Epistemological Motivations for Anti-realism*

Billy Dunaway

University of Missouri–St. Louis

forthcoming in *Philosophical Studies*

Does anti-realism about a domain explain how we can know facts about the domain? Often one of the chief motivations for denying realism about a subject-matter is *epistemological:* that knowledge, or justification, will be difficult to come by if realism is correct. But this epistemological problem for realism only benefits the anti-realist if her view has additional resources to explain the epistemological desiderata.

Here I will outline a framework for answering this question, and sound a pessimistic note on the possibility of showing anti-realism to be preferable to realism on epistemological grounds. I will focus on anti-realism about normativity to make this point.[1] There has been a recent emphasis on the epistemological problems realism about the normative faces, and a corresponding assumption that, if these problems are genuine, they constitute a prima facie motivation

*Thanks to an anonymous referee, as well as audiences at the University of Oxford Moral Philosophy seminar, University of Sydney, Australian National University, the St. Louis Ethics Workshop, the Higher Seminar in Theoretical Philosophy at Uppsala University, Saint Louis University, and the London Institute of Philosophy Language, Epistemology, and Metaphysics Seminar for helpful discussion of previous versions of this paper.

[1]Recent work by Berker (2014), Setiya (2012), and Tropman (2014) has also raised this question.

for adopting anti-realism.[2] I aim to show that the most simple and natural way of making good on the claim that anti-realism about the normative solves epistemological challenges in fact fails. This gives reason to be cautions about blanket assumptions of any distinctive epistemological features of anti-realism.

This failure of anti-realism to be in a different position epistemologically from realism has an easy diagnosis. The core of anti-realism is a commitment to grounding normative facts in normative belief and practice—and, thereby forging a close tie between our normative beliefs and the facts that would make them true. This creates the appearance that our beliefs about normative facts cannot, on the anti-realist view, be mistaken except in perhaps very weird cases. So it seems like a small and manageable promise to turn this feature into an epistemological difference with realism, which allows normative facts to diverge widely from our beliefs about them, since the facts (for the realist) are *independent* of what we think about them. But the promise is in tension with another aspect of anti-realist theorizing, which is *idealization:* normative facts are not necessarily tied to actual belief and practice, but rather to belief and practice in certain idealized circumstances. Each aspect of anti-realist theorizing is not new. But their joint epistemological consequences have been grossly under-appreciated.

My aim in this paper is not to show that *every* attempt to derive an epistemological advantage from an anti-realist starting point must fail. Instead, since the epistemology of anti-realism has been largely unexplored (as opposed to the

[2]See for example Mackie (1977), Harman (1986), and Street (2006) for different versions of an epistemological challenge to realism. Schafer (forthcoming) defends the realist from the most damning versions of these challenges, but concedes that anti-realists might have a better time explaining the presence of epistemically desirable features of normative belief.

epistemology of realism[3]), I will mostly focus on one very straightforward and natural way to understand the alleged connection between the metaphysics of anti-realism and the epistemological virtues that have been claimed for it. I do not wish to claim that this is the only way the anti-realist epistemological project might be realized. Rather I will make the case that it is one especially plausible route for the anti-realist to follow, before showing that it fails to deliver on its ambitions. In closing I will sketch how these problems will be faced by other ways to develop the anti-realist's epistemological claims.

1 Anti-realism

1.1 Constructivism

Why does anti-realism appear to have an epistemological advantage over realism? There are many ways to be an anti-realist. There is no guarantee that every version of anti-realism will fare the same in the epistemological realm. So we should focus only on versions which, have a prima facie case to be especially promising for epistemological reasons. (Whether they are promising in other respects is not a question I will focus on in detail.)

I will call the version of anti-realism I will focus on *Constructivist* because of similarities it bears to views that have been called versions of Constructivism in the literature. It is important to be clear, however, that this is less a conceptual claim about the nature of Constructivism or anti-realism than a label for a view which appears, for reasons I will make clear below, to be well-suited to do the

[3]Though again see Berker (2014), Setiya (2012), and Tropman (2014) for exceptions.

epistemological work that realism allegedly fails to do.

The form of anti-realism about normativity I will focus on holds that it is our normative *attitudes*—or, in the specific instance I will focus on, *beliefs*—that determine the normative facts. Again, it deserves emphasis that I am not legislating usage of the term 'anti-realism' or 'Constructivism' here; there certainly could be versions of Constructivist anti-realism that do not fit the schematic characterization I will use here.[4]

The guiding idea behind my semi-stipulative understanding of Constructivism is that what is distinctive about the view is a particular stance on the metaphysics of normativity. Here I will formulate the metaphysical view as a grounding thesis, in the form of a claim about the (non-normative) facts that ground, in the metaphysical sense, normative facts. Lots of meta-ethical views hold that normative facts are grounded in natural facts, of some kind.[5] But only the Constructivist view I am interested in here takes the grounding base for normative facts to involve natural facts that primarily concern non-normative beliefs.

Thus as I will use the term 'Constructivism', it labels a view that accepts the following claim:

> If it is a fact that one ought to ϕ in circumstance c, then this fact is grounded (at least in part) by one's believing that one ought to ϕ in c.

[4]See Southwood (Forthcoming), Street (2010) for conceptions of Constructivism that begin with a similar starting point with, diverge in other respects from, the conception I will be working with here.

[5]See for example Railton (1986) and Schroeder (2007) for examples of non-Constructivist naturalistic grounding claims. Here and throughout I use talk of 'grounding' to pick out a relation of metaphysical dependence. Beyond that, 'grounding' talk is neutral between different conceptions of the relation, as well as a metaphysical understanding in terms of determination or constitution instead.

This thesis specifies a set of facts that are *grounded*—normative facts of the form *one ought to* ϕ—and the facts which constitute part of the *grounding base* for normative facts—facts of the form *one believes that one ought to* ϕ. The grounding base contains *normative beliefs*, since they have a content which is specified partly in normative terms. Moreover on this (partially stipulated) understanding of Constructivism, it is a belief whose content is *identical* to the fact that it grounds.

While Constructivism as I am understanding it is a metaphysical claim, as it concerns the grounds of normative facts, it is this feature which makes it a promising candidate for resolving epistemological problems. Since the view claims that normative beliefs ground normative facts, these beliefs *ground their own truth.* That is, if I have the normative belief that I ought to ϕ, my belief can be a part of the grounds of the fact that I ought to ϕ. As it were, the belief makes itself true. It is this aspect of Constructivism that makes it especially promising as a route for explaining why the belief that I ought to ϕ is epistemologically privileged in some way.

Much more needs to be said to turn this observation into a convincing argument that the view has epistemological advantages that other views lack. (I will spell out one reason for optimism on this point in the next sections.) But even without going into the details here, we can gesture at the strategy behind my semi-stipulative use of the term 'Constructivism': if a version of anti-realism with this feature cannot secure epistemological advantages, then other versions of anti-realism that relax the relationship between normative belief and the normative facts they ground will likely be even less promising on this front. If the grounding

base doesn't include beliefs that ground their own truth, then alleged epistemic advantages will be even harder to acquire.

Some further clarifications and qualifications are necessary at the outset.

First, Constructivism is primarily a metaphysical claim—it is primarily a claim about what grounds normative facts. It is tempting to draw epistemological consequences from this metaphysical view—I have gestured at why this is tempting by highlighting the dual role of normative belief on this account—but success is not guaranteed. While the metaphysical status of normative facts is secured by the Constructivist view by definition, the epistemological status of normative beliefs on the view is not trivial; the statement of Constructivism as I have given it does not by itself contain any epistemological terms. Any connections between the Constructivist's metaphysics and her epistemology will be substantive, and must be established by argument.

Second, the relationship between normative fact and normative belief on the Constructivist view is a relationship between *token* normative beliefs and *token* normative facts. That is, Sally's normative belief *that Sally ought to tell the truth now* on January 25, 2016 might ground the fact that Sally ought to tell the truth on January 25, 2016. But the corresponding relationship between general facts need not hold as well: it needn't be that the belief in the proposition that one ought to tell the truth (which may be instantiated by more than one person) grounds the general fact (if it is a fact) that one ought to tell the truth. This is because the token normative belief only *partly* grounds the general normative fact. The general normative fact that one ought (always) to tell the truth is a consequence of a set

of token normative facts: that one ought to tell the truth on January 25, 2016, and that one ought to tell the truth on January 26, 2016, etc. Roughly Constructivism can be thought of as a theory of the grounds of *atomic* normative facts; non-atomic normative facts are logical consequences of the atomic facts so constructed.[6]

The third—and most important—clarificatory point concerns the additional components of the grounding base for normative facts. As I have characterized Constructivism, normative beliefs are, at the very least, *part* of the grounding base for normative facts. This is not a fully specific characterization of a view, because it does not say whether there are additional components to the grounding base, and if so, what they are. We can begin, for the sake of illustration, by supposing that every token belief that one ought to ϕ fully grounds the that that one ought to ϕ. Call this the *Full Grounding View*. The Full Grounding View can be schematically represented as follows, where N is a normative fact, b_N is the belief in N, and $\Leftarrow$ represents the grounding relation:

$$N \Leftarrow \{b_N\}$$

According to this view, the grounds for a normative fact are *exhausted* by the fact that it is believed.

