

Précis of *Knowing Better: Virtue, Deliberation, and Normative Ethics*

DANIEL STAR

Boston University

How is it possible for there to be genuine normative reasons that ordinary virtuous agents are able to rely on to determine what they should do, given that they are, generally speaking, ignorant of fundamental ethical principles of the kind that interest philosophers, and the reasons such principles purport to specify? It is naturally tempting to suppose that they are able to do this by responding to reasons that are of a non-fundamental or derivative kind. The fact that a friend is experiencing grief is a reason for you to console her, but we do not expect the basic principles of a normative ethical theory to refer to grief. The fact that an act would be a lie is a reason not to do it, but not even Kantian normative ethicists would view such a claim as fundamental; it might be said that all lies involve disrespecting the rational capacities of the person being lied to, but, even if this were true, it would not be facts concerning (actual or potential disrespecting of) rational capacities that we expect ordinary virtuous agents to be directly responding to when they avoid lying.

Not everyone will accept that there is a genuine philosophical problem here. On the one hand, some dogmatic ethicists, both consequentialist and non-consequentialist, will insist that the genuinely virtuous do know, or are committed to, the correct ethical theory (perhaps “implicitly”, which is an easy thing to say, but a difficult thing to make sense of). Others, especially those drawn to utilitarianism, will claim that we should reject the idea that the virtuous reliably act on or for the right reasons. Still others will endorse a particularism about ethical principles that holds on to the traditional thought that the virtuous reliably act on the right reasons, but gives up on the central project of modern normative ethics.

I recognize that some philosophers will wish to respond in one of these ways, and I do say some things in *Knowing Better* about why I think this would be a mistake (in Chapter 1, especially), but the main aim of the book is not to persuade people who wish to get off the boat this early on that they should not do so, but rather to provide a satisfactory response to what I hope many philosophers will take to be a genuine problem. A couple of convictions stand behind this project. The first is that ordinary virtue is often on a more secure footing than philosophical theorizing—it seems odd to me when moral philosophers purport to *believe* very general normative theoretical principles, as distinct from purporting to believe, say, that they should normally keep their promises.

Unwarranted dogmatism of this kind is inconsistent with the virtue of epistemic humility, and risks giving normative ethics a bad name (this is compatible with thinking that we are justified in making such theoretical claims when doing philosophy *as if* one believed them). The second is the conviction that normative ethics, of the kind that aspires to articulate and defend sophisticated principles, is a worthwhile enterprise, slowly but steadily making genuine progress—from this perspective, particularism about principles is best viewed as a position of last resort.

One thing is clear: taking the main problem of the book on its own terms, it is not enough to simply say that there are derivative, non-fundamental reasons that ordinary virtuous people directly respond to in a way that generally allows them to act as they should, for we can ask: how are the virtuous able to *reliably* respond to such reasons (as reasons), when they are ignorant of fundamental normative reasons and associated ethical principles? I contend that an account of normative reasons that I have previously defended with Stephen Kearns, *reasons as evidence*, is both illuminating with respect to understanding the metaphysical relationship between derivative and fundamental reasons, and uniquely well-placed to explain how we might reliably respond to derivative reasons when ignorant of fundamental reasons. Since, on this account, practical reasons are facts that are evidence that one ought to do one thing or another, many very mundane facts count as reasons, and since evidence is the golden standard for reliability from the deliberative perspective (when one doesn't already know the answer to a question one is asking), it is difficult to see how any other account of reasons could, on its own, deliver a better story about reliability. At least, it is difficult to see how this might be done without effectively giving up on solving the main problem of the book, and returning to the way of attempting to avoid it that would have us believe that it is of the essence of virtue to be directly responsive to fundamental reasons.

After providing a general outline of my solution to the main problem of the book in Chapter 1, I present, in Chapter 2, a new argument for *reasons as evidence* that is connected to the thought, just mentioned, that evidence is the golden standard for reliability from the deliberative perspective, in the absence of knowledge concerning what one ought to do. The first premise of this argument is that if there is a class of facts that can very generally play the role of being good guides for deliberation for agents that are ignorant of the fundamental reasons for action, then facts in this class also get to be reasons for action. I call this premise the *Razian Insight*, due to the fact it can be located in Joseph Raz's well-known work on practical authority. The second premise is that facts that are evidence that one ought to do some act or other constitute just such a general class of facts. I consider alternative candidates for such a class of facts and argue that none of these candidates (for example, facts concerning one's desires, and evaluative facts) are as well-suited to play the role in question. To the extent that other facts appear to be good candidates this, I contend, is because they converge with evidence concerning what one ought to do; where there is no such convergence it will always be more rational to follow evidence concerning what one ought to do.

In Chapter 3, I turn from reasons to virtue. The chapter begins by arguing against the idea that the virtuous always do the right act, all things considered (I won't rehearse the argument here, suffice to say I use a new thought experiment to back up other arguments in the literature; the thought experiment has the virtue, I think, of indicating how it is that virtuous acts do not even coincide with subjectively right acts, where rightness is taken to be relative to the evidence one possesses). I suggest that we should instead think of

the virtuous as doing their best to respond appropriately to normative reasons, and of individual virtues as consisting in stable dispositions to respond to reasons of various delineable types. I propose that individual virtues be analyzed in terms of reasons and non-normative properties, and discuss some examples. Once the account of reasons defended in the previous chapter is put together with the new account of virtue, the practical virtues can be understood to be mostly epistemic in nature (the will to follow through on our ethical judgments in intention and action escapes this reductive program). Finally, I consider a problem for this evidence-of-ought account of virtue, the solution to which requires positing that the virtuous possess a body of mundane (non-sophisticated) ethical knowledge. Fortunately, this is an independently attractive idea.

