividers. Suppose we've established (1)-(1,000,000) and (1,000,002) and look to avoid the contradiction by denying agglomeration. It seems that, barring cases of duress, a subject who fails to do what she subjectively ought to could be held responsible for that failure or subject to criticism. We shouldn't blame or criticize someone for carrying out all of the acts described in (1)-(1,000,000) if we'd blame or criticize them for failing to carry them all out. (Does it make sense to say that someone would avoid criticism iff they sentenced in each case but would not avoid criticism if they sentenced everyone?)

The fourth option is the last option for the dividers. It looks as if they have good grounds for challenging the bridge principle.[37] Consider the mineshaft case.[38] Ten miners are working together underground, but Agnes doesn't know whether they are in Shaft A or Shaft B. She knows that it's raining and that if she does nothing all will die at the bottom of the mine. She has three options: block shaft A, block shaft B, partially

---

[37] See Wedgwood (2013) for discussion. Wedgwood offers a revision of the Enkratic Requirement that is supposed to avoid the difficulties that arise for this bridge principle. The principle holds only when Φ is a fine-grained option. In our case, it is not.

[38] Discussed in Parfit (2011). The case is a variant on Regan (1980) and Jackson (1991).

block both shafts. If she completely blocks the opening to A, there will be enough air inside for 10 to live but B will flood and kill anyone in that shaft. If she blocks B completely, there will be enough air inside for 10 to live, but A will flood and kill anyone in that shaft. If she partially blocks both openings, she will be guaranteed to save 9 miners wherever they happen to be but 1 will be killed. In this sort of case, Agnes knows that she objectively ought to either block Shaft A or B, but doesn't know which shaft it is. Thus, she knows that she objectively ought not partially close the shafts. Still, it seems to many commentators that Agnes subjectively ought to partially close the shafts. This suggests that the bridge principle is mistaken.

Wedgwood notes (rightly, I think) that the Enkratic Requirement is a genuine requirement only when the subject's beliefs pertain to fine-grained options (e.g., blocking shaft A completely). Thus, it's plausible that the bridge principle is only superficially similar to the Enkratic Requirement and shouldn't be accepted without some very compelling further argument. If dividers reject the bridge principle, are they out of the woods?

No, not quite. If they deny the bridge principle, this allows them to reject (1,000,002). They could say that while there is a sense in which Agnes succeeded in doing everything she should do, there is a second sense in which she failed to do what she should do. Suppose we say that. If Agnes then asked whether she should make reparations, what would the divider say?

The divider should ask the subjectivists what they'd say. The subjectivist might offer a version of Repair, suitably modified to fit with their theoretical orientation. They might say that upon discovering that she sentenced an innocent person to prison, Agnes has a reparative duty to right some past wrong. On this view, it turns out to be true, as the objectivist says, that Agnes failed to do what she ought to do, but crucially the truth of this depends upon the fact that this new information came to light. Alternatively, the subjectivist could insist that since Agnes did no wrong, violated no norm, and never failed to do what she ought to do. If so, no reparation is called for. The subjectivist would say that the objective facts that figure in objectivist accounts of norm have to be excluded from the framework so that they don't figure in our account of rights, obligations, or duties to repair or make reparation.[39] The accused would have the right to a fair trial, but not to escape sentencing if innocent. Those responsible for sentencing would have the responsibility to provide an adequate trial, but no responsibility to discharge reparative duties in response to harms that result from trials that were conducted as they should have been.

The problem with the first response is that it leads us right back to the problem we were trying to avoid by denying the bridge principle. If the subjectivist says that Agnes has a reparative duty because she failed to do something she ought to do by virtue of sentencing the accused in each case, the subjectivist is offering us an account

---

[39] The most completely worked out version of this view is found in Zimmerman (2008). He is a subjectivist in the sense that he thinks that what a subject ought to do depends upon her evidence and not upon any further facts that don't supervene upon that evidence. He is not a divider. For further defenses of the subjectivist debater, see Kiesewetter (forthcoming) and Lord (forthcoming b).

on which (1,000,004) comes out as true. As we're assuming that the subjectivist wants to offer an account of 'ought' on which 'ought' satisfies agglomeration, this proposal won't do. The truth of (1,000,004) would force us to give up at least one of (1)-(1,000,000) but the subjectivist has no good grounds for giving up any of these claims. The new information doesn't change the fact that Agnes had adequate evidence at the time of sentencing and doesn't change the fact that she still has adequate evidence for believing of each conviction that it was appropriate.

