

NORMATIVITY WITHOUT CARTESIAN PRIVILEGE

Amia Srinivasan
All Souls College, Oxford

1.

A dialectical relationship between doubt and certainty characterises Cartesian epistemology: radical doubt about the outer world is quelled by the epistemic security of the inner. In viewing the phenomenal realm as a safe harbour from error, much contemporary philosophy inhabits the Cartesian worldview.¹ But recent arguments from empirical psychology and armchair epistemology alike suggest that the Cartesian presumption of privileged phenomenal access should be rejected. If these arguments are sound, we are left to conclude that there is no realm to which we enjoy privileged access. For if we do not have such privileged access to our own minds, then presumably we do not have it to anything whatsoever.

My question is this: what does rejecting Cartesianism imply for our relationship to the normative realm? My answer is that it implies that this relationship is more fraught than many would like to think. Without privileged access to our own minds, there are no norms that can invariably guide our actions, and no norms that are immune from blameless violation. This will come as bad news to those normative theorists who think that certain central normative notions—e.g. the ethical ought or epistemic justification—should be cashed out in terms of subjects' mental states precisely in order to generate norms that are action-guiding and immune from blameless violation. Meanwhile Anti-Cartesianism might come as good news to those normative theorists who resist cashing out norms in terms of mental states. For Anti-Cartesianism implies that no norms—however closely tied to the mental—can be perfectly action-guiding or totally immune from blameless violation. More generally, once we have accepted that our relationship to our own minds lacks the perfect intimacy promised by Cartesianism, we are, for better or worse, left with the view that the normative realm is suffused with ignorance and bad luck.²

2.

Following Williamson (2000, ch. 4), I'll call a condition *C* *luminous* just in case whenever one is in *C* one is in a position to know one is in *C*; *absent-luminous* just in case whenever one is not in *C* one is in a position to know one is not in *C*; and *transparent* just in case *C* is both luminous and absent-luminous. Thus a condition *C* is transparent just in case one is always in a position to know whether one is in *C*. I can now more precisely state the form of Anti-Cartesianism of interest to me here:

ANTI-CARTESIANISM: There are no transparent conditions³

Anti-Cartesianism is not merely the claim that for any condition we can sometimes fail to know whether we are in it. Even Cartesians will agree to that: the possibility of inattentiveness or negligence is sufficient to secure it. Anti-Cartesianism is the stronger claim that for any condition we can sometimes fail to be even in a *position* to know whether we are in it.⁴ For no condition will assiduous attention suffice to guarantee knowledge of whether it obtains.

Non-transparency is an unsurprising feature of most conditions. Take the condition of my car being parked on Main Street, *CAR*. It is possible for the condition to obtain without my being in a position to know that it does. For I might have insufficiently strong evidence (only a vague memory of having parked the car on Main Street), or misleading evidence (a false report of its having been stolen). Like countless other conditions—*that it's raining in Paris, that Peter is coming to the party, that the President is in New York*—I can fail to know whether *CAR* obtains simply because the world and my position in it preclude such knowledge. The radicalism of Anti-Cartesianism lies in its insistence that what is true of *CAR* is true of all conditions. This includes even paradigmatically Cartesian conditions like *feeling cold, possessing evidence that p, believing that p* and *its seeming to one that p*.

Why endorse Anti-Cartesianism? Williamson's argument (2000, ch. 4) is roughly this. Consider the condition of *feeling cold*. Since feelings of cold come in degrees of strength, it is possible to be in what we might call a *liminal case*⁵ of feeling cold: a case of feeling cold that is nearly not a case of feeling cold. Imagine that *S*, a normal agent like ourselves, is in such a liminal case α of feeling cold: if *S* were just slightly warmer than she actually is, it would be untrue that she feels cold. As it happens, in α *S* has the true belief that she feels cold. Is this true belief knowledge? Williamson argues not. For *S* (a normal human) has dispositions to believe that are not perfectly fine-grained. That is, there plausibly exists a very similar possible case β in which *S* has the untrue belief that she feels cold. Knowledge requires safety from error: for *S* to know in α that she feels cold, she cannot have in a sufficiently similar case an untrue belief that she feels cold. But there exists such a

sufficiently similar case of untrue belief: β . As such, in α S cannot know that she feels cold. Feeling cold is thus not luminous. Equally, we could imagine a case in which S is in a liminal case of *not* feeling cold; by analogous reasoning, S would be unable, in such a case, to know that she was not feeling cold. And so feeling cold is not absent-luminous, either. In general, safety from error precludes knowledge of whether one is in a condition C whenever one is in a liminal case of C's obtaining or not obtaining.

Many have found Williamson's argument unconvincing, suspecting some philosophical sleight-of-hand.⁶ So it's worth noting that Alison Gopnik (1983), Peter Carruthers (2011) and Eric Schwitzgebel (2008, 2011) all argue that the empirical psychology demands a rejection of Cartesianism. All three argue that experiments show that we can be mistaken, even in favourable conditions of sustained reflection, about our own mental states. And according to Schwitzgebel at least, these failures of Cartesianism are neither marginal nor unusual; he argues that our grasp on our mental lives—our visual imagery, our emotions, our tactile experience—is systematically *worse* than our grasp on the external world. So not only are we wrong to think we have privileged first-personal access, we cannot even retreat to the weaker Cartesian thesis that we have at least *greater* epistemic access to our minds than to external states of the world.

For my own part I think the most powerful reason to embrace Anti-Cartesianism is (not unironically) introspective. I sometimes find myself uncertain, even after careful consideration, about my own phenomenology: whether I'm angry or merely annoyed, whether I'm desirous or indifferent, whether I believe or am agnostic. Of course, the uncertainty at issue here is not an uncertainty about whether my phenomenology is *thus*: I'm always in a position to know that I'm feeling just *this*, the way I'm feeling. Instead, the uncertainty lies in the categorisation of my phenomenal experience under the appropriate concepts: *anger, annoyance, desire, indifference, belief, agnosticism*. My own introspective feelings of uncertainty deepen when we move to conditions of particular philosophical interest, such as *my having a credence x in p* or *p's having probability x on my evidence*. For these conditions, I very often feel that no amount of assiduous introspection will reveal whether they obtain. Perhaps not everyone finds her own phenomenology so recalcitrant to epistemic grasp. But to the extent that this kind of experience strikes one as familiar, one has reason, beyond the armchair and empirical arguments, to embrace Anti-Cartesianism.

Of course, embracing Anti-Cartesianism doesn't mean accepting the sceptical verdict that one is *never* in a position to know whether one is in various mental state conditions—any more than it means embracing the sceptical verdict that one can never know the location of one's car. One very often is in a position to know whether one's car is on Main Street, and one might very well often be in a position to know whether one is in various mental state conditions. We might say that conditions like CAR and feeling

could be *contextually transparent*, meaning that at certain contexts one can know whether they obtain. But contextual transparency is not the same as transparency *simpliciter*. According to Anti-Cartesianism, no conditions are transparent *simpliciter*.

While I think that there is compelling reason to embrace Anti-Cartesianism, those of a strongly Cartesian persuasion are invited to read this paper as offering further reason to cleave to the Cartesian orthodoxy—indeed, even to read it as a *reductio* of Anti-Cartesianism.

3.

Why think Anti-Cartesianism has any interesting implications for the normative sphere? Many normative theorists maintain that one or more fundamental normative notions—*epistemic justification*, *the subjective ought*, *rationality*—must be spelled out in terms of subjects' mental states. For example, epistemic internalists think that whether one is epistemically justified or epistemically rational supervenes only on one's mental states. Ethical subjectivists think that whether one ought to perform some action supervenes on the evidence that one possesses.⁷ An obvious question is: why think that these normative notions should be spelled out in this way? Why not think, as epistemic externalists do, that the justification of one's belief can supervene on non-mental facts, e.g. whether that belief is the product of a mechanism adequately hooked up to the world? And why not think, as ethical objectivists do, that the important sense of 'ought' is the one that supervenes on *all* the facts, not only the ones within one's epistemic ken? An obvious answer is this: only mental states possess a crucial property, namely transparency,⁸ and only norms that feature transparent states can meet some basic desiderata of norms. (This answer is explicitly rejected by 'mental state' internalists, who accept the mental supervenience claim but are neutral or negative about Cartesianism; my focus here is only on the 'access' brand of internalists, i.e. those internalists who accept both the mental supervenience and transparency theses.)^{9,10} Insofar as this diagnosis of the motivations for internalism and subjectivism is correct, then Anti-Cartesianism undermines the case for tying the normative to the mental.

Consider a norm that is not tied to the mental: the classic act-consequentialism norm that says that one ought to perform whichever action would maximise the good.¹¹ A familiar complaint about the norm is that one isn't always in a position to *know* which action would in fact maximise the good.^{12,13} One's evidence might, after all, be insufficient or misleading on this score. A similar complaint is made about externalist norms in epistemology, for example the reliabilist norm that says that one's belief is justified only if it was formed by a reliable mechanism. One isn't always in a position

to *know* whether one's belief was formed by a reliable mechanism; again, one's evidence on that score might be insufficient or misleading.