The Full Grounding View is incoherent. Some agents believe that they ought to lie at 2 pm on January 26, 2016, and they believe that they ought to tell the truth at 2 pm on January 26, 2016. There are various reasons why agents might have these contradictory beliefs. Some hold these beliefs due to simple irrationality:

[6]There are additional complications concerning what goes into the full grounding base for universal facts (such as that *one ought always to tell the truth*) and negative facts (such as that *it is not the case that one ought to lie today*). I will gloss over these complications here.

they do not recognize that both beliefs cannot be true, or do not give up one of the beliefs upon realizing this. Others might accept a (perhaps false) theory about why this combination of beliefs is rationally permissible. For instance they might buy into a global rejection of classical logic[7], or they might believe that such beliefs are the proper response to a case where there is a moral dilemma over whether to ϕ. Finally some might be the victims of unfortunate epistemic circumstances: they conceive of ϕ-ing in two ways, or have two "guises" for ϕ-ing. For instance, one might believe that they ought tell John that he is inconsiderate (because they believe that they ought to tell the truth) and at the same time believe that they ought not to tell John that he is inconsiderate (because they believe this not to be what John wants to hear).[8] Under conditions like these, it is very easy for an agent to *both* believe that she ought to ϕ *and* believe that she ought not to ϕ. She might do this without any awareness of rational impermissibility (or higher-order beliefs according to which believing contradictions is permissible), if she holds these beliefs under different guises of ϕ-ing.

The upshot is that the Full Grounding View is incoherent because it entails that it is sometimes true that one ought to ϕ, and true that it is not the case that one ought to ϕ.[9] The existence of *beliefs* in incompatible claims is understandable, and non-ideal agents sometimes have them. But the Full Grounding View turns these beliefs into contradictory normative facts. A fully plausible version of

[7]Williamson (2007: Ch. 4)

[8]See Salmon (1986) for a theory of guises and belief-formation. Here it is best to think of examples like this as involving an agent who thinks that they ought to do *this*, where the demonstrative refers to the truthful speech-act, and who also thinks that they ought to do *that*, where the demonstrative refers to the speech-act that John wants to hear.

[9]Here I assume the very plausible premise that—perhaps excepting for rare cases of difficult moral dilemmas—if you ought to ϕ, it follows that it is not the case that you ought not ϕ.

Constructivism will have to be more sophisticated than the Full Grounding View. But this version of Constructivism will be useful to keep in mind, when we turn to evaluating the epistemological properties of Constructivism.[10]

The solution to these problems is a move to a *Partial Grounding View:* one according to which normative belief partly grounds normative fact, but in addition holds that there are other conditions besides the relevant beliefs that constitute the *full* grounding base. These additional conditions must ensure ensure logical consistency (and other forms of coherence) among the normative facts they ground.

Existing versions of Constructivism in the literature make this very move: for instance Street says that normative truth is determined by which beliefs would *survive scrutiny*, holding other normative beliefs fixed. As a starting point, it is natural to take scrutiny survival to be a counterfactual feature of a belief. The belief B survives scrutiny only if it satisfies the following:

> SCRUTINY B is such that, if one were to become logically and probabilistically coherent, fully aware of one's other normative commitments, and fully factually informed, one would retain the belief B.

But plausibly more than this is needed to avoid incoherence in all cases: the normative belief b_N, which is a belief in the normative claim N might satisfy SCRUTINY—which is to say one would keep if one were to subject it to scrutiny

[10]Setiya (2012: 120) argues against something like the Full Grounding View on the grounds that, if we are reliable in forming normative beliefs, we will implausibly converge in what we think about normative matters. I will not rely on this style of criticism, because (as Setiya is aware) the implausible convergence can be avoided by relativizing normative facts to believers (see Schafer (2014)). Since I do not wish to take a stand on whether the Constructivist's commitments regarding relativity are plausible or not, I will leave this criticism to the side.

in light of one's other commitments in the relevant circumstances. But if one were instead to scrutinize one's belief $b_{N'}$ in the normative fact N' (which we can suppose is incompatible with N), one might also be disposed to keep it, too. So which normative fact obtains in this case? One option is to find a further grounding condition in addition to the relevant belief's satisfying SCRUTINY: Street says that in conflicts of this kind, the normative fact is determined by the values that are "most deeply" held by the agent. (Street, 2008: 234-5) On this version of the Partial Grounding View the normative fact N will be fully grounded in {b_N, b_N *satisfies* SCRUTINY, b_N *is most deeply held*}.

Another option is to allow the normative facts to be indeterminate once if there is no unique fact that is determined by the facts about which of an agent's normative beliefs satisfy SCRUTINY.[11] In cases where there is no determinate fact as to whether N obtains, this is because there is a belief in an incompatible normative fact N' which is such that the belief $b_{N'}$ *also* satisfies SCRUTINY. So, when the normative fact N determinately holds, this version of the Partial Grounding View holds that it will be fully grounded in {b_N, b_N *satisfies* SCRUTINY, *No belief incompatible with* b_N *satisfies* SCRUTINY}.

I will not try to evaluate these proposals here. Instead I will work with the following slogan to characterize Constructivism: normative facts are fully grounded in *idealized* normative beliefs. These will be normative beliefs that satisfy the counterfactual condition specified by SCRUTINY, and moreover satisfy whatever additional conditions are needed for a logically coherent view.[12] The

[11] Schafer (2014: 90) suggests this approach.

[12] Note that on this formulation, a normative fact N is grounded in the fact that the corresponding

Full Grounding View can be written in schematic form as follows:

$N \Leftarrow \{b_N, b_N \text{ would survive ideal scrutiny}\}$

In closing it is worth reiterating two points: first, the Partial Grounding View is still schematic in some respects, and can be filled in with various conceptions of what ideal scrutiny is. Second, the schematic Partial Grounding View does not capture all of the types of view that have been labeled "Constructivism" about normativity in the literature. While there are ways to fill out a Constructivist-style view without taking the grounding base to consist in normative beliefs with the same content as the facts they ground, these alternative views will have additional hurdles to overcome when it comes to capturing advantages in normative epistemology. When evaluating the potential for an anti-realist view to be motivated by its capacity to explain knowledge of normative facts, the most promising place to look is at a view that takes normative beliefs that ground their own truth. The Partial Grounding View is the closest to a view that does this, complicated only to avoid the logical incoherence of a Full Grounding View that entails the existence of logically inconsistent normative facts.

1.2 Grounding in anti-realism

The view that normative facts are fully grounded in idealized normative beliefs would prima facie support some epistemological conclusions. Here I will sketch

belief is held—this is the fact b_N—plus the fact that b_N satisfies the idealizing conditions including SCRUTINY. There is alternative view, which holds that the normative fact N is grounded in a counterfactual fact—the fact that *if* an agent were to be in a state where all of her beliefs satisfy the idealizing conditions including SCRUTINY, she would have the belief b_N. This is an available view, but I will not focus on it here, since (for reasons that will become clear below) it makes the epistemological project harder for the Constructivist to satisfy, since it locates the grounding beliefs in counterfactual worlds.

these prima facie motivations for the Constructivist view, which follow from the logical properties of the grounding relation, plus the place of normative belief in the grounding base on the Constructivist view.

One relevant logical property of the grounding-relation is the following: if P grounds Q, then it is not possible to have P true and not Q. A proposition cannot obtain without its grounds. In addition, it is plausible that a grounded fact cannot occur without its grounds: if P grounds Q, then it is not possible to have Q without P. I will call these the Necessitation and Counter-Necessitation properties of grounding:

Necessitation If P grounds Q, then necessarily, if P obtains, Q obtains as well.[13]

Counter-Necessitation If P grounds Q, then necessarily, if Q obtains, P obtains as well.[14]

These claims are meant to be fully general: any grounding thesis will include a commitment to more than just Necessitation and Counter-Necessitation. But these are of special relevance to the anti-realist's epistemological aims: they constrain what modal space is like, and rule out some combinations of facts: generally, if P grounds Q, then (by Necessitation) Q can't be false if P is true. Given Constructivism, this means that if a normative belief b_N is held in the right conditions, then N holds as well. Moreover if generally (by Counter-Necessitation) P can't be false if Q is true, then if the normative fact N holds, its

[13] Rosen (2010: 118)

[14] Counter-Necessitation is more controversial than Necessitation, since it embodies some substantive assumptions about how grounding relates to multiple realizability. (See Schaffer (2015) for more on this issue.) On any approach to grounding (or a cognate notion) on which a fact is grounded in all of its possible realizers, Counter-Necessitation will hold.

grounding base which includes b_N holds as well. Thus given the Constructivist view about what appropriate instances of 'P' and 'Q' are, these general structural features are of potential significance for moral epistemology.