The second subjectivist response causes trouble for the divider. It looks as if the subjectivist who offers this response is now engaged in a genuine, substantive disagreement with the objectivist. Even if Agnes learned the identity of the individual who was framed, our subjectivist would insist that the objectivists would be wrong to say that Agnes did something that violated this person's rights, that she failed in her obligations to the accused, and wrong to say that Agnes has a reparative duty to right some past wrong. At just this point it looks like there is a substantive question that we need to settle. Upon discovering the identity of the innocent person she convicted, would Agnes' decision not to make reparations be justified? If you say that it would, you are taking up a substantive view and siding with a kind of subjectivist debater. If you say that it would not, you are taking up a substantive view and siding with the objectivist who thinks that justification turns on conforming to objectivist norms. It seems pretty obvious to me that Agnes would have no justification for failing to try to make reparations, but this interesting fact about justification would appear to be connected to the past violation of an objectivist norm (Only the Guilty), not a subjectivist norm. Thus, the debate turns out to be substantive and the objectivist view about justification turns out to have an important virtue that the subjectivist view lacks. The objectivist (and the objectivist alone?) has a coherent position that vindicates important intuitions about reparation.

4. Conclusion

I have defended two objectivist views, one about norms and another about justification. There is no simple case to be made for either objectivist view, but it seems that the objectivist framework best fits with some plausible claims about the kinds of obligations we're under and with some popular claims about the normative significance of our evidence. What emerges from this is a simple idea that will undoubtedly require further defense. For belief to play its role well, it has to provide us with reasons and put us in touch with the part of reality that consists of facts. A belief only does this when it constitutes knowledge. This is why Only Knowledge governs belief. There is little theoretical gain to be had by insisting that this point, if correct, tells us nothing about justification. Justification loses all of its theoretical significance if we characterize it as a status that a belief can have or fail to have quite apart from whether that belief conforms to all the norms that govern belief. This is why objectivism about norms and justification stands or falls together.

References

Alvarez, Maria. 2010. *Kinds of Reasons*. Oxford University Press.

Arpaly, Nomy. 2002. *Unprincipled Virtue*. Oxford University Press.

Benton, Matthew. 2011. Two More for the Knowledge Account of Assertion. *Analysis* 71: 684-687.

Bird, Alexander. 2007. Justified Judging. *Philosophy and Phenomenological Research* 74: 81-110.

Boghossian, Paul. 2008. *Content and Justification*. Oxford University Press.

Brogaard, Berit. 2016. Staying Indoors: How Phenomenal Dogmatism Solves the Skeptical Problem without Going Externalist. In B. Coppenger and M. Bergmann (ed.), *Intellectual Assurance*. Oxford University Press.

Brown, Jessica. 2008. Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning. *Nous* 42: 167-189.

Coates, Allen. 2012. Rational Epistemic Akrasia. *American Philosophical Quarterly* 49: 113-24.

Cohen, Stewart. 1984. Justification and Truth. *Philosophical Studies* 46: 279-95.

Comesana, Juan. 2016. Falsehood and Entailment. *Philosophical Perspectives* 29: 82-94.

Conee, Earl and Richard Feldman. 2004. *Evidentialism*. Oxford University Press.

Dougherty, Trent. 2011. Introduction. In T. Dougherty (ed.), *Evidentialism and its Discontents*. Oxford University Press.

Dutant, Julien. Forthcoming. Knowledge-Based Decision Theory and the New Evil Demon Problem. In F. Dorsch and J. Dutant (ed.), *The New Evil Demon*. Oxford University Press.

Easwaran, Kenny and Branden Fitelson. 2015. Accuracy, Coherence, and Evidence. In T. Gendler and J. Hawthorne (ed.) *Oxford Studies in Epistemology*. Oxford University Press.

Enoch, David. 2009. Wouldn't It Be Nice if p, Therefore, p (for a moral p). *Utilitas* 21: 222-4.

Fantl, Jeremy and Matt McGrath. 2009. *Knowledge in an Uncertain World*. Oxford University Press.

Feldman, Richard. 1988. Subjective and Objective Justification in Ethics and Epistemology. *the Monist* 71 405-19.

Foley, Richard. 2001. The Foundational Role of Epistemology in a General Theory of Rationality. In A. Fairweather and L. Zagzebski (ed.), *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Oxford: Oxford University Press.

Gibbons, John. 2010. Things That Make Things Reasonable. *Philosophy and Phenomenological Research* 81: 335–61

Goldman, Alvin. 1986. *Epistemology and Cognition*. Harvard University Press.

Greco, Daniel. 2014. A Puzzle about Epistemic Akrasia. *Philosophical Studies* 167: 201-19.

Hawthorne, John and Ofra Magidor. Forthcoming. Reflections on Reasons. . In D. Star (ed.), *Oxford Handbook of Reasons and Normativity*. Oxford University Press.

Hornsby, Jennifer. 2007. Knowledge in Action. In A. Leist (ed.), *Action in Context*. De Gruyter, pp. 285–302.

Huemer, Michael. 2007. Compassionate Phenomenal Conservatism. *Philosophy and Phenomenological Research* 74: 30-55.