Such considerations might prompt one (and indeed have prompted many) to endorse rival norms that are more closely tied to the mental—that are appropriately *subjectivised* or *internalist*. For example, we might say that one ought to perform whichever action one *believes* would maximise the good, or whichever action would *probably* maximise the good. And we might say that one's belief is justified only if its content is *probable on one's evidence*. But if Anti-Cartesianism is true, then these emendations won't solve the problem with which we began—namely, that we aren't always in a position to know what the relevant norms demand of us. For it follows from Anti-Cartesianism that one is not always in a position to know which action one *believes* would maximise the good, or which action would *probably* maximise the good. And it similarly follows that one is not always in a position to know whether one's belief is probable on one's evidence. Indeed, if Anti-Cartesianism is true, then the same thing, *mutatis mutandis*, could be said of any norm whatsoever.¹⁴ Thus a complaint that is usually thought to be the special bugbear of objectivist and externalist norms turns out, if Anti-Cartesianism is true, to be a problem for all norms.

I'll be more precise. A norm is a universal generalisation about how an agent *S* is obligated or permitted to act, of the schema $\langle S \text{ (ought)/(is permitted) to } \phi \text{ iff } C \rangle$,¹⁵ where ϕ ranges over the basic actions available to *S*, and *C* is the norm's *triggering condition*. For example, objectivist act consequentialism can be (roughly) expressed as the norm $\langle S \text{ ought to } \phi \text{ iff } \phi\text{-ing would maximise the good} \rangle$. Here, the relevant triggering condition is *that ϕ -ing would maximise the good*. Similarly, the internalist epistemic norm mentioned above can be expressed as the norm $\langle S \text{ is permitted to believe } p \text{ iff } p \text{ is probable on her evidence} \rangle$. Here, the relevant triggering condition is *that p is probable on *S*'s evidence*.

Anti-Cartesianism implies that whatever a norm's triggering condition, it is non-transparent. For an obligatory norm, such as the consequentialist norm, this means that a competent agent who knows the norm cannot always know which of the basic actions available to her would result in conformity with the norm.^{16,17} For a permissive norm, such as the epistemic internalist norm, it means that a competent agent who knows the norm cannot always know whether a given basic action would result in conformity with the norm.¹⁸ Generalising across both obligatory and permissive norms, we can call a norm *lucid* just in case a competent agent who knows the norm is in a position to know of every basic action available to her whether it would be in conformity with the norm.^{19,20} Thus Anti-Cartesianism implies:

ANTI-LUCIDITY: there is no norm such that a competent agent who knows the norm is in a position to know of every basic action available to her whether it would be in conformity with that norm

It's important to note that Anti-Lucidity doesn't mean that there are no norms such that competent agents can *sometimes* (or even *often*) be in a position to know of every available action whether it would be in conformity with the norm. Just as Anti-Cartesianism is compatible with some conditions being *contextually* transparent, Anti-Lucidity is compatible with some norms being contextually *lucid*.

If there are no Cartesian conditions, then there are no lucid norms. A lack of lucidity is typically associated with objectivist and externalist norms. But if Anti-Cartesianism is true, then what is often thought to be a parochial problem turns out to be a global one.

But is it really a problem? After all, some consequentialists insist that we ought to do whatever would maximise the good while happily admitting that we sometimes fail to be in a position to know which action that might be.²¹ Kant seems to have embraced a similar conclusion about the categorical imperative: we ought to act in accordance with it, but we are hardly guaranteed to know whether we are so doing. (Indeed he seemed to think that we were *never* in a position to know this.) Epistemic externalists concede that we aren't always in a position to know whether we are acting in accordance with our epistemic obligations or prohibitions, but they also think that in many such cases we're excused for our errors.²² Thus some will find little to worry about in Anti-Lucidity.

Others, though, will feel differently. Others will feel that Anti-Lucidity threatens the very project of articulating an adequate normative theory. How so? When subjectivists and internalists explain the importance of lucidity for norms²³—and thus why, on their view, norms must be tied to the mental—they typically appeal to one of two considerations. The first is that a normative theory should be *action-guiding*. The second is that a normative theory must imply that competent agents can exercise *control* over their normative performance in a way that secures an alignment between the facts about what agents ought or are permitted to do (the *deontic* facts) and the facts about what agents are blameworthy for doing (the *hypological* facts).²⁴ I will discuss these considerations, respectively, in sections 4 and 5. In section 4 I will suggest that, on a commonsense notion of guidance, the desideratum that norms be guiding can in fact be met without lucidity. Thus insofar as some are motivated to tie the normative to the mental in order to achieve action-guidance, Anti-Cartesianism is plausibly not a threat. But in section 5 I will suggest that the desire to align the deontic and hypological does not fare so well in light of Anti-Cartesianism. Insofar as this desire motivates internalists and subjectivists, Anti-Cartesianism poses a significant problem.

4.

One often hears from subjectivists and internalists the demand that norms be *action-guiding*, *followable* or *useful*.²⁵ This demand usually issues

from an insistence on the ‘first personal’ role of norms, as things that not only provide us with third-personal metrics for assessing others and states-of-affairs, but also as things that tell us how to think and act—that guide and advise us. We might think that this demand is one that non-lucid norms fail to meet. After all, if we can’t always know whether a norm’s triggering condition obtains, how could we follow, be guided by or use such a norm? Thus Jim Pryor writes, explicating the internalist demand that epistemic norms be tied to the mental:²⁶

If a belief-guiding recipe [of the form ‘In circumstances C, believe *p*’] is to be *usable* in deciding what to believe, then the circumstances C it refers to must be circumstances such that the subject *can tell* whether they obtain, when he’s following the recipe (2001, 116).

Frank Jackson argues against objectivist ethical norms along similar lines:

[T]he fact that a course of action would have the best results is *not in itself a guide to action*, for a guide to action must in some appropriate sense be present to the agent’s mind. We need, if you like, a story from the inside of an agent . . . and having the best consequences is a story from the outside (1991, 466–7, italics added).

As does James Hudson:

An objective theory lays down conditions for right action which an agent may often be unable to use in determining her own behaviour. In contrast, the conditions for right action laid down by a subjective theory guarantee the agent’s ability to *use them to guide her actions* (1989, 221).

And Allan Gibbard:

It is clear enough why we should want a theory of what kinds of acts are right in the subjective sense. Such a theory offers moral guidance: even when we know we are ignorant of the relevant facts, we can use the theory, together with what we think we do know, to decide what acts to avoid on moral grounds. Why, though, should we want a theory of what kinds of acts are right in the objective sense? Such a theory offers no guidance when we know we are ignorant of relevant facts; in that case we need rules for acting without full information (1990, 43).²⁷

If followability and action-guidance are adequacy conditions on norms, then Anti-Cartesianism seems to imply that there are no adequate norms. After all, if no norms are lucid, then no norm is such that one is always in a position to draw on one’s knowledge of the norm²⁸ to deduce which of the available actions would be in conformity with the norm. For any obligatory norm, one can know the norm and yet fail to be in a position to know which

basic action one must perform in order to conform with it. And for any permissive norm, one can know the norm and yet fail to be in a position to know of the available basic actions which, if any, are permitted by the norm. Insofar as one thinks that norms ought to be action-guiding or followable, Anti-Cartesianism might seem to pose a serious problem.²⁹

A different version of the action-guidance demand focuses on the notion of *acting for a reason*. Here is Bernard Williams: “If there are reasons for action, it must be that people sometimes act for those reasons” (1981, 102).³⁰ The claim is that, conceptually, reasons are the sort of things that agents can act *for*, and that can *explain* their actions; and, thus, that only facts that are within the agents’ epistemic ken can serve as reasons. This version of the action-guidance demand is compatible with Anti-Cartesianism. To see how, take the norm R: <S ought to phi iff S has all things considered reason to phi>. If Anti-Cartesianism is true, it follows that S is not always in a position to know what she has all things considered reason to do. So S cannot use R to guide her actions, in the sense of using her knowledge of R to deduce which of the actions available to her would be in conformity with R. But this is compatible with the claim that facts must be within S’s epistemic ken to constitute reasons for S. Another way of putting the point is this: if Anti-Cartesianism is true, then there are no non-trivial general truths about what *kinds* of facts constitute our reasons. (We cannot say, for instance, that facts about our *beliefs* or our *evidence* always constitute reasons for us, since neither our beliefs nor our evidence are always within our epistemic ken.) But which facts, at a given context, constitute our reasons might nevertheless be a function of what we are in a position to know in that context.