Chapter 4 is concerned with the role knowledge plays with respect to deliberation and action in general. I begin with the idea, associated with Timothy Williamson's *Knowledge and its Limits*, that the norm of practical deliberation and action is knowledge, and I argue that *reasons as evidence* provides a much needed supplement to that idea. The combination of these two views provides an attractive picture of when it is rationally appropriate for facts to figure in practical deliberation as pro tanto considerations, and when it is rationally appropriate for them to be the basis for action (the idea that knowledge is the norm of practical deliberation and action does not distinguish between and settle these two separate issues). Facing up to a challenge to this view, which has its origin in Hume's thought that one cannot deduce an 'ought' from an 'is', forces us to accept that appropriate practical deliberation requires not just non-ethical knowledge, but essentially ethical knowledge as well. So the solution to the main challenge raised in this chapter is the very same solution as that provided for the challenge raised regarding the theory of virtue defended in the previous chapter. Given the main aim of the book, it is important that we not take ordinary ethical knowledge to be the same kind of knowledge that we aim for as an ideal in normative ethics. A concern that the favored account of reasons is talking past accounts of reasons that focus primarily on fundamental reasons is also addressed, in a manner that aims to shed new light on the nature of fundamental reasons.

References

- Kearns, Stephen and Star, Daniel (2008). "Reasons: Explanations or Evidence?" *Ethics* 119 (1): 31–56.
- (2009). "Reason as Evidence." *Oxford Studies in Metaethics* 4: 215–42.
- (2011). "On Good Advice: A Reply to McNaughton and Rawling." *Analysis* 71 (3): 506–508.
- (2013). "Weighing Reasons." *Journal of Moral Philosophy* 10 (1): 70–86.
- (2015). "Weighing Explanations," in Iwao Hirose & Andrew Reisner (eds.), *Weighing and Reasoning: Themes from the Philosophy of John Broome*. Oxford: Oxford University Press, 232–51.
- Raz, Joseph (1986). *The Morality of Freedom*, Oxford: Oxford University Press.
- Star, Daniel (2015). *Knowing Better: Virtue, Deliberation, and Normative Ethics*, Oxford: Oxford University Press.
- Williamson, Timothy (2000). *Knowledge and its Limits*, Oxford: Oxford University Press.

Replies to Cuneo, Driver, and Littlejohn

DANIEL STAR

Boston University

I am very grateful to Terence Cuneo, Julia Driver, and Clayton Littlejohn for finding time in their no doubt busy schedules to engage with *Knowing Better*, and for each doing so in a charitable, yet critical fashion. They raise a good number of challenges, all deserving of careful consideration, but I must, of course, be somewhat selective in my replies.

Reply to Cuneo

Cuneo focuses on the subject of the nature of normative reasons. He contrasts the account of reasons I argue for in my book with reasons primitivism, and raises the possibility that the account I favor is not better placed than this alternative is when it comes to addressing the main puzzle of my book, which Cuneo helpfully calls the “practical-theoretical puzzle.” Furthermore, he raises the possibility that I haven’t even managed to demonstrate that my preferred account does solve the puzzle, given all that I have to say about the relationship between fundamental and derivative reasons.¹ Although these are serious concerns, Cuneo describes his goal as being one of raising challenges for the book’s project, rather than providing full-fledged objections. I will suggest that he does not actually provide good reasons for thinking the book fails in either of the ways he mentions, although, in both cases, he raises very interesting questions in his discussion—these are questions that deserve more serious consideration than I will be able to provide here, yet they are questions that, so far as I can see, are not of central importance to the main concerns of the book.

Stephen Kearns and I dubbed the account of reasons we favor, *reasons as evidence* (Kearns and Star 2008, 2009). The claim at the heart of this account of reasons is the following.

R: Necessarily, a fact *F* is a reason for an agent *A* to ϕ if and only if *F* is evidence that *A* ought to ϕ (13)

Rather than rehearse the arguments for R that Kearns and I earlier provided, or focus on responding to criticisms of R, I present a new argument for this central claim in Chapter

¹ I am reversing the order in which he presents his two main challenges, here and below.

2, and put it to use in order to present a solution to the main puzzle the book addresses (a puzzle that the coauthored papers do not focus on).² It is not a goal of the book to establish that it is the *only* possible solution to this puzzle. To have set myself this task would have been to send myself on a fool's errand. I do, however, contend that other philosophers have not yet adequately faced up to the problem, and I would be delighted if my book sparked efforts by philosophers who defend alternative accounts of reasons to face up to it. Furthermore, I present some initial reasons for thinking that they will have a difficult time of it (see, especially, 17–18). Cuneo's comments regarding the alternative of accepting reasons primitivism suggest to me that I should either have dedicated more space to this last task, or scaled back the rhetoric in a couple of places (however, I do admit in the book that I do not possess an argument from elimination that would demonstrate that other accounts of reasons cannot solve the puzzle [17]).

Both in the book and in the earlier papers, R is generally discussed as an informative necessary truth, and not as a claim about the essence of reasons. Kearns and I noted that most of the arguments we were able to come up with, as well as most of the objections we had encountered, concerned this biconditional claim, but pointed out that we could argue for a stronger identity claim, from R, by inference to the best explanation (2009, 219). In my book, I similarly say, when introducing R, that for most purposes the biconditional claim is all one need focus on, and that I also do happen to think that *reasons as evidence* might best be construed as an account of the real definition of a normative reason (14–15). Cuneo reads me as putting more emphasis on the second idea than I intended to, since it was not, for the most part, necessary, with respect to either the premises or conclusions of my arguments (an exception lies near the end of the book, in section 5 of Ch. 4, as Cuneo rightly notes). And, despite using “in virtue of” language on a couple of occasions, I nowhere meant to rely on a metaphysical notion of grounding; I meant merely that, if and when we interpret R as providing a real definition, we will be offering an explanation of what reasons are (just as one might say the scientists who discovered that water *is* H₂O provided us with an explanation of what water is).