The significance of the Constructivist grounding claim specifically lies in the appearance of a kind of modal reliability for normative belief that it entails. It not only guarantees that some normative beliefs will be *true*, for instance when it is true that I ought to ϕ is true, and is moreover true because I hold a belief that I ought to ϕ that survives ideal scrutiny. Thus the view in addition guarantees that some normative beliefs are true for a very specific reason: they are held in conditions that guarantee that the belief will be true. Thus it is natural to say that some true normative beliefs won't owe their truth to an accident of luck: there is a very straightforward explanation of why that belief is true, which is found in the metaphysical claim that is distinctive of Construcitivism.

2 Motivating anti-realism

This is just a sketch of why Constructivism would appear to have a distinctive advantage in epistemology. Much more needs to be said about the epistemological side of the equation before it can be turned into a potential argument for the epistemological benefits of anti-realism. We can now ask in more detail: what would it take for this modal reliability in normative belief to constitute an epistemological motivation for accepting anti-realism? There are no epistemological terms in the thesis itself: it does not say that normative beliefs are justified, undefeated, or count as knowledge. So the connection between the metaphysics

of Constructivism and its epistemology would appear to require further spelling out.

2.1 A first attempt: trivial connections

But this might not be obvious—in fact some of the literature can be read as denying that the modal reliability thesis needs to be cashed out in familiar epistemic terms in order to be turned into an epistemological advantage for the Constructivist. For instance, here is one passage from Sharon Street commenting on the epistemological problems for realism, which she aims to solve with Constructivism:

> Either the realist is forced to embrace a skeptical conclusion—acknowledging that our normative judgments are in all likelihood hopelessly off track, having been fundamentally shaped in their content by forces that bear no relation to the independent normative truth—or else the realist must hold that an astonishing coincidence took place—claiming that as a matter of sheer luck, evolutionary pressures affected our evaluative attitudes in such a way that they just happened to land on or near the true normative views among all the conceptually possible ones. Both of these claims are implausible, however. (Street, 2008: 208-9)

This passage contains a number of terms that indicate the alleged failure of realism to account for a certain kind of modal reliability. I won't engage with Street's

arguments for this conclusion here;[15] rather the important point for present purposes is the terminology she uses to make the allegation. In this passage she uses the terms 'off-track', 'coincidence', and 'sheer luck' to describe the options for the realist's normative beliefs. But nowhere does she define these terms in familiar epistemological notions like knowledge, justification, defeat, and the like.[16] Rather on one reading, 'off-trackness', 'coincidence' and the like are epistemological vices in themselves. That normative belief for the realist is 'off-track' in Street's sense is *itself* an epistemological defect, and is not a defect because being 'off-track' is connected to the absence of knowledge, or to epistemic defeat.

So it is possible for the Constructivist to use analogous notions to motivate her own view: if the realist cannot explain why normative belief is not (for example) off-track, then the Constructivist might claim an advantage on the basis of an explanation for why normative beliefs do possess the relevant property of not being off-track. But on this strategy the argument is *not* that Constructivism is preferable to realism on epistemological grounds because it explains the presence of knowledge, or the absence defeaters, for normative beliefs. Rather the connection between the modal connection between normative belief and normative fact according to Constructivism is *trivial*. No further account of why the failure is of epistemological relevance is necessary on this approach. The question for this approach is how much work it can do in showing that epistemology favors certain brands of anti-realism over realism.

[15]For an extended discussion, see Dunaway (2016).

[16]For more discussion of the relationship between Street's arguments and familiar epistemological terms, see Clarke-Doane (2012, forthcoming b).

It is undeniable that Constructivism in the form of the Partial Grounding View is able to explain certain modal relationships between normative beliefs and normative facts that the realist cannot explain (or, at least the realist cannot explain the connections in the same way[17]). In a world where one holds a normative belief that one ought to ϕ, where the belief is also a belief that survives ideal scrutiny, it follows from Constructivism (plus the NECESSITATION property of grounding) that it is a fact in that world that one ought to ϕ. Analogous connections between normative belief and fact will hold across modal space.

An epistemological motivation for Constructivism is, however, aimed at *convincing* us that Constructivism is true.

A proper motivation for the view would not merely show that Constructivism has a feature Constructivists think constitutes an epistemological virtue. Rather it should show that Constructivism has a feature which is recognizable as an epistemological virtue whether Constructivism is true or not. Only then can the Constructivist claim to have motivated her view in the any helpful sense—that is, to have shown that Constructivism is independently appealing in a way that realism is not, and thereby have shown that previously uncommitted theorists have a reason to adopt constructivism as their view in virtue of its epistemological properties. There are lots of conceivable modal connections between normative belief and normative fact.

Thus the Constructivist can decisively show that her view establishes some

[17]Insofar as Constructivism has some revisionary consequences for first-order normative claims (see Street (2008)), it pretty much follows that realism won't explain precisely the same modal connections.

modal connection between the two, but does not successfully motivate her view by treating these modal connections as trivially epistemologically valuable. Whether these modal connections matter, epistemologically, is something someone tempted toward realism need not be convinced of. A successful epistemological motivation for anti-realism will need to rely on a substantive and non-trivial connection between the modal properties of normative belief and epistemological virtues that a realist will recognize.[18]

2.2 A second (and better) attempt: knowledge and defeat

The properties that will count as uncontroversially valuable in an epistemological sense include (among others) *knowledge* and the absence of *defeat*. We have seen that the Partial Grounding View entails a certain modal relationship between normative beliefs, and the normative facts that make these beliefs true, will hold. The most plausible route for the Constructivist to pursue, then, is to connect this modal relationship with the modal properties of knowledge and defeat.

I will say more about the details of the modal features of epistemically relevant notions shortly. But first it is worth emphasizing the roles these notions do, and the roles they do not, play in the overall dialectic of this paper. I envisage

[18] A similar point applies to the arguments in Setiya (2012: 96), who also targets a non-standard epistemic constraint (although, unlike Street, he explicitly labels it a part of the anti-luck condition on knowledge). According to this condition, knowledge requires not only using a reliable method, but its not being an accident that the method one used was reliable. Aside from doubts about whether this is indeed a necessary condition, the condition is dialectially ineffective for anti-realists. It is much more stringent than ordinary anti-luck conditions on knowledge, since some beliefs that are true in all nearby worlds (and thereby have no bad companions) can be formed by methods that are only accidentally reliable. But this condition will not be very successful at convincing realists that their view faces epistemological difficulties. At best they will take Setiya's discussion to show that they should use something weaker to characterize the anti-luck condition on knowledge. See Schafer (forthcoming: §3) for more on this point.

the modal features of knowledge, defeat, and the like, as constituting the most promising area where Constructivism can earn epistemological credentials.

This is a methodological assumption, and not a principled epistemological claim, to the effect that the modal dimensions to knowledge and justification are the *only* relevant epistemological virtues. There might well be others; I am only assuming here that if Constructivism fails to explain why the modal dimensions of knowledge and defeat are satisfied, then it would appear even less promising as a candidate to explain why other dimensions to knowledge and justification are met. Of course this possibility might be bore out; I will do nothing here to argue against it. Thus the conclusion I will argue for here is somewhat limited in scope: I will argue that the Constructivist faces serious difficulties in taking the most promising route to deriving epistemological benefits from her view. There may well be other routes; but we should both worry that the apparent epistemological advantages to anti-realism are illusory, and any attempt to recover the advantages will require serious work.

We needn't be too stringent about the criteria for success here. In order to show that her view has desirable features with respect to epistemological problems, the Constructivist does not necessarily need to show that *every* normative belief is a piece of knowledge, or avoids defeat. Even an anti-realist who can claim that her view has advantageous consequences with respect to normative knowledge can still maintain that some normative beliefs are false. For example on the Partial Grounding View, it is possible that our actual normative beliefs fail to match what our beliefs formed under conditions of idealized scrutiny would be. Some of these

beliefs will be false, and hence not knowledge.[19] The standards for success are not so high: she only needs to show that some pervasive obstacles to knowledge, or general sources of defeat do not arise on her view.