What about the related idea that a normative theory should consist in both a *criterion of right* and a *decision procedure*? Many consequentialists draw this distinction in response to the worry that the objectivist consequentialist norm is insufficiently action-guiding.³¹ The consequentialist norm is said, then, to be merely a ‘criterion of right’, a specification of what makes a given action right or wrong,³² while it is to the decision procedure that agents should appeal for action-guidance.³³ Candidate decision procedures range from following the dictates of commonsense morality (Sidgwick 1907, 361; Moore 1903/1993, 206–7; Hare 1981, 121) to employing some sort of probabilistic calculus. But whatever the decision procedure, it follows from Anti-Cartesianism that it won’t be lucid. For any decision procedure, it is possible that a competent agent who knows the decision procedure will be unable to use it to guide her actions. Thus Anti-Cartesianism threatens to make the distinction between a criterion of right and a decision procedure unprincipled.³⁴

That all said, it’s not entirely obvious that, at least as we ordinarily deploy the concepts, *followability* and *guidance* really do require lucidity. Consider the following norm: <When setting the table for Passover, one ought to set as many places as there will be Seder guests plus one>. Surely this norm counts

as followable and action-guiding. But it is plainly non-lucid. For there are occasions when one is simply not in a position to know how many people will come to one's Seder: a guest might fail to show up, or a surprise guest might arrive. In such cases, one is unable to know how many places to set through a deduction that uses the Seder norm as one of its premises. But it seems at least somewhat odd to dismiss the Seder norm as non-guiding or unfollowable.³⁵ In many situations, after all, one *does* know how many Seder guests one will have, and one is able to use this knowledge to deduce how many places to set. Here, at least, mere contextual lucidity seems to suffice for followability and guidance.

This consideration might prompt us to refine our demand: norms need not be lucid simpliciter, but they must be *contextually* lucid. That is, there must exist at least *some* contexts in which agents can use the norm to guide their behaviour. This refined demand might however feel too minimal. Recall that the importance of action-guidance is taken by many as a reason to favour internalist and subjectivist norms over their externalist and objectivist rivals. But if all that action-guidance requires is mere contextual lucidity, then many externalist and objectivist norms will satisfy that desideratum. For example the reliabilist norm—that one's belief is justified only if it was formed by a reliable mechanism—is contextually lucid. In many cases, agents are able to know whether their beliefs were formed by reliable mechanisms. (Consider our perceptual beliefs.) Thus appealing to this refined notion of action-guidance will not tell decisively against objectivist and externalist norms. A further refinement is to say that for norms to be adequately guiding, they must be contextually transparent *at the relevant contexts*. Then one might say of the reliabilist norm that it is non-transparent at precisely those contexts that we care about—i.e. contexts in which agents have insufficient or misleading evidence about the reliability of their belief-forming mechanisms, as with the brain-in-a-vat or Bonjour's clairvoyants.³⁶ We might think that these are the contexts that are made salient when we ask the questions that drive normative theorising: "How should I act?" or "What should I believe?" In this way the subjectivist or internalist could appeal to action-guidance as a desideratum that tells against rival norms.

But does action-guidance and followability even require contextual lucidity? Consider an opaque norm, for example, a norm that applies to all ticket-holders in a fair lottery: <One ought to buy a Ferrari today iff one's ticket will be drawn as the winner tomorrow>. Plausibly, one is never in a position to deduce whether buying or not buying a Ferrari would be in conformity with the norm, since one is never in a position to know whether one's ticket will be a winner. Nonetheless, one is generally in a position to know that one's ticket will *probably* lose—and thus that one can *probably* conform to the norm by abstaining from buying a Ferrari today. Despite its opacity, the Ferrari norm seems plausibly to count as guiding and followable. More generally, it seems to me that for a norm N intuitively to count as guiding at

a context, it must have some non-arbitrary relationship to a norm N^* that *is* lucid at the context. (If N is itself lucid at the context, then the non-arbitrary relationship to N^* might be the identity relationship.) The difficult question is: what relationships, apart from identity, count as non-arbitrary for this purpose?³⁷

Whatever the correct answer to that question—that is, just what it is that makes a norm followable—we have reason to think that norms that count as ‘action-guiding’ or ‘followable’ in the mundane sense are not necessarily norms that are lucid, nor even contextually lucid. If so, then Anti-Cartesianism does not pose a threat to action-guidance as a desideratum of norms. It does, however, pose a threat to the subjectivist/internalist strategy of appealing to action-guidance as a way of casting impunity on rival norms. For if lucidity isn’t required for action-guidance, then action-guidance isn’t obviously a reason, after all, to prefer subjectivist or internalist norms to their objectivist or externalist rivals. If, however, action-guidance does require lucidity, then Anti-Cartesianism implies that internalist and subjectivist norms are on all fours with externalist and objectivist norms. Those who wish to tie the normative to the mental because of action-guidance considerations are thus faced with a dilemma.

5.

The second consideration that is advanced in favour of lucidity as a normative desideratum is the putative intimacy between the *deontic* facts—the facts about what agents are obligated or permitted to do—and the *hypological* facts—the facts about whether agents are blameworthy for their actions. Ross writes that “the notion of obligation carries with it very strongly the notion that the non-discharge of an obligation is blameworthy” (Ross 1939, 163). In other words, agents who violate their obligations are *ipso facto* blameworthy. If so, argues Ross, some version of ethical subjectivism must be right: only if the circumstances that determine our obligations are transparent to us can it be that our failures to meet our obligations are indeed blameworthy.³⁸ Making an analogous case for epistemic internalism, Steup writes that: “[a] belief that is epistemically justified is a belief that is epistemically permissible, a belief for which the subject cannot justly be blamed” (1999, 312). Again, the thought is that the triggering condition of the justification-norm must be transparent to agents, as otherwise it would be possible for agents to blamelessly acquire unjustified beliefs.^{39,40}

Some normative theorists will reject out of hand Ross’ claim that normative violation implies blameworthiness on the grounds that one can blamelessly violate any norm whatsoever through incapacity (e.g. severe cognitive disability) or through non-culpable ignorance of the genuine norms (e.g. being brought up in a morally backwards environment).⁴¹ For the sake

of argument I want to grant that incapacity and non-culpable normative ignorance can excuse normative violation.⁴² So we should amend Ross' claim as follows:

ALIGNMENT VIEW: if a competent agent knows a norm and yet violates it, then she is *ipso facto* blameworthy

According to the Alignment View view, it is impossible for a competent, normatively knowledgeable⁴³ agent to violate a normative obligation or constraint without thereby being blameworthy. Put differently, the Alignment View holds that the *only* possible excuses for normative violation are either incapacity or ignorance of the norms.⁴⁴

If the Alignment View is right, then it must be the case (barring normative nihilism) that there exist some norms such that competent, normatively knowledgeable agents can never blamelessly violate them. Indeed one might suppose, as many internalists and subjectivists do, that internalist and subjectivist norms are just these sorts of norms. But if Anti-Cartesianism is true, then this supposition seems mistaken. Plausibly, no norms can satisfy the Alignment View. Or so I shall argue.

Aristotle wrote that acts “are thought involuntary, which take place by force or by reason of ignorance” (1980, 1110b 17). In what way can ignorance, like force, render an action outside the scope of one's voluntary control? Suppose Fred is choosing whether to donate some sum of money to Charity A or Charity B. Suppose further that the norm appropriate to this choice is a norm that says that one ought to donate to whichever charity is more effective—and that Fred knows this.⁴⁵ Finally, suppose that after thorough research, Fred's total evidence suggests, falsely, that Charity B is more effective than Charity A. Fred, eager to comply with the norm—or perhaps just disposed to do so—gives his donation to Charity B. He thereby fails to do what he ought have done. Intuitively we want to say: *Fred was here a victim of bad luck; what happened was out of his control; he is not to be blamed.*

This case shares something in common with the canonical cases discussed in the moral luck literature.⁴⁶ Consider for example George who, trying to save his ailing dog, treats her with a medicine he justifiably believes will save her. In fact the medicine is toxic; the dog dies. Intuitively we want to say: *George was here a victim of bad luck (more so the dog); what happened was out of his control; he is not to be blamed.* But what exactly is it that is supposed to be outside Fred and George's control such that neither is to be blamed? Both undertook voluntary, conscious actions—donating to Charity B, giving the dog the medicine.⁴⁷ Moral luck theorists typically say that what was not under George's control was the *result* of his (voluntary) action; what we have here is a case of “resultant luck” (Nagel 1979) or “outcome luck” (Scanlon 2008). But we can't quite say that in Fred's case, since it is not the result or outcome of Fred's donation, but the *very act of donating*, that is

(we are supposing) the moral violation. In Fred's case, we cannot factorise out so simply the voluntary action and its involuntary effects. What we lack is some way of describing the bad luck that is present in both George and Fred's cases.

What is not under Fred's control is whether his voluntary action constitutes a moral violation. That is, Fred's action is voluntary under the (basic) description *donating to Charity B*, but not under the (non-basic) description *committing a moral violation*. The same thing can be said of George's action: it is voluntary under the (basic) description *giving medicine to his dog*, but involuntary under the (non-basic) description *committing a normative violation*. Now, since we don't get to decide what the norms are, there's a sense in which it is never within our control whether a given action constitutes a normative violation; it is outside of my control that, if I torture an animal for fun, I will have done something I ought not. But when we know which basic action(s) a norm calls on us to do or not do, then it is within our control whether whatever we choose to do falls under the description *committing a normative violation*. Given that I know that torturing a cat would be a normative violation, it's within my control to commit one or not.