In any case, Cuneo is concerned to point out that R, understood merely as a necessarily true biconditional, is compatible with reasons primitivism. This is fair. I do not wish to deny, that in terms of *prima facie* logical possibility, R and reasons primitivism could both be true. As a matter of fact, I think most actual reasons primitivists, such as Derek Parfit and T. M. Scanlon would deny R (at the very least, they don't think that whenever there is evidence that one ought to do an act, there is a reason to do that act), and if I were able to convince *them* to adopt R without giving up on reasons primitivism, I should consider that a victory! If *I* came to accept reasons primitivism, while continuing to hold on to R, would I take it that my book project was unnecessary? Not at all. The point of the book was to solve the practical-theoretical puzzle, and R plays an essential, central role in my solution to the puzzle. I therefore agree with the spirit of Cuneo's upbeat concluding sentence. But that sentence is in tension with the earlier sentences where he asserts that the challenge he is pressing is that R is no more than “a way to resolve the practical-theoretical puzzle”, and that “Here is a reason to think that [this] is not [true]” (a similar statement of the challenge appears in the introduction to his paper).

² In our first two coauthored papers (Kearns and Star 2008, 2009), we considered a number of criticisms, and in Kearns and Star 2011, 2013a, and 2013b we provided detailed responses to four of our critics, but we have not yet responded to all of the criticisms of our view that are in print.

I don't think he provides such a reason. R may be "purely extensional" (his words), but it may be informative and important, nonetheless.

All this being said, I also happen to think that I can argue against reasons primitivism, which Cuneo takes to be the position that there is nothing informative that can be said about why some facts are reasons (yet some facts are reasons).³ I present one argument for favoring R, as a definitional claim, over reasons primitivism, in a forthcoming paper (Star, Forthcoming). This argument concludes that reasons primitivism is essentially redundant. I won't here rehearse the part of the argument that focuses on the role of reasons in deliberation, but I will say this: once R is on the table—that is, the biconditional claim is taken to be true—and we are comparing the strengthened, real definition interpretation of R simply with reasons primitivism, I don't think that arguing for reasons primitivism would achieve anything interesting. Reasons primitivism is taken to be the view that nothing informative can be said about why some things are reasons, R now offers to provide an explanation of why some things are reasons, and metaphysical redundancy is generally considered a bad thing. Reasons primitivism is, from this perspective, an unnecessary fifth wheel. Or, to put things in a more conciliatory way, the dialectic should at least be admitted to be one where the onus is on the reasons primitivist at this point to provide arguments for adopting reasons primitivism alongside R. The redundancy at issue looks even worse when the defender of *reasons as evidence* offers to provide an explanation of the strength of reasons in terms of epistemic probability (Kearns and Star 2008, 2009, 2013b, and Ch. 4 of my book). The reasons primitivist agrees that reasons come with strengths, so the primitivist's commitments have an extra dimension to the one Cuneo discusses. If it is admitted that the strength of reasons can be accurately specified in terms of the strength of evidence, then the primitivist either ends up with an unattractive mixed picture (that a fact is a reason is primitive, but the strength of this primitive thing can be explained in different terms), or for every property of having a certain evidential strength there is a corresponding non-natural property with magnitude, where the properties are always instantiated together, and their values are always the same.

Here is an additional argument against reasons primitivism. The solution to the practical-theoretical puzzle involves recognizing a distinction between fundamental and derivative reasons, where fundamental reasons are those reasons picked out by the very general normative ethical principles we articulate and defend when doing normative ethics. Reasons primitivists must surely accept a distinction of this kind. I think they should do this because they should allow that ordinary virtuous agents are sometimes ignorant of fundamental reasons, yet are still able to respond to reasons. But even reasons primitivists who reject this last idea generally accept that ordinary instrumental reasoning involves chains of instrumental reasons: I have a reason to go to the shop because the fridge is empty and I have a reason eat dinner tonight (and I have a reason to eat dinner tonight because. . .). Assuming they admit just this much, reasons primitivists actually don't think that *all* reasons are incapable of being explained. Fundamental reasons are

³ The reasons primitivist may also wish to rely on a substantive non-naturalist metaethics, but this is not something discussed in Cuneo's paper. If one's main motivation were to find a place for positive non-naturalist or robust realist metaethical claims, then I think one could aim to combine R with a non-naturalist account of necessarily true, fundamental ethical principles or laws. Such principles would refer to fundamental duties or goods, but one can think that what explains why fundamental duties or goods are reasons is that they are *invariably* evidence that one ought to do acts of some type (see Ch. 4).

unexplainable, but derivative or instrumental reasons can be explained in terms of primitive reasons and contingent facts (such as that the fridge is empty), by use of a transmission principle of some kind.⁴ The essence of derivative or instrumental reasons is not primitive, even for the primitivist. In effect, this makes for a disjunctivist account of reasons. *Reasons as evidence* avoids this problem: certainly there are some reason facts that play a fundamental explanatory role in relation to contingent ought facts, but these fundamental reason facts are as much evidence that one ought to act in a particular way as the derivative reasons are. This gives us some reason to think that *reasons as evidence*, being a more unified account of normative reasons in general, is a better theory than the reason primitivists' ultimately disunified theory of reasons.