2.3 Risk, knowledge, and defeat

I have already emphasized that Constructivism, in the form of a Partial Grounding View, has the resources to explain some striking modal connections between normative beliefs and normative facts. The most natural place to look for a connection with familiar epistemological properties is to the modal dimensions of knowledge and defeat. One such feature, which I will focus on here, is the *risk* of false belief. A belief that is at risk of being false has a modal property: roughly, there is a nearby world where that belief is false. Once we make some needed refinements about exactly what this amounts to, it will be very plausible that the presence of this kind of risk—or something closely related to it—will be incompatible with knowledge. And the Partial Grounding View will have an identifiable task: to show that because normative beliefs ground their own truth, normative beliefs are (for the most part) not susceptible to the kind of risk that is inconsistent with knowledge.

Begin with some suggestive examples: someone staring at 2 pm at a broken clock with its hour hand pointing directly at '2' doesn't know what time it is, even if she has (on the basis of her staring at the clock) a true belief about the time.[20]

[19] In addition some normative beliefs are subject to defeaters, even for the Constructivist: for instance when you receive misleading evidence that fish cannot experience any pain, you might thereby acquire a defeater for your belief that it is morally wrong to eat fish (since the basis for this belief has been defeated).

[20] Russell (1912), also see the cases in Gettier (1963).

And a natural way to think about what she doesn't know is that her belief, though true, is only luckily true and is at risk of being false: there are nearby possibilities where she forms the belief in the same way, but winds up believing falsely. (For instance she might easily have looked at the same clock at 1 pm.) Thus knowing requires satisfying at minimum an anti-luck condition, which is a modal property of a belief.[21] Lucky beliefs have nearby counterparts which are false—we can call these nearby beliefs that are incompatible with knowledge *bad companions*. Bad companions must, at a minimum, be held in nearby worlds, and be false. A few additional details on bad companionship are in order here.[22]

First a bad companion for the belief in P need not be a belief in the *same* proposition, P. All that is required is that it be a belief in a sufficiently similar proposition. For instance: if I am merely guessing in response to queries about large sums, I might correctly guess that 164 + 682 = 846. But this doesn't mean the corresponding belief isn't lucky to be true: by virtue of simply guessing, I will have similar beliefs in nearby worlds that are false, since for instance in some nearby world I falsely believe that 164 + 682 = 823.

A second feature of bad companionship: *how* one comes to the false belief in a nearby world matters for bad companionship. I might know that Sally is in town because I happened to run into her at the store. But I could easily have not seen her at the store, and if I had not, I would have believed that she was traveling in Spain (because a normally reliable friend told me that she was there this morning). This nearby false belief doesn't mean that I don't know Sally is in

[21] Unger (1968), Williamson (2000), Pritchard (2004).
[22] For more on this notion, see Dunaway and Hawthorne (2017).

town—I saw her, after all. So it isn't a bad companion for my actual true belief. Bad companions need to be formed via a process that is sufficiently similar to the process by which their actual world counterparts are formed, in order to be truly epistemically malicious.

If a belief has bad companions, it follows in present jargon that the belief is not knowledge. This is because beliefs with bad companions are, even if actually true, at risk of being false being false in the relevant way.

Bad companionship is connected to epistemic defeat as well. One way to acquire a defeater for a belief is to learn that it is at risk of being false, and hence lucky in a way that is incompatible with its being knowledge. So if one learns that one's belief has a bad companion, one will thereby acquire a defeater for that belief.

The notion of bad companionship gives the anti-realist a target for establishing her epistemological credentials in a compelling way. She can show it follows from the anti-realist view that normative beliefs are not systematically accompanied by bad companions, and thereby subject to the kind of risk that prevents them from being knowledge, and gives rise to defeaters. Of course as I emphasized earlier, there might be additional dimensions to epistemological virtues besides the absence of bad companions. But this seems like the most promising route for a non-trivial connection between the Constructivist's metaphysics and epistemology, since it would appear that the metaphysics of the Constructivist view is perfectly suited to explaining exactly this kind of modal reliability. Whether this appearance reflects epistemic reality is the question I take up for the rest of this

paper.

3 The Full Grounding View and bad companions

Recall the Full Grounding View, which we dismissed in §1 on logical grounds: this is the view that the full grounds of the normative fact that one ought to ϕ are the fact that one holds the belief that one ought to ϕ. This view, while it has the disadvantage of being incoherent, would plausibly succeed in ruling out the existence of bad companions for normative beliefs. So it is helpful to start with the Full Grounding View, in order to get a feel for how the Constructivist might try to exploit the resources of her view to show that bad companions will, in general, not be present for normative belief.

First, on the Full Grounding View, any normative belief will be true in the world in which it is held. According to this view, when I believe that I ought to ϕ, the complete grounding base for the fact that I ought to ϕ is thereby instantiated. So, by Necessitation, it is true that I ought to ϕ, and my normative belief is true. Moreover there won't be any nearby worlds where my belief is false, either. This is for the same reason that guarantees that my actual normative belief is true. When I believe in a nearby world that I ought to ϕ, the complete grounding base for the fact that I ought to ϕ is thereby instantiated. So, by Necessitation, it is true that I ought to ϕ, and my normative belief is true in the nearby world as well. In all nearby worlds, then, my normative beliefs are true. My actual normative beliefs have no bad companions, since none of the candidates for companionship are false.

Of course the Full Grounding View is a non-starter. We need to move to a Partial Grounding View, which not only grounds normative fact in normative beliefs, but in additional facts as well, namely facts about which normative beliefs survive ideal scrutiny. Idealized normative beliefs are not only those that survive the process of acquiring full information, reflection, and the like, but also additional (and so far unspecified) idealizing constraints that ensure full coherence. These additions are necessary for a viable Constructivist view. But we need to ask whether, once we move to a Partial Grounding View, we are able to keep the straightforward epistemological benefits of the Full Grounding View.

At a purely formal level, the reason why a Partial Grounding View cannot directly claim the same epistemological benefits is straightforward: a world where a normative belief that I ought to ϕ is held is not ipso fact a world where it is true that I ought to ϕ, since the belief by itself does not suffice for the normative fact. On a Partial Grounding View, there are additional grounds besides the belief itself which ground a normative fact. Since one might hold the belief that one ought to ϕ, but the belief not survive ideal scrutiny, the belief can be false. There is no guarantee of an absence of bad companions for normative belief in the same way. By moving away from the Full Grounding View, we need to find an alternative way to derive the absence of bad companions for normative beliefs.

4 A false start: the Grounds Transfer principle

One natural thought is that the Full Grounding View and the Partial Grounding View share one important feature, which matters to the epistemology of Con-

structivism. The feature in question is the epistemic status of the grounding base. In particular, the grounding base on both the Full and Partial Grounding Views consists in *psychological* facts—facts about what an agent believes, or would believe under particular counterfactual circumstances that are specifiable in psychological terms. Since the epistemology of psychology is not under threat here—we have no reason to suppose that these psychological facts, whether actual or counterfactual, are unknowable—both versions of Constructivism ground normative facts in facts which are uncontroversially knowable, by ordinary empirical methods. And so we might conclude that the Partial Grounding View entails that normative facts themselves are knowable, because they are grounded in easily knowable psychological facts.[23]

More generally, at the heart of this strategy is the Grounds Transfer principle:

Grounds Transfer If P grounds Q, and one knows P, then one knows (or can easily come to know) Q.

Of course this general principle, even if true, might be of very little help for

[23]This line of thought is suggested in Schafer (2014: 90), where he endorses a principle which entails that I can know that I ought to ϕ whenever I know I have an idealized belief that I ought to ϕ. Knowing the grounds of my obligations is sufficient for knowing the normative facts about what my obligations are.

The general version of Schafer's principle is: "The judgment/assertion that P is warranted just in case the judger/asserter is a position to know that P is true$_S$ relative to his normative perspective." (p. 90) Here 'true$_S$' marks the relational version of truth which Schafer employs, that holds between a proposition and an agent's perspective. Schafer here uses the notion of 'warrant', since for technical reasons this general norm cannot connect relative truth with knowledge: knowledge of a proposition entails that the proposition is true simpliciter, and so a principle which claims that an agent can have knowledge of a proposition whenever that proposition is true$_S$ from that agent's perspective would make any knowable proposition true simpliciter. So Schafer's norm in general form concerns the non-factive epistemological property of warrant instead. However in the first-personal case, this technical worry does not arise: if it is true relative to my perspective that I ought to ϕ, there is no barrier to my claiming that I *know* that I ought to ϕ, and hence that it is true simpliciter that I ought to ϕ. In what follows, I will use first-personal normative claims when discussing views similar to Schafer's to avoid these complications.

normative epistemology. If we never, or very rarely, know what our normative beliefs are and whether they survive ideal scrutiny, then the Grounds Transfer principle will have nothing to say about normative knowledge. Since its antecedent would go unsatisfied, Constructivism as view about what grounds normative facts would be in no position to take advantage of it. But it is very natural to add an additional claim that these grounding facts are in fact easy to know: we are, in general, in a position to know which of our beliefs are also idealized normative beliefs.[24] It is not obvious that this assumption is true, but I will grant it to investigate the Grounds Transfer principle more carefully.[25]

We are granting, then, that facts about which of one's beliefs are idealized are easy to know. Whether this claim is helpful to the Constructivist's epistemological ambitions is a question of whether knowledge that a normative belief survives ideal scrutiny eliminates any potential bad companions for the normative belief. Beliefs about what one ought to do—normative beliefs—are different kinds of beliefs from beliefs about what one's own normative beliefs are, and beliefs about

[24]Schafer (2014: 90).