A non-lucid norm is one such that agents cannot always exercise control over whether what they do constitutes a violation of that norm. What lies beyond the sphere of control is luck; thus one's performance vis-à-vis a non-lucid norm is susceptible to a kind of bad luck. Fred, despite trying his best or being disposed to conform to the moral norm governing his donation, ends up violating it through ignorance of the relevant non-normative facts. The same is true of George. They both violate the relevant norm through what we might call *bad performance luck*. In general, for any non-lucid norm, it is possible to end up violating it through bad performance luck.^{48,49} Thus Anti-Cartesianism implies:

PERVASIVE PERFORMANCE LUCK: There are no norms such that an agent's performance vis-à-vis that norm is immune from violation through bad performance luck

In his postscript to "Moral Luck", Williams wrote that "what is not in the domain of the self is not in its control, and so is subject to luck" (1993, 20). Williams' Cartesian identification of the realm of the self with the realm of agential control offered both ethicists and epistemologists a safe harbour from performance luck. But if Anti-Cartesianism is correct, the harbour is hardly safe.

Pervasive Performance Luck threatens the Alignment View.⁵⁰ In the Fred and George cases we intuitively wanted to say that because their actions were "a matter of bad luck", Fred and George aren't blameworthy for their normative violations. After all, they tried their best to do what was normatively required of them, and only failed to do so because of bad luck. Plausibly,

what matters for our hypological status is the quality of our will, not the way that our will is manifested as a function of that which lies beyond our control.^{51,52} If so, then any norm that is violable through bad luck is also a norm that can be violated blamelessly.⁵³ It follows from Anti-Cartesianism then that all norms can be blamelessly violated by normatively knowledgeable, competent agents.⁵⁴ Thus we have:

BLAMELESS VIOLABILITY: Any norm can be blamelessly violated by a competent agent who knows the norm

If Anti-Cartesianism is true, then the Alignment View must be false: there cannot be any norms that are immune from blameless violation by competent, knowledgeable agents. It follows that ethical subjectivism and epistemic internalism cannot be defended on the grounds that they secure the Alignment View. For no norms do.⁵⁵

What's interesting about Anti-Cartesianism in this respect is that it gives us a novel reason to check our impulse to align the deontic and hypological facts. As I mentioned, the more standard reasons have to do with the elusiveness of the *normative*: non-culpable ignorance about the general normative truths, and the possibility of normative incapacity. But here we have a more mundane reason to reject the desire to align these two sets of facts: the pervasive possibility of *non-normative* ignorance. It follows from Anti-Cartesianism that we must accept some inevitable measure of alienation between the deontic and the hypological. Even in a sort of normative utopia—where all agents were competent, everyone knew the genuine normative truths, and everyone was maximally motivated or disposed to fulfil their normative obligations—we would still have deontic violations. Even in such a utopia, normative performance would not lie entirely within our control.

6.

I have suggested that insofar as one is motivated to tie the normative to the mental out of a desire to articulate norms that are strongly action-guiding or that secure a tight connection between the deontic and the hypological, Anti-Cartesianism poses a problem. If Anti-Cartesianism is true, then there are no norms that are always action guiding or that are never blamelessly violable. If we insist—as many internalists and subjectivists do—that norms be perfectly guiding and immune from blameless violation, we are forced to embrace normative nihilism: the view that there are no norms.

An obvious retreat for internalists or subjectivists who want to accept Anti-Cartesianism is to weaken their demands. They do not demand norms, they might say, that are *always* action-guiding, nor do they want norms that *perfectly* satisfy the Alignment View. They simply want norms that are

sufficiently action-guiding, and secure a *sufficient* intimacy between the deontic and the hypological (cf. Prichard 1949/2002, 94). Action-guidance and deontic-hypological intimacy come in degrees: some norms score higher on these counts than others, just as some norms are simpler or more intuitive than others. The internalist or subjectivist might then say that degree of action-guidance and degree of deontic-hypological intimacy are just two desiderata for normative theory choice, along with other familiar desiderata such as simplicity, intuitiveness, and so forth. The argument for internalist and subjectivist norms, then, isn't that they satisfy action-guidance and deontic-hypological intimacy requirements while externalist and objectivist norms do not. It's rather that internalist and subjectivist norms, in the final analysis, have the better balance of theoretic virtues.

I want to say two things about this sort of response. The first is that if internalists and subjectivists want to embrace it, then certain dialectical moves are no longer available to them. For the simple objection that objectivist or externalist norms aren't always action-guiding or can be blamelessly violated is, if Anti-Cartesianism is true, a *tu quoque*. Subtler objections—that treat these issues as a matter of degree, not kind—are called for. Given how often complaints about action-guidance and blameless violation are used as a decisive, principled reason to reject externalist and objectivist theories, such a dialectical shift would not be insignificant.

The second thing I want to say is this. Anti-Cartesianism comes in different flavours. Williamson's anti-luminosity argument seeks to establish only the minimal thesis that even mental state conditions can be non-transparent. But recall that Schwitzgebel argues for a stronger claim, on empirical grounds: namely, that we are often in a *better* position to know external world conditions than we are to know our own mental state conditions. If that's so, then there is no *a priori* guarantee that internalist and subjectivist norms will turn out to possess a higher degree of action-guidance or deontic-hypological alignment than externalist or objectivist ones. Indeed, it looks like the question is largely an empirical one. We can't rule out from the armchair the possibility that norms not tied to the mental—i.e. objectivist and externalist norms—might better meet the aspirations that motivate many internalists and subjectivists.

7.

Without the comfort of Cartesian privilege, our relationship to the normative realm is, I have suggested, fraught with ignorance and luck. What sort of positive conception of the normative does this leave us with? Here we might turn to the Ancients, especially Aristotle. In the *Nicomachean Ethics*, Aristotle tells us that *Eudaimonia* “needs the external goods as well; for it is impossible, or not easy, to do noble acts without proper equipment” (1980,

1099a31–33). To be good is to do well, Aristotle says, and to do well requires good fortune.

We do not generally share the Greeks' comfort with luck; we are often determined to think that our normative standing, if nothing else, is within our control. In his essay "On Misunderstanding the *Oedipus Rex*" (1966), E.R. Dodds laments how this modern discomfort with luck yields poor interpretation of Sophocles' play. According to the common reading that is Dodds' target, Oedipus' downfall is due to a moral failing—his pride or overconfidence, perhaps, a 'fatal flaw'. Not only is that reading unwarranted by Sophocles' text, Dodds argues, but it is also unsupported by Aristotle's use of *Oedipus* as a canonical example of tragedy. Aristotelian *hamartia* (literally, 'missing of the mark')—of which Oedipus, along with Thyestes, is the central example for both Aristotle (*Poetics* 13) and Plato (*Laws* 838c)—should not be interpreted as 'fatal moral flaw' but rather as ignorance of a non-moral fact that leads to wrongdoing. In Oedipus' case, this is ignorance of the identity of his parents. In the case of Thyestes, it is ignorance that what seems to be butcher's meat is actually the flesh of his children. The natural moral laws prohibit parricide, incest and cannibalism. But these laws are non-lucid. One cannot always be guided by them, and one can violate them through bad performance luck. The tragedy of Oedipus, then, is not of a great man brought down by moral weakness. It is rather of a great man brought down (to put it with thundering banality) by normative non-lucidity.

Like us, the Ancient Greeks recognised a distinction between those moral harms performed by bad performance luck and those performed under agential control. Nonetheless, the all-important moral state of *pollution* is one that taints even the unlucky:

Thyestes and Oedipus are both of them men who violated the most sacred of Nature's laws and thus incurred the most horrible of all pollutions; but they both did so without *πονηρία* [*poneria*, vice], for they knew not what they did—in Aristotle's quasi-legal terminology, it was a *ἁμάρτημα* [*hamarthma*], not an *ἀδίκημα* [*adikhma*, moral injustice]. That is why they were in his view especially suitable subjects for tragedy. Had they acted knowingly, they would have been inhuman monsters, and we could not have felt for them that pity which tragedy ought to produce. As it is, we feel both pity, for the fragile estate of man, and terror, for a world whose laws we do not understand (Dodds 1966, 39–40).

Anti-Cartesianism invites us to return to a more tragic outlook of the normative. Another reason for such a return is itself normative. The impulse to eradicate luck from the normative sphere often issues from a laudable discomfort with the role that luck plays in other spheres of human life, especially its intimate connection with social and political inequality. Understandably, then, we hope to establish that *one* realm of human life is immune from luck – that is, the normative itself. But this philosophical longing might in the

end be morally neutering: a revolt against injustice in theory, rather than in practice. In the place of such a longing, we might instead have a post-Cartesian normative theory that is more attentive to the fragility of our human estate.