Cuneo also suggests, in passing, that explanation-based account of reasons provide an alternative type of account that might fair just as well as *reasons as evidence* in solving the practical-theoretical problem. Here he overlooks that I argue against such accounts (17–18, as well as Kearns and Star 2008, 2015). Explanans appear only after explananda when we look at how explanations are arrived at in practice, whereas judgments concerning one's reasons typically occur before judgments concerning what one ought to do when we look at how reasons and oughts are related to each other in deliberation, yet this is exactly the opposite of what one would predict if reasons were themselves parts of explanations of ought facts. In addition, the argument I just made regarding lack of unity with respect to reasons primitivism when it comes to instrumental reasons will apply *mutatis mutandis* to explanation-based accounts of reasons. Finally, simply adding a transmission principle to a basic theory of this kind to generate derivative reasons won't be enough to solve the practical-theoretical puzzle, since it won't tell us how people are able to reliably respond to derivative reasons, as reasons, in ignorance of the fundamental reasons that ground the derivative reasons, since they cannot do this by making use of the transmission principle, as this would require them to form inferences on the basis of knowledge of fundamental reasons (a point I also make in the book).

I turn now to Cuneo's other main challenge. I acknowledge in my book that pretty much any theorist of reasons can assert that there are both derivative and fundamental reasons, but claim that it is not enough to assert this, for "an adequate explanation of how these two types of reasons are related" must also be provided (12). This, I claim, is a desideratum for solving the main puzzle of the book. Cuneo thinks that any effort to meet this desideratum, as I construe it, must involve specifying "how the ought-facts and fundamental reasons are related such that by grasping the former, the virtuous agents reliably track the latter." However, I never said the desideratum in question requires any

⁴ Here is a candidate, in the same spirit as R (even if its scope is meant to be more limited), from Niko Kolodny (Forthcoming): "*General Transmission*: If there is reason for one to *E*, and there is positive probability, conditional on one's *M*-ing, that one's *M*-ing, or some part of one's *M*-ing, helps to bring it about that one *E*'s nonsuperfluously, then that is a reason for one to *M*, whose strength depends on the reason for one to *E* and the probability, so long as the reason for one to *E* is not explained by an application of *General Transmission* to reason for one to achieve some distinct *E*". Of course, any particular specification of a transmission principle like this one will be controversial, but I provide this example here in order to indicate that the appeal of relying on epistemic probability (Kolodny makes it clear in his paper that he has this type of probability in mind) to explain the strength of derivative reasons is not limited to defenders of *reasons as evidence*. It is also notable that Stephen Finlay (2014, 92), in the midst of an excellent defense of an explanation-based account of reasons, contends that we might explain the strength of *all* reasons in terms of probability (without indicating that this brings his view much closer to *reasons as evidence*, even though he criticizes this view in the same chapter of his book).

such explanation, and I don't think it does. Cuneo asserts that I think that "the virtuous track the fundamental reasons by knowing the derivative ones", but I don't think this is true, and I don't say that the virtuous are, generally speaking, in the fundamental reason-tracking business (I instead say they are, in general, in the derivative reason-tracking business).⁵ When we think about virtuous agents who are not particularly interested in philosophical reflection, they may have no commitment to there being reasons more fundamental than those they generally take seriously (although, unlike some philosophers who are interested in the "one thought too many" problem, I don't think ordinary virtuous agents are positively committed to more fundamental reasons *not* existing).

I should admit that there is one place in my book that may have given Cuneo the impression that I think that the virtuous track fundamental reasons, and that is in the paragraph where I write, "I can respond to the fact that an act would cause someone pain *indirectly* by responding *directly* to the fact that the act would be a lie, without for a moment taking into account the fact that it would cause pain." (22–23). I didn't mean this to suggest that virtuous agents need to take themselves to be responding, even indirectly, to fundamental reasons, nor that this is required to meet the just mentioned desideratum. It's worth noting that in the next paragraph of the book I point out that a limitation of putting my point in terms of indirect responses is that in cases of misleading evidence that one ought to do an act, there is no fundamental reason present to be responding to.⁶ Still, it might have been best to altogether avoid putting things this way in the first place.

Cuneo is quite right to say that R concerns a relation between reasons and (putative) ought-facts, and does not mention the relation between derivative and fundamental reasons. He points this out, and then says, "The practical-theoretical puzzle, however, concerns a relation not between reasons and ought but between derivative and fundamental reasons." It isn't at all clear to me why he thinks this is the concern, or even a concern, of the puzzle of my book, since I never described the puzzle in the way this sentence suggests I did. I'm not even sure what the puzzle, as Cuneo construes it when he says this, is meant to be. In saying this, I don't wish to deny that there are interesting questions in the vicinity. For instance: how can philosophers manage to work out what our fundamental reasons are if we start off generally responding to derivative reasons? This is not the problem of my book, even if I do suggest, in Ch. 1, that the method one will wish to focus on in answering this question is simply the widely practiced philosophical method of seeking reflective equilibrium between particular moral judgments and general principles, but where this is understood to involve the use of pre-theoretical moral

⁵ In passing, I wish to register a minor disagreement with one way that Cuneo describes fundamental reasons. He writes, "Utilitarians maintain that your neighbor's being doubled-over in pain is a [fundamental] reason for you to call for help when and because calling for help would maximize utility," but I take it that utilitarians will not wish to explain why particular facts provide us with pro tanto reasons by an appeal to maximizing utility (although they will rely on that notion when providing a criterion to determine which acts are the acts you ought to do), but will instead do this by referring to facts concerning pleasure and pain. This is because a reason to relieve pain still needs to be a reason to relieve pain when it is outweighed, where it being outweighed is explained by the fact that acting on it will not maximize utility.