[25]Berker (2014: 243) points out that the Constructivist's resources for explaining normative knowledge appear to only work when we focus on *first-personal* normative beliefs. (See also Chrisman (2010).) Since the belief that *you* ought to ϕ is grounding in what *you* believe about ϕ-ing, we can no longer be assured, merely by the facts about what grounds normative fact, that my belief that you ought to ϕ is true.

One advantage to Grounds Transfer is that it presents a response to this objection. In principle there is no insurmountable difficulty in knowing someone else's psychological states, though how I come to know them might be different from the route by which I come to know my own psychological states. If knowledge transfers over grounds, as Grounds Transfer claims, then third-personal psychological knowledge can give rise to knowledge of third-personal obligations.

Since I will end up rejecting Grounds Transfer, this response to Berker will not in the end be satisfactory. The Constructivist may have to concede that her epistemological project is more limited than originally promised: instead of showing that Constructivism explains how all the normative facts can be known, it is better advertised as showing how each agent can come to know the obligations that apply to her. This would still be an interesting advantage over the realist, if it could be achieved, and I will pursue the question of whether Constructivism can even achieve this more limited goal in the next section.

whether they are idealized. So the absence of bad companions for one kind of belief—*viz.*, beliefs about one's own normative beliefs—might fail to guarantee the absence of bad companions for normative belief simpliciter.

In fact this possibility is very likely to be realized. There are two especially common ways for this failure to occur, each of which has the upshot that Grounds Transfer is false. Nothing in the metaphysics of Constructivism rules out the existence of bad companions.

1. First case: no correlation. We are granting that one can know the grounding base for normative facts, namely facts about which normative beliefs survive ideal scrutiny. This means that (at the very least) one's beliefs about which normative beliefs survive ideal scrutiny are true across all nearby worlds. One is not at risk at having false beliefs about these matters. This means that, in every nearby world, if one has a belief about which beliefs survive ideal scrutiny, then one has a true belief in a proposition which fixes a normative fact.

More formally: suppose one knows that the normative belief b_N survives ideal scrutiny. In every nearby world, then, if one believes that b_N survives ideal scrutiny, then it is true that b_N survives ideal scrutiny. And in each of these worlds, the fact that b_N survives ideal scrutiny entails (by the Partial Grounding View and Necessitation) the normative fact N. So in each nearby world where one believes that b_N survives ideal scrutiny, N is true in that world as well.

But this doesn't mean that one's normative beliefs in these nearby worlds will reflect the normative facts. One might not be aware, across all nearby worlds, of the connection between the (known) facts about which beliefs survive ideal

scrutiny and the normative facts about what one ought to do. Nothing in this view explains why there won't be a nearby world where one truly beliefs that b_N survives ideal scrutiny, but declines to treat this as relevant to the normative facts (as it were, one treats one's beliefs in a state of ideal scrutiny as irrelevant to what the normative truths are). Thus one comes to believe the negation of N. So one's normative belief in this nearby world is false. It is a bad companion for one's actual normative beliefs, and the existence of the bad companion is fully compatible with one's beliefs about which beliefs survive ideal scrutiny being known.

More concretely: suppose it is true that one ought to ϕ. This means that the ground for this normative fact holds: one's belief that one ought to ϕ survives ideal scrutiny. If one *knows* that one has the belief, and that it survives ideal scrutiny, then in all nearby worlds, one has no false beliefs about whether this belief survives ideal scrutiny. But just because one has these true beliefs concern the *grounds* of normative facts in nearby worlds, one could very easily form, in such a world, that it is not the case that one ought to ϕ. This happens when one doesn't know about the connection between beliefs that survive ideal scrutiny and normative facts. And in such cases we get bad companions for true normative beliefs.

The failure of the Grounds Transfer principle in this case is analogous to the way in which a simple closure principle, which claims that knowledge transfers across mere implication, fails.[26] One can know the simple mathematical fact

[26] *cf.* Hawthorne (2004)

that 2 + 2 = 4, which (suppose) entails the Continuum Hypothesis. But this doesn't guarantee that one knows the Continuum Hypothesis, since one has no knowledge of the logical relationship between the two claims. Similarly the Grounds Transfer principle ignores the possibility that one can fail to know normative facts, even though one knows the facts which ground them, because one can fail to know that the grounding relationship obtains.

2. *Second case: withholding belief.* Sometimes one's belief that P has no bad companions because, in some of the nearby worlds, one has no beliefs at all on P-related matters. (Since bad companions are by definition false beliefs, a world where one has no relevant beliefs at all must be worlds that contain no bad companions.[27]) Take our earlier assumption that we can know the psychological facts about what our normative beliefs are, and whether they survive ideal scrutiny. This means that our beliefs about these psychological facts are not false in nearby worlds. But it doesn't mean that we will have beliefs about these psychological facts in all these nearby worlds. It is perfectly compatible with my knowing which normative beliefs survive ideal scrutiny that I do not have beliefs about about ideal scrutiny in some nearby worlds at all.

So knowing the grounds of normative facts doesn't imply that in *all* nearby worlds one has true beliefs about the grounds of normative facts. There are worlds where the psychological ground of a normative fact obtains, but one has not beliefs about it. But in these worlds, it will be extremely likely that one will

[27] We could make the conditions more stringent by including absence of belief as sometimes entailing bad companionship. See Manley (2007). But here this will make the Constructivist's epistemological project even harder, so I will ignore this complication.

continue to have *normative* beliefs about what one ought to do. And since one's beliefs about the grounds of normative facts will be absent, one won't be able to utilize true belief about the grounds of normative facts to form true beliefs about the normative facts themselves. Thus in these worlds there is no reason to expect that one will form true normative beliefs—after all, one won't have true beliefs about the facts which determinate the normative facts. So one's normative beliefs in these worlds can be false, and constitute bad companions for one's actual normative beliefs. What would ensure that one's normative belief is true, when one doesn't have any opinion about the psychological facts that ground it?

More concretely: suppose it is true that one ought to ϕ. This means (by Counter-Necessitation) that its ground holds, so it is also true that one's belief that one ought to ϕ survives ideal scrutiny. Moreover we are spotting the Constructivist the assumption that one knows this, so it also follows that one has no false beliefs in nearby worlds about whether one's belief that one ought to ϕ survives ideal scrutiny. In some of these nearby worlds, however, one will have no beliefs at all about whether the belief that one ought to ϕ survives ideal scrutiny. This does not prevent one from forming, in such a world, a belief about whether one ought to ϕ or not. And since by hypothesis one has no beliefs about the relevant idealized normative beliefs in such a world, one will form beliefs about whether one ought to ϕ with no beliefs about the grounds this normative fact to guide them. In some worlds like this one's normative belief will be false, and constitute a bad companion for one's actual normative beliefs.[28]

[28]More formally, we can suppose that in the actual world the belief b_N survives ideal scrutiny, and one knows that b_N survives ideal scrutiny. The question is whether the normative belief b_N is itself

The Grounds Transfer principle is false. The modal relationship between normative belief and fact secures the impossibility of normative facts that come apart from ideal normative beliefs. But for *beliefs* about the normative facts to be accurate in nearby worlds, one needs more than grounding: one needs awareness of the grounds in addition; otherwise one's normative beliefs will have no guidance that prevents them from falling into error.[29]

It is worth at this point briefly rehearsing a point made by Selim Berker which applies to an obvious fix to this issue. As we noted earlier, the analogy with the transfer of knowledge over logical entailment suggests that the problems with bad companionship for normative beliefs go away for agents who *know* that the Constructivist's grounding thesis is true. That is, take agents who know that the fact that one ought to ϕ is grounded in the fact that the belief that one ought to ϕ would survive ideal scrutiny. We have seen that knowledge doesn't transfer over the grounding relation alone, but it is much more plausible that for such agents, their knowledge does transfer over a *known* grounding relation. So for agents who know that the Partial Grounding View is true, and also know that a particular normative belief survives ideal scrutiny, it is very plausible that these agents also have some normative knowledge. So the Constructivist's epistemological advantage appears to be back in play, once we focus on agents who know the appropriate non-normative facts.

knowledge—this doesn't follow immediately from the fact that one knows that b_N survives ideal scrutiny. If one knows that b_N survives ideal scrutiny, then one has no false beliefs about the subject in nearby worlds. But some of these worlds will be worlds where one has no beliefs at all about whether b_N survives ideal scrutiny. In some of these worlds, one will still have normative beliefs about whether N holds, and some of these beliefs will be false. So the normative belief b_N in the actual world has bad companions, and is not knowledge.