Huemer, Michael. Forthcoming. Defending a Brain in a Vat. In F. Dorsch and J. Dutant (ed.), *The New Evil Demon*. Oxford University Press.

Hughes, Nick. 2014. Is Knowledge the Ability to ϕ for the Reason that P? *Episteme* 11: 457-62.

Hyman, John. 1999. How Knowledge Works. *Philosophical Quarterly* 49: 433-51.

Ichikawa, Jonathan and Ben Jarvis. 2013. *The Rules of Thought*. Oxford University Press.

Jackson, Frank. 1991. Decision-Theoretic Consequentialism and the Nearest and Dearest Objection. *Ethics* 101: 461-82.

Kiesewetter, Benjamin. Forthcoming. You ought to ϕ only if you may believe that you ought to ϕ. *Philosophical Quarterly*.

Lasonen-Aarnio, Maria. MS. Enkrasia or Evidentialism: Learning to Love Mismatch.

Littlejohn, Clayton. 2011. Defeating Phenomenal Conservatism. *Analytic Philosophy* 52: 35-48.

Littlejohn, Clayton. 2013. The Russellian Retreat. *Proceedings of the Aristotelian Society* 113: 293-320.

Littlejohn, Clayton. 2014. The Unity of Reason. In C. Littlejohn and J. Turri (ed.), *Epistemic Norms*. Oxford University Press.

Locke, D. 2015. Knowledge, Explanation, and Motivation. *American Philosophical Quarterly* 52: 215-33.

Lord, Errol. Forthcoming. Defeating the Externalist's Demons. In F. Dorsch and J. Dutant (ed.), *The New Evil Demon*. Oxford University Press.

Lord, Errol. Forthcoming b. Acting for the Right Reasons, Abilities, and Obligation. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*. Oxford University Press.

McKinnon, Rachel. 2013. The Supportive Reasons Norm of Assertion. *American Philosophical Quarterly* 50: 121-35.

Mantel, Susanne. 2013. Acting for Reasons, Apt Action, and Knowledge. *Synthese* 190: 3865-3888.

McGlynn, Aidan. 2014. *Knowledge First?* Palgrave Macmillan.

Mitova, Veli. Forthcoming. Truthy Psychologism about Evidence. *Philosophical Studies*.

Nolfi, Kate. 2015. How to be a Normativist about the Nature of Belief. *Pacific Philosophical Quarterly* 96: 181-204.

Parfit, Derek. 2011. *On What Matters*. Volume I. Oxford University Press.

Preston-Roedder, Ryan. 2014. A Better World. *Philosophical Studies* 168: 629-44.

Regan, Don. 1980. *Utilitarianism and Cooperation*. Oxford University Press.

Ryan, Sharon. 1991. The Preface Paradox. *Philosophical Studies* 64: 293-307.

Sayre-McCord, Geoff. MS. On a Theory of a Better Morality.

Sepielli, Andrew. Forthcoming. Subjective and Objective Reasons. In D. Star (ed.), *Oxford Handbook of Reasons and Normativity*. Oxford University Press.

Smith, Martin. 2016. *Between Probability and Certainty*. Oxford University Press.

Smithies, Declan. 2011. The Normative Role of Knowledge. *Nous* 46: 265-88.

Steglich-Petersen, Asbjorn. 2013. Truth as the Aim of Justification. In T. Chan (ed.), *The Aim of Belief.* Oxford University Press.

Sutton, Jonathan. 2007. *Without Justification*. MIT University Press.
Tucker, Chris. 2010. Why Open-Minded People Should Endorse Dogmatism. *Philosophical Perspectives* 24: 529-45.
Turri, John. 2011. The Express Knowledge Account of Assertion. *Australasian Journal of Philosophy* 89: 37-45.
Unger, Peter. 1975. *Ignorance*. Oxford University Press.
Way, Jonathan and Daniel Whiting. Forthcoming. If You Justifiably Believe that You Ought to ΦYou Ought to Φ *Philosophical Studies*.

Wedgwood, Ralph. 2002. Internalism Explained. *Philosophy and Phenomenological Research* 65: 349-69.
Wedgwood, Ralph. 2013. The Right Thing to Believe. In T. Chan (ed.), *The Aim of Belief.* Oxford University Press.
Wedgwood, Ralph. 2013b. Akrasia and Uncertainty. *Organon F* 20: 484-506.
Whiting, Daniel. 2013. Nothing but the Truth: On the Norms and Aims of Belief. In T. Chan (ed.), *The Aim of Belief.* Oxford University Press.
Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.
Williamson, Timothy. Forthcoming. Justifications, Excuses, and Skeptical Scenarios. In F. Dorsch and J. Dutant (ed.), *The New Evil Demon*. Oxford University Press.
Worsnip, Alex. 2015. Belief, Credence, and the Preface Paradox. *Australasian Journal of Philosophy* 94: 549-562.
Zimmerman, Michael. 2008. *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge University Press.