⁶ Although I think practical reasons, and not just epistemic reasons, can sometimes be misleading, I am concerned in the book to argue that the practical reasons that virtuous agents respond to are not systematically or generally misleading (this is on the assumption that there can be virtuous agents; if normative nihilism is true, it follows from R that all reasons are misleading, but it should not be surprising to learn that virtue is a sham if nihilism is true).

knowledge, and not just “intuitions” (see, especially, 10–12, 19). I also say more about the relationship between derivative and fundamental reasons in the discussion of particularism in section 6 of Ch. 1, as well as in the closely related section 5 of Ch. 4.

Perhaps Cuneo’s concern is that it is not clear why we need fundamental reasons at all if responding to them is not something that the virtuous are generally in the business of doing. If this is the concern, I would repeat the following claims from the book: (1) it’s important for understanding the relation of normative ethical theory to ordinary agents that we take fundamental right-makers and wrong-makers to be reason providing (a solution to the puzzle must maintain that fundamental reasons really are reasons for people to do various acts, even if they are not responding to them); (2) being virtuous is not the same as being maximally virtuous, and being maximally virtuous may require knowledge of the fundamental reasons (my central claims about virtue are merely claims about what being virtuous requires); (3) being virtuous doesn’t require one to have all the answers to difficult moral questions, and this is why we need normative ethics (I think we are pre-philosophically committed to thinking ordinary virtuous people possess epistemic humility in relation to the question of what they ought to do, and I argue that virtue comes apart from right action in Chapter 3); (4) even if, contrary to the last claim, being virtuous were enough to ensure that a person would always do the right thing, philosophers would still be interested in more fundamental explanations of ethical truths (we value philosophy for the explanations it offers, and not just for the practical benefits uncovering certain explanations might provide).

In contending that other theories of reasons do not adequately explain the relationship between fundamental and derivative reasons, but R does, and in claiming that this is a desideratum for solving the main puzzle of the book, I was thinking that other theorists need both a metaphysical story that allows them to assert that there are non-fundamental reasons *and* an epistemological story that explains how we can reliably respond to these derivative reasons in ignorance of fundamental reasons, *and* that something in these stories also needs to tell us what happens when derivative and fundamental reasons interact (to avoid a double counting problem for agents who are familiar with both, for a start). A general recipe for the first is easy enough to provide; as I said above, all that may be needed is a transmission principle that explains when it is that there are derivative reasons (but this may involve introducing an unattractive lack of unity, as I noted). It is much more difficult to see how other accounts of reasons will answer the epistemological challenge. Finally, I’m not sure what other accounts of reasons will say about problems that can arise when we think about the possibility of double-counting reasons. It would be a mistake to add to the weight of a derivative reason to get lunch the weight of the fundamental wellbeing-related reason that would provide an ultimate justification for getting lunch, and thereby end up concluding that one has twice as much reason to get lunch as one would have thought one had before one landed on a theory of wellbeing. I contend that defenders of *reasons as evidence* can rely on the distinction between independent and dependent items of evidence that epistemologists are familiar with to solve this problem, given that practical reasons are (or are coextensive with) evidence concerning what one ought to do (33–34).

Reply to Driver

Driver’s work on virtue is well known, and I used one of the central ideas of her *Uneasy Virtue* (2001) as a foil in Chapter 3 of my book. It is therefore only natural that her

response to the book focuses on what I have to say there about virtue, although it is a nice feature of her response that she does not simply focus on differences between our views. She rightly points out that there is an interesting midway point between her own view and mine. And she carefully considers the question of whether my theory of virtue can really make sense of “virtues of ignorance”, especially modesty (it’s a strength of her own theory that it is well-placed to do so, assuming one thinks that modesty is a virtue and that it is correct to say that it requires ignorance). She also raises an interesting issue about higher-order evidence, and wonders whether my appeal to basic moral knowledge really helps in addressing the problem I confront in my book regarding virtue and misleading evidence.

Driver begins with a generally sympathetic and accurate discussion of the problem that motivates the book, where she says that Nomy Arpaly has “also” argued for a view that makes the moral worth of an action depend on the right-making features of action, but does not require that the agent appreciate the right-making status of these features (Driver says “those right making features”). I think it is important for me to note that Arpaly and I don’t, in fact, agree about the position the virtuous person is in with respect to right-makers, and that instead of addressing the puzzle I take myself to be focusing on in my book, Arpaly (at least the Arpaly of Arpaly and Schroeder 2014) actually denies one of its key assumptions, in one of the ways I mention it is possible to do in the Précis.⁷ It’s true that both of us agree that virtuous or morally worthy acts need to be responses to genuine normative reasons that do not need to rely on an understanding of ethical theory, however if we use “right-makers” (a term of art) to refer to the facts that fundamentally explain why acts are right or wrong, Arpaly thinks virtuous agents do need to be responding to (or morally worthy acts need to be responses to) facts of this kind, albeit not necessarily under the guise of “right-maker”, whereas I think they do not to be responding to fundamental reasons at all. In her book on virtue, Arpaly (2003) is keen to emphasize that some virtuous agents may believe that the reason facts they are responding to are *not* right-makers, and her discussion of Huck Finn is now rightly much discussed because of her insightful interpretation of it along these lines.⁸

Arpaly and Schroeder (2014) are, like me, keen to understand how it is that virtue can be reliable, given that ordinary virtuous agents are ignorant of ethical theory. They claim that virtuous agents respond directly to fundamental normative reasons via intrinsic desires that concern the facts that are reasons *correctly conceptualized*, using the same concepts that the normative ethicist with the correct theory would use when specifying fundamental reasons. This is what guarantees the reliability of the responses of the virtuous, so far as Arpaly and Schroeder are concerned. Virtuous agents need not think they are responding to fundamental reasons, nor even construe them under a normative or evaluative guise.