Acknowledgments

Special thanks to John Hawthorne, Timothy Williamson, Selim Berker, Derek Parfit, David Wiggins, David Velleman, Nishi Shah, Paul Boghossian, Crispin Wright, Jennifer Nagel, Jane Friedman, Kian Mintz-Woo, Robert Simpson, Fred Wilmot-Smith, Errol Lord, Neil Mehta and Ram Neta for reading and commenting on earlier drafts. Thanks also to Corine Besson, Akeel Bilgrami, Dave Chalmers, Cian Dorr, Daniela Dover, Jeremy Goodman, Anandi Hattiangadi, Harvey Lederman, Nakul Krishna, Paul Myerscough, Andreas Mogensen, Jim Pryor, Wlodek Rabinowicz, Regina Rini, Daniel Rothschild and Jeff Russell for helpful discussion of these issues. This paper draws on a paper I co-authored with John Hawthorne, “Disagreement Without Transparency: Some Bleak Thoughts” (2013), in Christensen D. and J. Lackey, (eds.), *The Epistemology of Disagreement* (Oxford: OUP). Earlier versions of it have been presented at the Princeton Varieties of Normativity workshop (2014), UCL (2014), Oxford (2014), the All Souls Epistemology Workshop, Oxford (2013), the DPhil Dissertation Seminar, Oxford (2013), the Objective and Subjective Oughts seminar, Oxford (2013), the Epistemology of Philosophy Conference, Cologne (2011), the Mind and World Conference, IUC-Dubrovnik (2011) and the Oxford Philosophy Faculty women’s event (2011). Many thanks to everyone in those audiences as well.

Notes

1. See Schwitzgebel (2011) for a discussion of philosophers who are committed to privileged first-person access. See also various essays in Coliva (2012) for defences of a Cartesian view of the phenomenal. I should say that it’s far from obvious that Descartes was himself in this sense Cartesian.
2. See Suikkanen (2008) and Wedgwood (2010) for a discussion of a related claim: namely, that if Williamson’s anti-luminosity argument is sound, then Dummettian realism about ethics is false: i.e. there must be some unknowable ethical truths.
3. Are there exceptions? Williamson (2000, 107) points out that impossible conditions are vacuously luminous. Take the condition of *its being the case that 2 plus 2 equals 5*. Since that condition never obtains, it is trivially true that whenever it obtains we are in a position to know it obtains. But even such conditions won’t be absent-luminous, and so won’t be transparent; plausibly, we can fail to be in a position to know that 2 plus 2 doesn’t equal 5—say, if we have strong

evidence that we're on a drug that prevents one from doing simple sums. What about a condition that obtains in every case, i.e. the necessary condition? Such a condition is vacuously absent-luminous, since it is trivially true that whenever it doesn't obtain we are in a position to know it doesn't obtain. But the necessary condition won't be luminous since it could be presented under a guise that prevents us from seeing its tautologous nature, e.g. "water is H₂O". What about conditions that are constitutively connected to doxastic conditions? Many have argued that if there are such conditions, then they are exceptions to Anti-Cartesianism (Weatherson 2004, Berker 2008 and Ramachandran 2009). If so, then a certain (constructivist, anti-realist) metanormative view—namely one that ties the normative facts to the facts about what we believe about the normative facts—might seem a refuge from the threat of Anti-Cartesianism. For an argument that that even constitutive conditions don't constitute counterexamples to Anti-Cartesianism—and thus that no metanormative view is immune from anything I say here—see Williamson (2000, 99 ff.) and Srinivasan (2015).

4. What is it to be in a "position to know"? This is what Williamson has to say: "To be in a position to know *p*, it is neither necessary to know *p* nor sufficient to be physically and psychologically capable of knowing *p*. No obstacle must block one's path to knowing *p*. . . . The fact is open to one's view, unhidden, even if one does not yet see it" (Williamson 2000, 95). The notion is admittedly vague. I don't think this creates damning problems for anything I say here, though perhaps I'm being overly optimistic.
5. By 'case' here I mean a centred possible world: a possible world with a designated subject and time.
6. For a defence of it against some of the most common lines of criticism, see Srinivasan (2015).
7. Many things are meant by the term 'subjectivism' in ethics. I am following Shelly Kagan (1998, 65 ff.) and various others in this usage. In particular, I'm not using 'subjectivism' to pick out a meta-ethical position.
8. For an alternative (and I think convincing) answer to the question as it relates to epistemic norms, see Schoenfield (this volume).
9. Thus I use 'internalism' throughout to refer to the traditional, 'access' variety of internalism (see, e.g. Chisholm 1977, ch. 6, section 5; BonJour 1985, chs. 1–2). For a discussion of mental state internalism, see Conee and Feldman (2001). For a discussion of varieties of internalism and their motivations, see Pryor (2001). While my focus here is not mental state internalism, I think a worthwhile question is: why, if we aren't committed to Cartesianism, should epistemic justification or rationality supervene only on mental states? Wedgwood (2002) argues that we should accept the supervenience thesis (about epistemic rationality, not necessarily epistemic justification) because complete folk psychological explanations of belief revision invariably feature non-factive mental states. Even if Wedgwood is right about the habits of folk psychological explanation, I'm unsure how much philosophical weight this can carry; one might well worry that folk explanations bottom out at certain non-factive mental states because the folk tacitly presumed that these states are transparent. Why else do folk explanations make epistemic behaviour, in Wedgwood's terms, "intuitively intelligible" and "unsurprising" (357)? Wedgwood also suggests that his view of epistemic rationality can vindicate the intuition that whether one's belief is rational is a matter of

control, not luck (355); this is because rationality consists in ‘basic rules’ that can be ‘followed directly’. Again, one might suspect that folk intuitions about what rules can and cannot be ‘followed directly’ presuppose a commitment to Cartesianism.

10. See fn. 55 for a discussion of subjectivist views that do not presuppose transparency.
11. I’m assuming here and throughout that there is a unique action that fits this description.
12. A stronger version of this claim holds that we are *never* in a position to know which basic action would maximise the good. See, e.g. Lenman (2000) and Kagan (1998). For an argument that all norms share this feature, see Prichard (1932/2002, 88) and Ross (1939, 149ff).
13. cf. Kagan (1998, 64); Lenman (2000); Prichard (1932/2002); Ross (1939, 146–167); Ewing (1947, 128); Gruzalski (1981); Smith (1988); Hudson (1989); Broome (1991, 128); Griffin (1996, ch. 7); Howard-Snyder (1997); Scanlon (2001, 2008); Mason (2003); Wedgwood (2003, 204); Timmons (2012, 10–11).
14. Pettit (2003, 99) makes a similar suggestion, at least for leading non-consequentialist theories.
15. I mean the norm “S ought to phi iff C” to be read as “If C, S ought to phi, and if not-C, S ought not phi”. I mean the norm “S is permitted to phi iff C” to be read as “If C, S is permitted to phi; if not-C, S is not permitted to phi”. For the sake of simplicity I am setting aside non-biconditional norms.
16. In the concealed question sense of “knows”, i.e. “knows what the norm *says*”. A norm need not be true to be in this sense known.
17. I use the term ‘basic action’ to mean an action that a competent agent knows how to perform and also is in a position to know whether she is performing it. It follows from Anti-Cartesianism that there are no true generalisations about which actions are basic, since there are no actions (under whatever guise) such that an agent is always in a position to know whether she is performing that action under the relevant guise. Nonetheless, there are actions that are basic at a given context; when I talk about ‘basic actions’ I mean the actions that are basic at the particular context under discussion.
18. The relationship between obligations and constraints is a complicated matter. Obligations are naturally expressed as oughts (“you ought to help people in need”), while constraints are naturally expressed as permissions (“you’re permitted to lie only if it is necessary to avoid serious harm”). However, constraints can also be expressed as oughts (“you ought not lie unless . . .”), raising the question of why obligations and constraints are different. To vindicate the distinction, we can note that *helping people in need* is a unified action, while *not lying* isn’t; the former is an action while the latter is merely an omission. And in general, one can conform to a constraint by *not doing anything*, while one can conform with an obligation only by doing *something*. Of course, we might wonder whether there really is an action/omission distinct. For the purposes of this paper, I will assume there is, and thus that there is a correspondingly meaningful distinction between obligations and constraints.
19. I am assuming here that a competent agent who knows which basic action she ought to perform in order to conform with an obligatory norm can thereby come