[29] Schroeder (2007) makes a similar point.

But here there is reason to be cautious for anyone who hopes to turn this into a distinctive epistemological advantage for the relevant Constructivist view: this line of reasoning will work for *anyone* who accepts a grounding thesis for the normative.[30] Nothing in the idea that knowledge transfers across a known grounding base depends on the details of how the Constructivist grounds the normative. If this strategy is helpful to the Constructivist, then, even a realist who holds some alternative view about the grounds for normative facts, can likewise claim that *if* someone were to know this grounding thesis, then they will have normative knowledge available to them. So there is no distinctive epistemological advantage for the Constructivist here.[31]

A similar point arises when we consider agents who form normative beliefs not on the basis of their knowledge of the grounding claim, but instead on the basis of a direct inference from the grounding fact. For example, one might come to believe that one ought to ϕ by inferring it directly from the claim that the belief that one ought to ϕ survives ideal scrutiny. This doesn't require knowing, or even believing, that the second claim grounds the first, and so can be knowledge-producing even if Grounds Transfer is false. Instead, it is knowledge-producing if the following principle is true:

Grounds Transfer* If P grounds Q, and if one knows P and infers Q from P,

[30] Berker (2014)

[31] Perhaps the proponent of the Partial Grounding View can argue that, while realists can avail themselves of the same line of thought *if* one were to know the realist's view about what grounds normative facts, the realist's grounding thesis will be harder or even impossible to know. This is an alternative route that is worth exploring. But it deserves emphasis that it faces significant hurdles, since the Constructivist's grounding thesis appears quite counterintuitive with respect to ordinary moral thought, and consequently does not look like the kind of thing that will be easier to know than alternative appearance-saving proposals about the grounds of normative facts. Schafer (forthcoming: 80) discusses some of the hurdles for this line of argument.

then one knows Q.[32]

If knowledge requires the absence of bad companions, Grounds Transfer* is especially appealing as a way to avoid the difficulties that beset the original Grounds Transfer principle. A nearby false belief constitutes a bad companion for one's actual belief only if it is formed by a relatively similar process as the actual belief (*cf.* §2.3). We might then add that believing a normative claim N on the basis of the fact that N survives ideal scrutiny is a very different process from believing N on other grounds. Thus someone who knows N survives ideal scrutiny and believes N on this basis will have a normative belief with no bad companions. Given Necessitation, knowing the grounds and inferring N on this basis will be sufficient to ensure that belief in N has no bad companions as well—any false beliefs on the subject will be formed by another process.

So Grounds Transfer* avoids some of the pitfalls of the original Grounds Transfer by explaining how the Constructivist grounding claim can be part of an explanation of normative knowledge without requiring knowledge of how normative facts are grounded. But there are still reasons to worry that it will not be helpful to the Constructivist's claim to have an epistemological advantage over realism.

Grounds Transfer* relies on a very strong criterion for which beliefs are candidates for bad companionship: they must be the formed by the *same* process of inferring normative claims from beliefs about beliefs that survive ideal scrutiny. But in general the standards for bad companionship are not so strong: in our

[32]Thanks to an anonymous referee for suggesting this alternative.

example of guessing mathematical truths in §2, the bad companions for one's actual belief are not formed by the exact same process, but rather one which results in one forming a different (and false) mathematical belief. The lesson from this is that we shouldn't individuate belief-forming processes too finely in an attempt to characterize the relationship between luck and the absence of knowledge. But this is what makes Grounds Transfer* seem helpful to the Constructivist: the sketch in the previous paragraph is appealing only if we assume that beliefs formed by an inference from facts about which beliefs survive ideal scrutiny can serve as bad companions. This is a very narrow construal of a belief-forming process, and one we should be wary of using.

When we are asking whether, if Constructivism is true, then agents who form normative beliefs by the inferential process suggested by Grounds Transfer* have normative knowledge, this worry is especially pressing. There are many similar belief-forming processes that do not necessarily yield normative knowledge, because in some circumstances these processes do not even yield true beliefs. For instance: one might infer N from the fact that your better-informed self believes N. But even if the Partial Grounding View is correct, it is possible that your better informed self believes N, and yet N is false. In some circumstances, one will form this false belief in N in a nearby world. There will be significant pressure to treat it as a bad companion.

There are other examples of similar belief-forming processes that, in some environments, yield false beliefs. Not all of these involve inferring normative claims from beliefs about something other than ideal scrutiny. Some falliable

processes involve using different inference patterns—for instance, inferring from the fact that the belief in N survives ideal scrutiny that N is objectively true for all agents at all times. This is a stronger claim that is false on the Constructivist view.

Thus the conditions under which the Constructivist can claim that such knowledge is present are remarkably strong, requiring agents to adhere to a very specific belief-forming process across all nearby worlds. Anti-realists should not view this as much of an advantage over their competitors.[33]

Admittedly, there is something appealing in the thought that if, as the Constructivist claims, normative facts are grounded in our attitudes, then we should have better epistemic access to these facts. In particular it is tempting to think that we are in a position know normative facts by inferring them from our knowledge of their grounds, only when these grounds are not independent of us and our attidudes, as the Constructivist claims. But this idea is difficult to make sense of in a way that both relies on true epistemic principles about the relationship between grounding and knowledge, and yet also is unavailable to realists.

5 A more promising route: Automatic Truth

Here is another strategy. Recall what made the Full Grounding View rule out some bad companions for normative belief: it holds that a normative belief is the *only* constituent of the grounding base for the relevant normative truth. Whenever the belief b_N is held, it grounds the normative fact N—and hence it grounds the

[33]In addition, a version of Berker's point still applies: realists can make equally good use of Grounds Transfer*. Whatever the realist holds that there are some grounds for normative facts. For exactly the same reason as before, Grounds Transfer* will entail that agents who know these grounds, and infer normative facts from them, will have normative knowledge.

fact that makes it true. This rules out bad companions.

The Partial Grounding View says that the normative belief b_N is *part* of the grounding base for the normative fact N. There are additional components—what I have called the facts that go into making b_N a belief that survives ideal scrutiny. But it is very natural to think that, for any normative belief that is accompanied by the fact that it survives ideal scrutiny, the epistemology of that belief will be the same as the epistemology for normative beliefs on the Full Grounding View. That is, if a normative belief b_N survives ideal scrutiny, then it will ground the normative fact N and hence will ground its own truth. This appears to be the same feature of the Full Grounding View that made it appealing for epistemological reasons. It is worth investigating whether this appearance holds up under a closer look.

At the heart of this motivation is a connection between the fact that, on the Constructivist view, some normative beliefs ground their own truth, and the epistemologically relevant property which we have labelled freedom from bad companions. The Automatically True principle states this connection:

> Automatically True If one has a normative belief in P that survives ideal scrutiny, and P is automatically true because it is grounded in the ideal scrutiny-surviving belief in P, then the belief in P is free of bad companions.

We have already seen that an analogue of Automatically True is very plausible on the Full Grounding View. On the Partial Grounding View, some normative beliefs—namely those that survive ideal scrutiny—also ground their own truth. If Automatically True is correct, then this should also mean that

they are free of bad companions in virtue of the fact that they ground their own truth.

The status of AUTOMATICALLY TRUE represents a very specific framing of a question that is generally of interest to the epistemology of Constructivism. It could be argued that Constructivists have generally assumed that the distinctive grounding claim for normative facts entails that normative beliefs have a privileged epistemic status. For instance Sharon Street makes the following claim about the epistemic status of normative belief on the Constructivist view, without much in the way of argument that the relevant epistemological claims follows from the Constructivist metaphysics:

> The breaking of our bones *is* bad, in other words, and we're well aware of this. But the explanation is not that it is true independently of our attitudes that the breaking of our bones is bad and we were selected to be able to notice this; the explanation is rather that we were selected to *take* the breaking of our bones to be bad, and this evaluative judgement withstands scrutiny from the standpoint of our other evaluative judgements ... (Street, 2006: 154)

Here Street claims that we are "aware" of a normative fact (the fact that breaking our bones is bad), and appears to say that this follows directly from the Constructivist's distinctive claim about what makes this normative fact obtain (that we take the breaking of our bones to be bad, and this evaluative judgment withstands scrutiny).