⁷ Driver doesn’t refer to Arpaly and Schroeder (2014), but I am assuming that the fact that this is a more recent book than the earlier book she refers to makes it not unreasonable for me to take it to be representative of Arpaly’s views in this area. One review of my book also suggests I should consider Arpaly’s views (Bommarito, Forthcoming). I’m glad Driver has provided me with the opportunity to do so, even if the discussion here must be brief.

⁸ For my part, I would say that Finn, in the relevant part of the story, is responding appropriately to some non-misleading evidence concerning what he ought to do (the reasons that apply to him that he disowns but fortunately acts on), but has misleading higher-order evidence about this evidence, as well as misleading evidence concerning what he ought to do (misleading reasons).

For what it's worth, this seems to me to be an unattractive way to approach virtue, since it must either leave too little room for philosophical reflection to make progress beyond the ethical judgments ordinary agents start with, or make it too difficult for non-philosophers to act in a morally worthy way. Given the evidence we have concerning debates in contemporary normative ethics, with respect to possible candidates for fundamental reasons, I think too much in the way of conceptual sophistication is required of virtuous agents on this theory.

I would contend that the correct conceptualization required on this theory cannot be present in actual agents, for desires to utilize, in the absence of extended philosophical reflection, but even if it could be present absent such reflection, the mere possession of such correct conceptualization would not guarantee the kind of reliability needed for virtue. This is, I think, the most important point of disagreement between us. Suppose a Kantian theory that claims facts about respecting rational capacities are most basic is correct, and Paula understands this theory well enough to possess the concepts required for picking out fundamental reasons, and uses these concepts to reliably pick out rational capacities, but she follows the theory with her desires—responds to the relevant reason-facts in the way required by Arpaly and Schroeder—merely because she *plumps* for tracking facts using these concepts (she is not led to accept the relevant theory in a rationally justifiable way by considering arguments, nor via some insight in the form of a non-inferentially justified belief, or what have you). Paula might very easily have instead *plumped* for utilitarianism, let us suppose. It seems very odd to think that, absent any warrant or epistemic justification for where she ends up, Paula now has what is required for virtue, whereas if she had instead *plumped* for the utilitarian account of fundamental reasons she would utterly fail to be virtuous. And, if Sharmin *plumps* for relying on one set of concepts over the other without any awareness or consideration of ethical theories at all (when she is otherwise like Paula), this cannot put her in a *better* position than the position Paula is in. Our disagreement here concerns a manner in which *luck* might be thought to undermine virtue. Arpaly and Schroeder think that just happening to have the correct concepts, and appropriately employing them, will be enough, whereas I think it really matters *how* one has ended up repeatedly employing the correct concepts (is one in a good epistemic position when employing them?).⁹

In relation to the issue of luck, Driver's own view is more extreme than Arpaly and Schroeder's, since she contends that a character state will count as a virtue if it systematically (and not just with respect to the behavior of the particular agent who possesses it) brings about good effects in the world. Everyone might think greed is a vice, but it might still be a virtue. And, as I say in the book (noting that I am not the first to point out that this is an implication of her view), kindness would be a vice if, unbeknownst to any of us, an evil demon tortures people on another planet whenever we are kind to each other (but doesn't if we are not kind). This radically detaches virtue from the ordinary agent's perspective, removing any requirement, in the form of a necessary condition, that virtues be reasons responsive (although Driver rightly says some virtues may nonetheless be reasons responsive, as a matter of contingent fact), and very much makes the status of being virtuous a hostage to a type of luck that it is impossible to forecast.¹⁰

⁹ Thanks to Timothy Schroeder for confirming the above points of disagreement in correspondence.

¹⁰ This type of luck is quite different than the kind of luck that is involved in receiving, or not receiving, a decent moral education, which we should all agree virtue can be a hostage to.

Despite putting things in these strong terms, I should say that this is also one of the most important accounts of virtue that have been defended in recent times, and Driver deserves credit for her defense of it. As she points out in her essay, part of her motivation for doing so when she wrote her book was the prevalence of the kind of intellectualist Aristotelian approach to virtue that we both find unattractive. She was not then considering the possibility that one might tie virtues to reasons and evidence, but not to fundamental reasons. As Driver indicates, I'm very much concerned to defend an account of virtue that isn't intellectualistic in the way she was concerned the Aristotelian account was, hence am not the kind of target she had in mind when she wrote her book.

Regarding Driver's point that a "mixed" account of virtue can claim that the production of good effects can be a necessary condition for a trait to count as a virtue, without being a sufficient condition for it to do so (reasons responsiveness will also provide a necessary condition, on one version of this view), I concede that she is right to point to the availability of this option, and that it is *prima facie* somewhat attractive. I didn't say anything about it in my book, although it deserves serious consideration. However, it was reasonable for me not to view this account as a target, because I wanted to specifically argue against theories that tie virtue to right action or virtue *rather than* reasons responsiveness, in order to clear space for my presentation of a reasons responsiveness view. Furthermore, my own view may be consistent with the mixed view, if one thinks epistemic norms are ultimately to be given a practical consequentialist (perhaps rule consequentialist) grounding. There exist somewhat persuasive arguments against this view of epistemic norms, but the question of what, if anything, does ground epistemic norms is clearly a live and interesting philosophical question.