- by inference to know of all the other available actions that she is obligated not to perform them.
20. For the sake of simplicity, I speak in terms of knowledge, which conveniently brings together truth and justification into a single mental state. But one can say everything I want to say in factorised terms—i.e. if Anti-Cartesianism is true, it follows that one can either (a) have a justified but false belief that a given condition *C* obtains (or doesn't obtain), or (b) not be in a position to have a true and justified belief about whether *C* obtains. A *lucid* norm in these terms is a norm such that a competent agent who has a justified true belief about what the norm says is thereby in a position to have a justified true belief of every basic action available to her whether it would be in conformity with the norm.
 21. See e.g. Moore (1903/1993, 198–9) and (2005, 99).
 22. See e.g. Williamson (2000, 2007); Sutton (2007).
 23. Again, I'm only talking about those subjectivists and internalists who insist that their favoured norms *are* lucid. For further discussion of subjectivists and internalists who do not insist on the importance of lucidity, see fns. 9 and 55.
 24. I take the useful term 'hypological' from Zimmerman (2002, 554). It comes from the Greek *hypologos*, meaning 'to hold accountable or liable'. Thanks to Selim Berker for suggesting this usage.
 25. One finds this desideratum invoked and discussed across normative philosophy, especially in response to consequentialist and virtue theories in ethics and externalist theories in epistemology. For a few examples see Williams (1981); Nagel (1989); Bennett (1998); Shapiro (1998); Hursthouse (1999); Cruz and Pollock (1999); Farrelly (2007); Thomson (2008); Bales (1971, 256–65); Darwall (1983, 30–31); Hudson (1989); Gibbard (1990, 43); Jackson (1991, 461–82); Korsgaard (1996, 8); Smart (1961/1973, 44, 46); Stocker (1990, 10); Timmons (2002, 3).
 26. This is not a line of reasoning that Pryor himself endorses.
 27. Some might find this series of quotations unnecessary; a footnote with the citations would have been better. I agree. But I put these quotations in because of the number of philosophers who have told me that subjectivists and internalists never make these sorts of claims. They do.
 28. To repeat, in the concealed question sense, i.e. knowing what the norm says.
 29. An anti-intellectualist might reject the action-guiding requirement on the grounds that this presupposes that what's involved in such action is the active, cognitive application of general rules to particular instances.
 30. cf. Raz on the "explanatory-normative nexus" (2011, 27).
 31. See e.g. Mill (1861/2001, ch. 2); Moore (1903); Sidgwick (1907, 413, 489–90); Smart (1961/1973, § 7); Bales (1971); Hare (1981, ch. 4); Parfit (1984, 24–9, 31–43); Railton (1984, 140–6, 152–3); Pettit and Brennan (1986); Hooker (2000).
 32. Analogously, Conee and Feldman (2011, 310–314) argue that their evidentialist theory of epistemic justification is not meant to furnish an evaluation of agents, since their theory implies that agents can act blameworthily while being justified. Instead, Conee and Feldman claim only to be formulating conditions under which an agent's beliefs at a time fit the available evidence in the way required for the achievement of justification. In this way their evidentialist norm is merely a 'criterion of justification', not a guide to virtuous action.

33. See Bales (1971) for an account and defence of this distinction between criterion of right and decision procedure.
34. This isn't to say that the distinction cannot be meaningfully drawn at all. Since some norms are lucid at more (and perhaps more important) contexts than others, the distinction between a criterion of right and a decision theory might turn out to be a distinction of degree. If, as some argue (cf. n 15), the objectivist consequentialist norm is lucid at no contexts—if it is an *opaque* norm—then the distinction between a criterion of right and a decision procedure could be drawn (in a principled fashion) along the opaque/non-opaque line.
35. Cf. Wedgwood 2002, 365–6.
36. BonJour (1980) describes two cases involving clairvoyants. In the first case, clairvoyant Samantha (who has no reason for or against believing she is clairvoyant) forms a reliable, true belief through her clairvoyance that the President is in New York, despite having overwhelming evidence that he is in fact in Washington, D.C. In the second case, clairvoyant Casper forms a reliable, true belief through his clairvoyance that the President is in New York, despite having overwhelming evidence that he is not clairvoyant. Both cases are supposed to be counterexamples to reliabilism.
37. Thanks to Ram Neta for focussing my attention on this point.
38. Though Ross, following Prichard, worries that even the subjectivist norm might be non-transparent because of worries to do with borderline cases.
39. For other sympathetic epistemologists, see e.g. Ginet (1975, ch. 3); Chisholm (1977); BonJour (1985, ch. 1); Cohen (1984); and Audi (2001). For other sympathetic ethicists, see e.g. Prichard 2002.
40. For a similar diagnosis of the motivations for internalism, see Alston (1989, essay 8); Plantinga (1993a, ch. 1); Zagzebski (1996, 39); Goldman (1999).
41. On the link between this kind of ignorance and blameworthiness, see Rosen (2004), Smith (1983), Zimmerman (1997). For a dissenting view see Harman (2011).
42. By 'normative ignorance' I mean ignorance of the norms, not the kind of ignorance that is at issue in this paper—i.e. ignorance of the non-normative but normatively relevant facts.
43. I use the term 'normatively knowledgeable' to mean knowledgeable of the particular norm at hand—not knowledgeable of all the normative truths (which might, after all, be thought to include the truth about which basic action would be in conformity with the norm).
44. The issue of which theorists endorse the Alignment View is complicated by the fact that, in both the ethical and epistemic realms, there are various different sets of deontic fact which one might want to align with the hypological facts. Thus an externalist, for example, might be happy to grant that *justification* or *rationality* tracks blamelessness, but then deny that this is the deontic notion crucial for knowledge (e.g. Plantinga 1993a, 1993b). Similarly, many ethicists distinguish between two (or more) deontic notions—objective vs. subjective oughts, reasons, obligations, rightness, and so on—only one of which they take to track blamelessness (Ross 1939, 161; Broad 1946, 110–11; Russell 1970, 14; Brandt 1959, 366; Smith 2010, 65; Feldman 1986, 46; Parfit 2011, ch. 7; Zimmerman 2008, ch. 4; Raz 2011, ch. 6). I am glossing over this issue by acting as if there is only one

- set of deontic facts, i.e. conformity with the norms. But my elision is of no great importance; my claim is that there is *no* set of deontic facts that can have the tie to hypological facts promised by the Alignment View. Attempts to proliferate normative concepts in order to inhabit a half-way house between subjectivism and objectivism (or externalism and internalism) thus fails.
45. In this section, when I speak of an agent “knowing” a norm, I mean that the agent knows what the norm says, but also that the agent knows the norm to be true. Of course, evaluating counterfactuals that involve agents knowing false norms will involve evaluating counterfactuals with false antecedents in a way that does not make them come out vacuously true.
 46. For the *loci classici*, see Williams (1976) and Nagel (1979).
 47. While we might doubt that *any* of our actions are voluntary because of qualms about free will, this won’t help us explain what Aristotle is getting at when he picks out cases of coercion and ignorance as special cases of involuntariness. Nagel (1979) brings together the kind of moral luck at issue in the Fred and George cases with other kinds of luck, including the bad luck of having been causally determined to act in a particular way. Nagel writes that when we take into consideration all the various ways in which our actions are affected by forces outside our control, “[t]he area of genuine agency, and therefore of legitimate moral judgment, seems to shrink under this scrutiny to an extensionless point” (1979, 66). I read Nagel’s argument as a *reductio* against the idea that morality is immune from the machinations of luck (cf. Gardner (2001), 127–28.) But I think Nagel’s argument overlooks the fact that many normative theorists accept the existence of the kinds of luck on which he focuses (circumstantial, constitutive, and causal), but nonetheless maintain that morality is immune from luck in one significant way—i.e. that, on a suitably subjectivised conception of morality, agents control (in some important sense) the normative status of their actions. Nagel offers us no systematic reason to think otherwise, except to point (like Williams) to our actual practices of allowing luck to infiltrate our moral judgments. My purpose in this section is to show that Anti-Cartesianism implies that the project of subjectivising normativity in order to exclude this type of luck is also doomed.
 48. One might worry that *permissive* norms constitute counterexamples to this generalisation. While conformity with obligatory norms requires one to perform some action, conformity with permissive norms can be trivially secured by doing nothing (see n. 18). Suppose Arthur is permitted to eat the cake iff the cake was baked for him. So long as Arthur doesn’t eat the cake, he is sure not to violate the norm; in this sense, it is always within Arthur’s control whether he conforms to the norm. Thus our performance vis-à-vis *permissive* norms might not seem to be susceptible to bad performance luck after all. If so, then what I have to say only applies to obligatory norms. But I think this is too quick. First, the brain-in-a-vat is clearly a victim of bad performance luck vis-à-vis the reliabilist norm, though it is a permissive norm. The BIV isn’t obligated to form a belief about whether he has hands; if he had refrained from forming a belief, he would have certainly conformed to the permissive norm. And yet, we intuitively think he is a victim of bad performance luck. The vital issue here might be that the BIV’s evidence is *misleading*; thus he is in a position to justifiably (if falsely) believe

that he is forming a belief via a reliable mechanism. Things might be different for a permissive norm for which this is not a live possibility—i.e. a norm whose triggering condition is such that one can never have misleading (but only insufficient) evidence about whether it obtains. It's a serious question whether there are any plausible norms that meet this criterion. We should also query the thought that permissive norms can be conformed to simply by "doing nothing". After all, in both the epistemic and practical spheres, doing *nothing at all* is hardly an option: I always already find myself doing or believing *something*. In any case, I will leave these complications aside.