As emphasized earlier, I am not addressing this claim directly. Rather I am

addressing a related version of it that has several advantages. Instead of the question of whether we are "aware" of normative facts, I am directly discussing whether we know, or are justified in believing, these normative facts. And I am focusing only on one condition on knowledge and justification, namely freedom from an epistemically relevant kind of of luck (the presence of "bad companions" in present lingo) that the Constructivist seems particularly well-equipped to explain.

Perhaps there are other, even more promising avenues for the Constructivist to pursue this line of thought, without reference to knowledge and bad companions. My goal here is not to explore all of them. Instead in asking whether Automatically True is correct I will be exploring whether the modal features of the grounding relation (Necessitation and Counter-necessitation) can explain why modal features of knowledge and justification are satisfied (and in particular whether it guarantees the absence of bad companions). If this fails, then we will have a roadmap of where the hurdles are that other routes to vindicating the epistemological advantages of Constructivism will need to overcome.

Automatically True is false. One way to see this is to focus on two senses in which a belief can be immune from error. In one sense, a belief is immune from error if a belief formed in the conditions in which it is actually formed cannot be false; it is formed in conditions which entail that the belief is true *in those conditions*. In another sense a belief is immune from error when bad companions are absent. This includes not only freedom from error in the actual circumstances the belief is formed in, but also freedom from error in nearby worlds. Some beliefs easily

have been formed in conditions that are different from the conditions in which it was actually formed, so these senses of immunity from error are not equivalent.

Begin with the first kind of immunity from error. The Partial Grounding View holds that the normative fact that one ought to ϕ holds only if the belief that one ought to ϕ survives ideal scrutiny. But *if* one holds the belief that one ought to ϕ in such a condition, then it follows that the fact holds as well: one holds the belief in a condition that guarantees the truth of the belief. So—on at least one way of precisifying how a belief can be "immune from error"—normative beliefs on the Partial Grounding View have this feature, when they are held in the right conditions.

But Automatically True says that normative beliefs, if they satisfy the relevant idealization requirements, are immune from error in another sense—they have no bad companions. Since being free of bad companions is a matter of being true in all nearby worlds, immunity from error in the actual world does not necessarily entail freedom from bad companions. This second type of immunity from error is, in the framework we are working with here, the epistemologically significant property. If all the Constructivist can claim about for normative beliefs is that that they are actually true when they survive ideal scrutiny, she will face the possibility that, ultimately, her view does nothing to address questions in epistemology.

And in fact it will be quite common for normative beliefs that survive ideal scrutiny—that is, beliefs which are guaranteed to be true in the actual world according to the Partial Grounding View—to have bad companions. Thus they

will lack immunity from error in the epistemologically important sense, and Automatically True will be false. The important point to notice is that many beliefs that actually survive ideal scrutiny are beliefs that could easily have failed to do so. In other words they are beliefs that are actually formed in states of perfect logical coherence, full information, etc., and thereby satisfy the counterfactual requirement Scrutiny. But the agent who forms such a belief might be formed a similar normative belief in nearby worlds in a state that is not fully logically coherent, or in a state where she lack certain information. One might, in one world, believe that one ought to ϕ, and do so in a way which makes the relevant belief satisfy Scrutiny, but in a nearby world have the same belief in a way that does not satisfy Scrutiny. The belief will be idealized in the first world, and ground its own truth *in that world*. But in the second nearby world, it will not be survive ideal scrutiny, and hence not ground its own truth in the nearby world. There is no guarantee that the belief is true in the second world, and so the possibility of bad companionship remains.

So Automatically True contains a subtle, but important error: it mistakes the immunity from error in the *actual* world for a complete absence of risk of error, simpliciter. Since normative beliefs will be free from the risk of error—that is, have no bad companions—only if there are no false beliefs in nearby worlds as well, the Partial Grounding View does cannot rely on Automatically True to claim any epistemological benefits.

Here it is worth revisiting one point from our discussion of Grounds Transfer from the previous section. It is tempting to save Automatically True by holding

that beliefs formed in worlds where one is not in a state of ideal scrutiny cannot bad companions, since they are formed by different processes.[34] There is a sense in which this is true: forming a belief in a state of full information and full coherence is strictly speaking a different way of forming a belief than forming a belief in a state of less-than-full information or less-than-full coherence. But this does not mean that these beliefs are not capable of serving as bad companions for beliefs that are formed in a state of ideal scrutiny. Bad companionship does not require strictly identical belief-forming processes, and beliefs formed in states that fall short of ideal scrutiny might be formed by similar enough processes to, in some circumstances, count as bad companions.

The foregoing arguments suggest that one way of filling out a minimally plausible version of anti-realist Constructivism—the Partial Grounding View—fails to deliver in any straightforward way on a simple epistemological goal that it appears especially well-suited for. A simple modal reliability condition on knowledge and justification, formulated here in terms of the absence of "bad companions", is a central epistemological property that is a appears at first glance to have a promising explanation for proponents of the Partial Grounding View. The arguments I have given here show that this appearance is misleading.

We can speculate further on what lessons to draw from this (somewhat limited) conclusion. One tempting additional conclusion is that the tendency to attribute any epistemological advantages to anti-realism about the normative is misguided. We have seen why one of the most natural ways to straightforwardly derive

[34]Thanks to an anonymous referee for suggesting that this point also might be raised in connection with Automatically True.

epistemological virtues for anti-realism fails; we might assume that anti-realists have implicitly relied either either on Grounds Transfer or Automatically True when claiming epistemological advantages for their view, and conclude on this basis that there is no way for anti-realism to make good on its epistemological ambitions.

Of course nothing of the sort straightforwardly follows from what we have argued for here. There might well be ways to derive epistemological advantages to anti-realism by focusing on a version of anti-realism that is distinct from the Partial Grounding View, or by focusing on conditions on knowledge or justification that cannot be connected to the absence of bad companions. These are possibilities that might be pursued by Constructivists who wish to rehabilitate the epistemological prospects for anti-realism. I will not explore them further here.

While these unexplored lines of defense are a promising note for the proponent of Constructivism, there are also other lines of attack on the view. One concerns the relationship between relativity and disagreement: as Setiya (2012) argues, Constructivists will have to hold either that idealized agents will converge on the same normative beliefs, or else that they believe different normative claims and do not disagree with each other. These may well be additional difficulties for the Constructivist, and we should weigh whether any alleged epistemological advantages would be worth these costs.

But I will not focus on these costs here, since, as the arguments above suggest, it is a difficult task to show that these epistemological advantages are present, at any cost. This is a distinctive challenge for anti-realists in general, and Con-

structivists in particular. There is no shortage of defenses of realists in response to the epistemological challenges that motivate anti-realism,[35] where these defenses are aimed primarily at showing that realists can explain why normative beliefs connect with the facts in a way that is epistemologically satisfying. Others are less sanguine about realists meeting these epistemological challenges, but like Setiya (2012), point out that Constructivism is less appealing on non-epistemological grounds.[36] What is less often discussed is whether, even if we grant that realism faces epistemological difficulties, these same epistemological problems arise for promising forms of anti-realism as well. With some natural assumptions in place, the results are not friendly to anti-realists.

There is another general methodological issue that we should also mention. So far I have focused on the question of whether the Constructivist can explain why certain normative beliefs are free of bad companions, and reached pessimistic conclusions about two natural routes to this conclusion. But none of this entails that, if the Partial Grounding View is correct, people *never* have normative knowledge or undefeated normative beliefs. If the Partial Grounding View is correct, it is perfectly possible that sometimes people will have normative beliefs that are free of bad companions. But this by itself doesn't show anything about the epistemic status of the Partial Grounding View itself—the standards for deriving epistemological benefits from a view require more than that the view simply fails to be inconsistent with the possibility of knowledge (or justification). Mere consistency with knowledge is cheap, and uninteresting for the purposes of

[35]See for example Berker (2014) and Enoch (2010).
[36]Enoch (2010: §5.1) summaries this point.

providing an epistemological motivation for a particular view of the metaphysics of normative facts. Thus in closing I will outline in a perspicuous way what I take the explanatory project for the anti-realist to be, if she is to claim epistemological benefits that are of genuine theoretical interest.

In closing I will discuss this issue in more detail, and use it to raise a more general problem for anti-realism two things. The specific problems we located for Grounds Transfer and Automatically True in the present framework are instances of a general problem all anti-realists will face. Those who wish to explore the epistemology of Constructivism outside of the framework of the Partial Grounding View and bad companionship will still need to face the same kind of issues, even if they arise in a different form.

6 Conclusion: methodology and idealization

Of course the proponent of the Partial Grounding View might consider additional fixes to rule out bad companions. The crux of the problem for vindicating Automatically True is that even normative beliefs formed in ideal circumstances can easily fail to have been formed in the relevant ideal conditions. One can, for instance, easily fail to have full information even if one in fact has it. By idealizing the conditions in which a belief grounds its own truth, mere risks of not being ideal become epistemically significant, as they introduce the risk of false belief.