When Driver later asks whether my reliance on moral knowledge to avoid the problem of misleading evidence really works, she might, in effect, be overlooking this consequentialist option with respect to epistemic norms (although I suspect she might have some sympathy for the option itself). Her concern is that it may be possible that certain putative virtues, despite the fact that they might well appear to be undergirded by commitments that we are tempted to view as instances of moral knowledge, could turn out to be positively harmful in the very long term. I'm now suggesting that if the consequentialist strategy with respect to epistemic norms works at all (and isn't terribly revisionary with respect to what we presently take those norms to be, which would be a huge mark against such a strategy, most of us would think), then even a virtue that might seem to lead to bad consequences when considered in isolation, really shouldn't be considered in isolation from all the other practical virtues. (If we put to one side this option, I would respond to her later objection by biting the bullet, as she suggests I might, in line with my adoption of the form of perspectivalism I discuss when responding to Littlejohn below.)

I am certainly not driven in anything I have written by opposition to consequentialism *per se*, as I am agnostic as to whether a consequentialist or a non-consequentialist theory of right action will ultimately turn out to be correct (and I recommend this agnosticism to everyone, although I don't think it entails that one shouldn't pick a side for the purposes of testing out philosophical arguments and developing sophisticated theories). I discussed Driver's (early) account of modesty, and argued, as others also have, that we should reject it, but my only motivation for doing this was that, if one took it that only her account of virtue could handle the virtues she calls "virtues of ignorance" (modesty being the most important of these), there would appear to be just one

significant advantage to adopting her account of virtue instead of the one I favor. I say just one advantage, because of the counterintuitive consequences of her view that I mention above, and the fact that my account of virtue is compatible with consequentialism about right action, yet saves more of our pre-philosophical judgments concerning virtue (an alternative motive for adopting her account that I mention in the book, of simply aiming to be a *global* consequentialist, would sacrifice the way in which my account can redeem ordinary virtues for the sake of unnecessary theoretical tidiness).

In her discussion of modesty and other virtues that may seem to involve ignorance or evidence insensitivity (tolerance may be an example of the latter kind, she suggests, and this is certainly an interesting idea), Driver mentions a way in which she weakened her own view of modesty in her later work, moving away from the idea that ignorance is necessary for this virtue to the weaker idea that some important cases of modesty involve ignorance. I regret that I overlooked this feature of her later work. This weakening of her view does, I agree, make our views about modesty much more similar, but it would also appear to significantly weaken the basis for using the fact that her view can explain virtues such as modesty as a reason to prefer her more general account of virtue over my own.

Evidence collecting is, however else we might also categorize it, a practical activity. If modesty is a virtue (and I'm less certain that it is than I am that courage and kindness are virtues), and if it fits the recipe for defining practical virtues that I provide in the book (the precise, complex details of which I must omit here), then it need not be considered a narrowly epistemic virtue concerned with how we treat evidence when we possess it. The need for "willful ignorance" Driver rightly claims is part of my account of modesty is, I claim, something to be achieved through acts that avoid putting one in touch with certain kinds of evidence about oneself, and one's standing in relation to others, rather than in ignoring evidence when it is staring one in the face. This is why I wrote, "Suppose the modest person stumbles across evidence that she is excellent at some activity—e.g. she is shown a newspaper that she generally takes to be reliable, bearing a large headline announcing that she is the best pianist in the world. Does this now make it impossible for her to be modest? Perhaps not..." (84). I also took issue with Driver's early view that the modest person must be somewhat lacking in self-worth, because, as she now agrees, immodesty has more to do with inappropriate comparisons (as well as, I would add, inappropriate inferences, e.g. since one is the best pianist, one is a better person than all the people one went to school with).

Driver's suggestion that an analysis of certain practical virtues may require reference to higher-order evidence is extremely interesting. I'm not sure we need to mention higher-order evidence in order to understand modesty, but perhaps we do. This will partly depend on how exactly we define higher-order evidence (does it include evidence concerning whether one ought to engage in the evidence seeking and avoidance patterns of activity the modest person engages in, or merely evidence concerning the reliability of one's various capacities for assessing evidence that one possesses?). In any case, it might well be that one needs to keep track of evidence concerning one's own particular evidence-gathering and avoiding tendencies in order to succeed at being modest. I'm not sure what to say about Driver's comments regarding the virtue of tolerance, but I take her general point that one can think sensitivity to evidence is usually very important, yet also think that certain virtues may work against it in a way that invites discussion of higher-order evidence. This wouldn't, I should admit, fit with the complex formulae

provided in the book when I discuss some key virtues—they don't refer to higher-order evidence—but I appreciate her suggestion that it might conform with the spirit of the general account, as that matters to me much more. Driver's comments here are grist for the mill.¹¹

Reply to Littlejohn

Littlejohn's discussion of the part of my book project that is concerned with evidence and knowledge is astute, and it contains challenging arguments that press me to say more about an issue that I was worried about when I was writing the book (this will be news to him, which perhaps just goes to show how astute his discussion is). In essence, the issue here is one of whether or not to adopt a perspectivalism about reasons that would have it that all normative reasons are dependent on the particular epistemic positions that agents find themselves in. Rather than making my mind up about this issue at an earlier time, I decided to complete the book in a way that would be non-committal in the relevant respects, choosing my words carefully with this decision in mind in a few places. Although I could see that I would need to explicitly address the issue at some point down the road, not doing so at the time seemed justifiable because I was trying to write a book that was ecumenical with respect to a wide range of positions in normative ethics, and committing to either perspectivalism or objectivism would seem to halve the space of possible normative ethical theorists that my solution to the main problem of the book might appeal to.¹² (Still, if perspectivalism turns out to be needed to solve the main problem of the book, or is anyway correct, objectivist ethical theorists should instead opt for perspectivalist versions of their theories.)