49. One can of course imagine cases in which agents end up *conforming* to the relevant norms through *good* performance luck. Suppose that Fred, in a miserly spirit, gave his donation to Charity A, while justifiably believing that Charity B would have been the more effective option. Or suppose that George abstained, out of cruelty, from giving his dog the medicine he justifiably but falsely believed to be life-saving. In these cases, Fred and George conform to the relevant norms, but only through good performance luck. And these are both cases in which the agents fail to be in a position to know whether the relevant norm's triggering condition obtains. So we might think that Anti-Cartesianism implies that there are no norms that are immune from *good* performance luck, either. I think this is true, but not for the reasons that we might suppose. If there were lucid norms, they would be immune from bad performance luck. But such norms wouldn't be immune from good performance luck. Suppose that Fred is always in a position to know which of two charities is most effective. Thus, he is in a position to know that he ought to donate to Charity A. But he can nonetheless conform to the norm through good performance luck simply by not investigating the efficacy of the charities. Or Fred might know that he ought to donate to Charity A, but do it for the wrong reasons, e.g. because he relishes the idea of depriving Charity B. So the lucidity of a norm at best immunises it from bad performance luck, not good. So while it is true that, if there are no lucid norms, then all norms can be violated through bad luck and conformed to by good luck, it is only on the first implication that I will focus.
50. On the supposition, that is, that normative nihilism—the view that there are no genuine norms—is false.
51. The claim that such ignorance excuses normative violation might seem to fall to a simple counterexample: suppose that I promise to pick you up at the airport and, forgetting I made such a promise, fail to pick you up. Intuitively I am blameworthy for failing you, although I was in no position to know (once the promise was forgotten) that I should have. I take it that a good response here is that if I was genuinely *non-culpable* for forgetting—say I got knocked on the head, rather than being simply negligent—than I am indeed excused. But in any case I am going simply to assume that such ignorance excuses. (This case was given to me by Stewart Cohen; for similar cases see Heintz 1975, p. 456; Sher 2009, p. 25; and Raz, 2011, 126.)
52. Thus Hume wrote: "'Tis evident, that when we praise any actions, we regard only the motives that produced them, and consider the actions as signs or indications of certain principles of mind and temper. The external performance has no merit. We must look within to find the moral quality. This we cannot do directly; and therefore fix our attention on actions, as on external signs. But these actions are

- still considered as signs; and the ultimate object of our praise and approbation is the motive, that produced them” (1739/1978 III.II.1).
53. Isn't there an obvious family of exceptions? What about the norm: $\langle \Phi \text{ iff you would be blameworthy for not } \phi\text{-ing} \rangle$? Call this norm N. If one violates N, one is *ipso facto* blameworthy. But one isn't blameworthy for *violating* N. One is blameworthy for violating some *other* norm, N*, tacitly referred to by N. If there are no lucid norms, then there is no norm such that whenever one violates it one is always blameworthy *for violating that norm*. One might of course be blameworthy in virtue of violating some *other* norm.
 54. One might again worry (cf. n. 29) that this assumes a background intellectualism about the normative, viz. the view that practical and epistemic action consists in consciously applying general rules to specific cases. If one is an anti-intellectualist, one will dislike the idea that blamelessness vis-à-vis normative violation consists in 'trying one's best' to conform to the norms, as I suggested in the Fred and George cases. One will think instead that blamelessness consists in something more like acting on dispositions to conform to norms, or acting as one so disposed would act. Does my claim here presuppose that such a view is mistaken? Plausibly not. Suppose that blamelessness vis-à-vis a norm N consists in acting on a disposition to conform to N. If Anti-Lucidity is true, then plausibly any norm N is such that one can have a disposition to conform to it that is *masked* by non-culpable ignorance of the relevant non-normative facts. (For example, if one is disposed to be kind to all and only those deserving of kindness, that disposition can be unsuccessfully acted on in a case where one fails to be in a position to know whether someone is deserving of kindness.) Is it possible that there are norms such that one can be disposed to conform to them without its being possible for the relevant dispositions to be masked by non-culpable ignorance? For example, I might have a disposition to conform to the modus ponens norm $\langle \text{If you believe } p \text{ and you believe that } p \text{ implies } q, \text{ believe } q \rangle$ that operates completely outside my conscious cognition. I might be just set up such that my conformity with the norm is automatic. Still, one might think that when I, say, have the false but justified belief that the norm's triggering condition doesn't obtain—i.e. I falsely but justifiably believe that I don't believe p or I don't believe that p implies q —then this non-culpable ignorance might mask this disposition in a way that excuses my violation of the norm. Then we have a case of my blamelessly violating the norm despite being competent: a counterexample to the Alignment View. One might of course simply deny this possibility by insisting that any agent who violates the modus ponens norm is incompetent. Then the Alignment View will turn out to be vacuously true. For what it's worth, I find it hard to think that the person who allows their conscious beliefs to interact with their inferential dispositions in this way—someone who is trying to put his doxastic house in order—should be deemed incompetent. I also find it implausible that such a person is blameworthy.
 55. For a defence of subjectivism that does not cite the Alignment View, see Smith (2010) and Zimmerman (2014). Zimmerman rejects objectivism because of cases discussed by Regan (1980, 264–5), Jackson (1991, 462–3) and Parfit (2011, ch. 7). In these cases, agents aren't in a position to know which basic action would maximise the overall good, but they are intuitively blameless for performing an action that they are in a position to know will *not* maximise the good.

(And indeed, they would be intuitively blameworthy for pursuing one of the actions that might maximise the good.) The reason this counts against objectivism is that the objectivist cannot explain away the intuition that these agents do what they ought simply by saying that, because the agents tried their best to conform to the objectivist norm, they are blameless. For it is a feature of such cases that the agents *do not* try their best to conform to the objectivist norm. Thus subjectivists want to say: the best explanation for why these agents are blameless is that they did what they ought, viz. that objectivism is false. I am unsure what to say about these cases, but they are somewhat orthogonal to my project here. I have not offered an account of what it is to be blameless vis-à-vis some norm; I have only assumed that for every non-lucid norm, it is possible to violate it through performance luck, and thus blamelessly. There is a distinct question whether the violations in Regan-Jackson-Parfit cases are committed through performance luck, or are blameless for some other reason. I don't mean to be defending objectivism against such cases.

References

- Alston, W. (1989). *Epistemic Justification*. Ithaca: Cornell University Press.
- Aristotle. (1980). *The Nicomachean Ethics*. Ross, D. (transl.) Oxford: Oxford University Press.
- Audi, R. (2001). "An Internalist Theory of Normative Grounds". *Philosophical Topics*, 29: 19–46.
- Bales, R.E. (1971). "Act-utilitarianism: Account of Right-making Characteristics or Decision Making Procedure?" *American Philosophical Quarterly*, 8: 257–265.
- Bennett, J. (1998). *The Act Itself*. Oxford: Oxford University Press.
- Berker, S. (2008). "Luminosity Regained". *Philosophers' Imprint*, 8(2): 1–22.
- BonJour, L. (1980). "Externalist Theories of Empirical Knowledge". *Midwest Studies in Philosophy*, 5(1): 53–74.
- BonJour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Brandt, R.B. (1959). *Ethical Theory*. Englewood Cliffs: Prentice-Hall.
- Broad, C.D. (1946). "Some of the Main Problems of Ethics". *Philosophy* 21: 99–171.
- Broome, J. (1991). *Weighing Goods*. Oxford: Blackwell.
- Carruthers, P. (2011). *The Opacity of Mind: An Integrative Theory of Self-Knowledge*. Oxford: Oxford University Press.
- Chisholm, R.M. (1977). *Theory of Knowledge*, 2nd edn. Englewood Cliffs, NJ: Prentice-Hall.
- Coliva, A., (ed.) (2012). *Self and Self-Knowledge*. Oxford: Oxford University Press.
- Conee, E. and Feldman, R. (2001). "Internalism Defended". *American Philosophical Quarterly*, 38(1): 1–18.
- Conee, E. and Feldman, R. (2011). "Replies" in Dougherty, T. (ed.), *Evidentialism and its Discontents*. Oxford: Oxford University Press, 281–332.
- Cruz, J. and Pollock, J. (1999). *Contemporary Theories of Knowledge*, 2nd edition. Lanham, MD: Rowman & Littlefield.
- Darwall, S. (1983). *Impartial Reason*. Ithaca, NY: Cornell University Press.
- Dodds, E.R. (1966). "On Misunderstanding the 'Oedipus Rex'". *Greece & Rome, Second Series*, 13(1): 37–49.
- Ewing, A.C. (1947). *The Definition of Good*. New York: MacMillan.
- Farrelly, C. (2007). "Justice in Ideal Theory: A Refutation". *Political Studies*, 55(4): 844–64.
- Feldman, F. (1986). *Doing the Best We Can*. Dordrecht: Reidel.
- Gardner, J. (2001). *Relating to Responsibility*. Oxford: Oxford University Press.