The most obvious fix for the Constructivist is to claim that the only normative beliefs that are guaranteed to be free of bad companions are those that not only actually survive ideal scrutiny, but in addition those that are not at risk of not

being formed in such ideal conditions. These Call such a belief *robustly* ideal. If a normative belief b_N is robustly ideal, then it survives ideal scrutiny in the world w in which it is formed, and moreover in every world $w*$ that is near to w where b_N is formed is also a worlds where it survives ideal scrutiny. The partial Grounding View then guarantees that in every world $w*$, the normative fact N holds and hence b_N is true. So robustly ideal normative beliefs will be free of bad companions on the Partial Grounding View.

What does this do for the epistemological motivation for Constructivism? This is where it is worthwhile to reflect on the methodological question of what the Constructivist would need to show in order to gain an epistemological advantage over her rivals: arguably robustly ideal beliefs are of very little help on this front. We should treat Constructivism as an appealing view on epistemological grounds only if it can explain why normative knowledge is unmysterious and easy to come by. But robustly ideal beliefs are neither—very few of us are robustly ideal (for instance most people do not actually have full information, much less in all nearby worlds), and why we would be robustly ideal is not at all unmysterious (even if we did have complete information in all nearby worlds, why we should be this way is left totally unexplained).

Some agents might have beliefs that turn out to be robustly ideal, and thereby have normative knowledge. In general, any belief can, though an accident of luck, be formed in a world where there are no nearby worlds that contain bad companions.[37] These beliefs are not luckily true (they have no bad companions),

[37]See Lasonen-Aarnio (2010) for discussion of a related idea.

but they are lucky to be free from luck. When normative beliefs are robustly ideal, this has very little to do with Constructivism in the form of the Partial Grounding View or any other version—and much more to do with their good fortune to be robustly ideal. We shouldn't treat it it as a virtue of this version of Constructivism that it predicts robustly ideal beliefs as having an epistemological virtue. Rather absent further explanation as to why Constructivism should expect us to regularly form robustly ideal normative beliefs, this does nothing to separate it epistemologically from any competing view.[38]

Do these problems for Constructivist epistemology go away if we focus on epistemological virtues besides the absence of bad companions? Perhaps, but we need to be careful not to overlook lessons learned from focusing on bad companionship. We noted that a Full Grounding View—where a normative belief alone grounds its own truth—will have the desired epistemological benefits with respect to bad companionship. But this is of little help to the Constructivist project of motivating her view on epistemological grounds, since the Full Grounding View is incoherent. We need to introduce idealizations into the grounding

[38] A Constructivist could exploit a slightly weaker condition that robust ideality, but the same point applies. Some normative beliefs that are not held in ideal conditions are still true. These are beliefs about one's obligations that would continue to be held in the nearest world where one's beliefs do survive ideal scrutiny. So in principle the Constructivist could hold that some non-robustly ideal normative beliefs have no bad companions, because their content is the same as the nearest normative beliefs that do survive ideal scrutiny. Call such beliefs *guidedly ideal*. The important feature of guidedly ideal beliefs is that while they are not held in ideal conditions in all nearby worlds, they are still true in all nearby worlds.

In principle the Constructivist could claim that, while nothing in her view explains why moral knowledge is possible for beliefs that are merely ideal, beliefs that are guidedly ideal are candidates for knowledge. (Since guided ideality is not as stringent as robust ideality, this view could perhaps be said to be more helpful to the Constructivist's epistemological ambitions.) But again it is no surprise that it is possible that normative beliefs can be free of bad companions in special circumstances; in order for the Constructivist to tie this feature of guidedly ideal beliefs to a distinctive feature of her view, she would need to somehow show that Constructivism explains why we frequently hold guidedly ideal beliefs.

base, thereby making normative beliefs only part of the grounds of normative truths. Once these idealizations are introduced, the epistemic benefits of the Full Grounding View are much harder to secure.

The problem can be seen as an aspect of a tension between the securing epistemologically desirable features, and idealizations needed for plausibility in logical, metaphysical, and ethical respects.[39] Any view that ties normative belief and normative truth as closely as the Full Grounding View will have all kinds of epistemological benefits—including, but not limited to modal reliability and freedom from bad companionship. But no one actually holds such a view, for good reason. By idealizing the grounding base for normative facts, we open up the possibility of errors in normative belief. It is a large—and hitherto unanswered—question whether by introducing these idealizations we can keep any epistemological virtues of a simplistic but unrealistic view that directly ties normative fact to normative belief.

References

Selim Berker. Does Evolutionary Psychology Show that Normativity is Mind-Dependent? In Justin D'Arms and Daniel Jacobson, editors, *Moral Psychology and Human Agency: Philosophical Essays on the Science of Ethics*, pages 215–252. Oxford University Press, 2014.

Matthew Chrisman. Constructivism, Expressivism, and Ethical Knowledge. *International Journal of Philosophical Studies*, 18:331–53, 2010.

Justin Clarke-Doane. Morality and Mathematics: The Evolutionary Challenge. *Ethics*, 112(2):313–340, 2012.

Justin Clarke-Doane. Debunking and Dispensability. In Neil Sinclair and Uri Leibowitz, editors, *Explanation in Ethics and Mathematics*. Oxford University Press, forthcoming b.

[39]See Street (2008) for more on idealizations in Constructivism.

Billy Dunaway. Luck: Evolutionary and Epistemic. *Episteme*, 2016.

Billy Dunaway and John Hawthorne. Scepticism. In William J. Abraham and Fredrick D. Aquino, editors, *Oxford Handbook of the Epistemology of Theology*. Oxford University Press, 2017.

David Enoch. The Epistemological Challenge to Metanormative Realism. *Philosophical Studies*, 148(3):413–438, 2010.

Edmund L. Gettier. Is Justified True Belief Knowledge? *Analysis*, 23(6):121–123, 1963.

Gilbert Harman. Moral Explanations of Natural Facts–Can Moral Claims Be Tested against Moral Reality? *Southern Journal of Philosophy*, pages 57–68, 1986.

John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, 2004.

Maria Lasonen-Aarnio. Unreasonable Knowledge. *Philosophical Perspectives*, 24: 1–21, 2010.

J.L. Mackie. *Ethics: Inventing Right and Wrong*. Penguin, 1977.

David Manley. Safety, Content, Apriority, Self-Knowledge. *Journal of Philosophy*, 104:403–423, 2007.

Duncan Pritchard. *Epistemic Luck*. Oxford University Press, 2004.

Peter Railton. Moral Realism. *The Philosophical Review*, 95(2):163–207, 1986.

Gideon Rosen. Metaphysical Dependence: Grounding and Reduction. In Bob Hale and Aviv Hoffman, editors, *Modality: Metaphysics, Logic, and Epistemology*, pages 109–36. Oxford University Press, 2010.

Bertrand Russell. *The Problems of Philosophy*. Oxford University Press, 1912.

Nathan Salmon. *Frege's Puzzle*. Ridgeview, 1986.

Karl Schafer. Constructivism and Three Forms of Perspective-Dependence in Metaethics. *Philosophy and Phenomenological Research*, 89(1):68–101, 2014.

Karl Schafer. Evolutionary Debunking Arguments, Explanatory Structure, and Anti-Realism. In R. Debes and K. Stueber, editors, *Taking Sentimentalism Seriously*, pages 66–85. Cambridge University Press, forthcoming.

Jonathan Schaffer. Metaphysical Semantics Meets Multiple Realizability. *Analysis*, 73(5):736–751, 2015.

Mark Schroeder. *Slaves of the Passions*. Oxford University Press, 2007.

Kieran Setiya. *Knowing Right from Wrong*. Oxford University Press, 2012.

Nic Southwood. Constructivism about Reasons. In Daniel Star, editor, *Oxford Handbook of Reasons and Normativity*. Oxford University Press, Forthcoming.

Sharon Street. A Darwinian Dilemma for Realist Theories of Value. *Philosophical Studies*, 127(1):109–166, 2006.

Sharon Street. Constructivism about Reasons. In Russ Shafer-Landau, editor, *Oxford Studies in Metaethics, vol. 3*. Oxford University Press, 2008.

Sharon Street. What is Constructivism in Ethics and Metaethics? *Philosophy Compass*, 5(5):363–384, 2010.

Elizabeth Tropman. Evolutionay Debunking Arguments: Moral Realism, Constructivism, and Explaining Moral Knowledge. *Philosophical Explorations*, 17(2): 126–140, 2014.

Peter Unger. An Analysis of Factual Knowledge. *Journal of Philosophy*, 65:157–70, 1968.

Timothy Williamson. *Knowledge and Its Limits*. Oxford University Press, 2000.

Timothy Williamson. *The Philosophy of Philosophy*. Oxford University Press, 2007.