- Gibbard, A. (1990). *Wise Choices, Apt Feelings*. Oxford: Oxford University Press.
- Ginet, C. (1975). *Knowledge, Perception, and Memory*. Dordrecht: D. Reidel Publishing.
- Goldman, A. (1999). "Internalism Exposed". *Journal of Philosophy*, 96: 271–93.
- Gopnik, A. (1983). "How We Can Know Our Minds: The Illusion of First Person Knowledge of Intentionality". *Brain and Behavioral Science*, 16: 1–14.
- Griffin, J. (1996). *Value Judgment*. Oxford: Clarendon Press.
- Gruzalski, B. (1981). "Foreseeable Consequence Utilitarianism". *Australasian Journal of Philosophy*, 59: 163–76.
- Hare, R.M. (1981). *Moral Thinking: Its Levels, Method and Point*. Oxford: Clarendon Press.
- Harman, E. (2011). "Does Moral Ignorance Exculpate?" *Ratio*, 24(4): 443–468.
- Heintz, L.L. (1975). "Excuses and 'Ought' Implies 'Can'". *Canadian Journal of Philosophy*, 5(3): 449–62.
- Hooker, B. (2000). *Ideal Code, Real World: a Rule-consequentialist Theory of Morality*. Oxford: Oxford University Press.
- Howard-Snyder, F. (1997). "Rejection of Objective Consequentialism". *Utilitas*, 9: 241–48.
- Hudson, J.L. (1989). "Subjectivization in Ethics". *American Philosophical Quarterly*, 26: 221–29.
- Hume, D. (1739/1978). *A Treatise of Human Nature*. Selby-Bigge, L. A. (ed.), 2nd ed. revised by Nidditch, P. H. Oxford: Clarendon Press.
- Hursthouse, R. (1999). *On Virtue Ethics*. Oxford: Oxford University Press.
- Jackson, F. (1991). "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection". *Ethics*, 101: 461–82.
- Kagan, S. (1998). *Normative Ethics*. Boulder: Westview Press.
- Korsgaard, C. (1996). *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Lenman, J. (2000). "Consequentialism and Cluelessness". *Philosophy and Public Affairs*, 29: 342–70.
- Mason, E. (2003). "Consequentialism and the 'Ought Implies Can't' Principle". *American Philosophical Quarterly*, 40(4): 319–31.
- Mill, J.S. (1861/2001). *Utilitarianism*, 2nd edition. Sher, G. (ed.) Indianapolis: Hackett.
- Moore, G.E. (1903/1993). *Principia Ethica*. Baldwin, T. (ed.) Cambridge: Cambridge University Press.
- Moore, G.E. (2005). *Ethics*. Shaw, W. H. (ed.) Oxford: Clarendon Press.
- Nagel, T. (1979). *Mortal Questions*. New York: Cambridge University Press.
- Nagel, T. (1989). *The View from Nowhere*. Oxford: Oxford University Press.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.
- Parfit, D. (2011). *On What Matters*. Oxford: Oxford University Press.
- Pettit, P. (2003). "Consequentialism," in Darwall, S. (ed.), *Consequentialism*. Malden, MA: Blackwell Publishing, 95–107.
- Pettit, P. and Brennan, G. (1986). "Restrictive Consequentialism". *Australasian Journal of Philosophy*, 64: 438–55.
- Plantinga, A. (1993a). *Warrant: The Current Debate*. New York: Oxford University Press.
- Plantinga, A. (1993b). *Warrant and Proper Function*. New York: Oxford University Press.
- Prichard, H.A. (1932/2002). "Duty and Ignorance of Fact," in Ross, W.D. (ed.), *Moral Obligation*. Oxford: Clarendon Press, 18–39.
- Pryor, J. (2001). "Highlights of Recent Epistemology". *British Journal of Philosophy of Science*, 52: 95–124.
- Railton, P. (1984). "Alienation, Consequentialism, and the Demands of Morality". *Philosophy and Public Affairs*, 13: 134–71.
- Ramachandran, M. (2009). "Anti-Luminosity: Four Unsuccessful Strategies". *Australasian Journal of Philosophy*, 87(4): 659–73.
- Raz, J. (2011). *From Normativity to Responsibility*. Oxford: Clarendon Press.
- Regan, D. (1980). *Utilitarianism and Cooperation*. Oxford University Press.

- Rosen, G. (2004). "Skepticism About Moral Responsibility". *Philosophical Perspectives*, 18: 295–313.
- Ross, W.D. (1939). *Foundations of Ethics*. Oxford: Oxford University Press.
- Russell, B. (1970). "The Elements of Ethics". In Sellars, W., and Hospers, J. (eds.) *Readings in Ethical Theory*, 2nd edition. Englewood Cliffs: Prentice-Hall, 3–28.
- Scanlon, T.M. (2001). "Thomson on Self-Defense", in Byrne, A., Stalnaker, R. and Wedgwood, R. (eds.), *Fact and Value: Essays on Ethics and Metaphysics for Judith Jarvis Thomson*. Cambridge, MA: MIT Press, 199–215.
- Scanlon, T.M. (2008). *Moral Dimensions: Meaning, Permissibility, and Blame*. Cambridge, MA: Harvard University Press.
- Schwitzgebel, E. (2008). "The Unreliability of Naïve Introspection". *Philosophical Review*, 117: 245–273.
- Schwitzgebel, E. (2011). *Perplexities of Consciousness*. Cambridge, MA: MIT Press.
- Shapiro, S. (1998). "The Difference That Rules Make", in Bix, B. (ed.), *Analyzing Law: New Essays in Legal Theory*. Oxford: Clarendon Press.
- Sher, G. (2009). *Who Knew? Responsibility Without Awareness*. Oxford: Oxford University Press.
- Schoenfield, M. (2015). "Internalism without Luminosity". *Philosophical Issues*, 25(1): 252–272.
- Sidgwick, H. (1907). *The Methods of Ethics*, 7th edn. London: Macmillan.
- Smart, J.J.C. (1961/1973). "Outline of a System of Utilitarian Ethics", in Smart, J.J.C. and Williams, B.A.O. (eds.), *Utilitarianism: For and Against*. Cambridge: Cambridge University Press, 3–74.
- Smith, H. (1983). "Culpable Ignorance". *Philosophical Review*, 92: 543–571.
- Smith, H. (1988). "Making Moral Decisions". *Noûs*, 22: 89–108.
- Smith, H. (2010). "Subjective Rightness". *Social Philosophy and Policy* 27: 64–110.
- Srinivasan, A. (2015). "Are We Luminous?", *Philosophy and Phenomenological Research*, 90(2): 294–319.
- Steup, M. (1999). "A Defense of Internalism", in Pojman, L. (ed.), *Theory of Knowledge: Classical and Contemporary Readings*. Belmont, CA: Wadsworth, 373–84.
- Stocker, M. (1990). *Plural and Conflicting Values*. Oxford: Clarendon Press.
- Suikkanen, J. (2008). "Williamson and (Moral) Realism". PEA Soup blog. Posted 27th March, 2008. URL: <http://peasoup.typepad.com/peasoup/2008/03/williamson-and.html#more>
- Sutton, J. (2007). *Without Justification*. Cambridge, MA: MIT Press.
- Thomson, J. J. (2008). *Normativity*. Chicago and La Salle, IL: Open Court Press.
- Timmons, M. (2002). *Moral Theory: An Introduction*. Lanham, MD: Rowman and Littlefield.
- Timmons, M. (2012). *Conduct and Character*, 6th edn. Wadsworth: Cengage Learning.
- Weatherson, B. (2004). "Luminous Margins". *Australasian Journal of Philosophy*, 82: 373–383.
- Wedgwood, R. (2002). "Internalism Explained". *Philosophy and Phenomenological Research*, 65(2): 349–69.
- Wedgwood, R. (2003). "Choosing Rationally and Choosing Correctly", in Stroud, S. and Topolet, C. (eds.), *Weakness of Will and Practical Irrationality*. Oxford: Oxford University Press, 201–30.
- Wedgwood, R. (2010). "There Are Unknowable Moral Truths". PEA Soup blog. Posted 10th August 2010. URL: <http://peasoup.typepad.com/peasoup/2010/08/there-are-unknowable-moral-truths.html>
- Williams, B. (1976). "Moral Luck". *Proceedings of the Aristotelian Society*, 50: 115–135.
- Williams, B. (1981). *Moral Luck*. Cambridge: Cambridge University Press.
- Williams, B. (1993). "Postscript", in Staman, D. (ed.), *Moral Luck*. Albany: State University of New York Press, 251–8.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.

- Williamson, T. (2007). "On Being Justified in One's Head", in Timmons, M., Greco, J. and Mele, A. (eds.), *Rationality and the Good: Critical Essays on the Ethics and Epistemology of Robert Audi*. Oxford: Oxford University Press.
- Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge University Press.
- Zimmerman, M. (1997). "Moral Responsibility and Ignorance". *Ethics*, 107: 410–428.
- Zimmerman, M. (2002). "Taking Luck Seriously". *Journal of Philosophy* 99: 553–76.
- Zimmerman, M. (2008). *Living with Uncertainty*. Cambridge: Cambridge University Press.
- Zimmerman, M. (2014). *Ignorance and Moral Obligation*. Oxford: Oxford University Press.