Before I engage with Littlejohn's main arguments, a point of clarification is in order. Littlejohn asserts that when a speaker says that Agnes has a reason to ϕ this means no more and no less than that there is a reason for Agnes to ϕ . Here we disagree. I don't wish to assert the contrary claim that ordinary speakers invariably refer to different things when using these two locutions. However, I do think that: (1) reasonologists of various stripes may find it helpful to distinguish between reasons that are possessed by agents and reasons that agents fail to possess (e.g. a Humean might think of possessed reasons as explained by one's present desires, and unpossessed reasons as instead constrained by the desires one would possess if one were better informed, more consistent, more imaginative, etc.); (2) sometimes we use the "is"/"has" distinction in ordinary speech in a way that lines up with this distinction between possibly unpossessed reasons and possessed

¹¹ In his review of *Knowing Better*, Alex Worsnip (2016) raises different concerns about higher-order evidence, in relation to *reasons as evidence*, and how this account of reasons analyzes reasons for belief, in particular. I cannot address these concerns here, but hope to on another occasion.

¹² I wasn't worried about one thing that people might worry about with respect to perspectivalism (especially when it goes by the name "subjectivism"), and that is that it might undermine moral objectivity. If we reflect on the well-known distinction between expected-value consequentialism (perspectivalist) and actual-value consequentialism (non-perspectivalist), we see that the first ethical theory is no less objective than the other, in the sense of "objective" that concerns us when considering moral skepticism. I appeal to the existence of basic, non-inferential ethical knowledge in my book, and the perspectivalist can understand such knowledge to be *a priori* in nature (admittedly, if one has the view that non-inferential moral knowledge is arrived at *a posteriori*, perspectivalism may conflict with moral objectivity, although the devil will be in the details). I will return to *a priori* knowledge below.

reasons;¹³ (3) sometimes the “is”/“has” distinction is also used in ordinary speech to point to a parallel distinction between possibly unpossessed and possessed evidence (e.g. “there may be evidence that Jack is the thief, but none of us have it!”); and (4) the fact that there is a parallel between these uses of “is” and “has”, in relation to reasons and evidence, provides some, admittedly defeasible, support for *reasons as evidence* (which, as Kearns and I defend it, would say that possessed reasons line up with possessed evidence, and unpossessed reasons line up with unpossessed evidence).

I find this distinction helpful. Some of Littlejohn’s claims need to be adjusted if one wishes to respect it, as he does not wish to. However, so far as I can see, none of his arguments hinge on this fact, so I don’t want to make too much of it. The most important thing to note is that when Littlejohn speaks of “possessed” evidence or uses the “have” locution, he means (or is open to meaning) evidence facts that an agent is in a position to know, whereas, when I use the locution, I mean evidence facts that an agent in fact knows (to be precise, the agent knows the facts that are evidence), since I think it is best to reserve talk of “unpossessed” evidence for facts that an agent is in a position to know but does not know.¹⁴ I think there are many good reasons to use this distinction, but a particularly salient one is that it is an ingredient of the solution to the main puzzle of my book: virtuous agents possess and respond to derivative practical reasons, but need not possess the fundamental practical reasons that it is the business of the moral philosopher to uncover (nonetheless, these reasons are still reasons for such agents).

To add some flesh to this distinction, let me very briefly describe my present view as to how it fits together with R. Not all of the claims I will make in this paragraph are in my book, although many of them are; in any case, they harmonize with things I say there. One’s normative reasons are all facts (true propositions), and the facts that are one’s reasons are one’s evidence that one ought to ϕ . Some of these reason facts are known, and some of them one is merely in a position to know. Unpossessed reasons are those reason facts one is in a position to know, but do not know. Possessed reasons are known reason facts. What is it to be in a position to know a fact? With respect to ordinary contingent facts, being in a position to know them requires, roughly speaking, that they be present in one’s immediate environment—as when the keys are on the table, staring one in the face (an example used in Gibbons 2013)—or be able to be immediately recalled from memory, or be able to be reached by a simple piece of reasoning. With respect to fundamental *a priori* accessible normative and evaluative truths, being in a position to know the relevant truths is simply guaranteed by them being *a priori* accessible (I do not wish to claim that all *a priori* accessible truths count as truths one is in a position to know, but there are good reasons to think fundamental normative and evaluative truths have this status). Of course, to say that some truth is *a priori* accessible is not at all to say that it is easily accessible, and this will be important to bear in mind whenever I use the phrase “in a position to know”. Finally, let me reiterate a related claim about the strength or weight of reasons that I make in Ch. 4 of my book: we

¹³ There is a minor logical difference between the “unpossessed”/“possessed” distinction and the “is”/“has” distinction: if one has a reason to do some act, then there is a reason to do that act, but if a reason to do some act is possessed it isn’t also unpossessed (one’s unpossessed reasons to ϕ are reasons for one to ϕ with respect to which one doesn’t stand in the having relation).

¹⁴ This is not exactly how I put it in my book, because, as I have indicated, I hadn’t yet opted for perspectivalism. I made the less precise claim, “Plausibly, a fact F can only be evidence for me that *p* if an *idealized* counterpart of myself would have this evidence (that is, would stand in appropriate relations, cognitively speaking, to F and *p*).” (14)

can understand the weight of a possessed reason to ϕ as being equivalent to the relevant epistemic probability that I ought to ϕ , and the content of one's actual *a priori* knowledge of ethical truths will play a crucial role in determining this probability when we are talking about practical reasons. If we are thinking about the weight of unpossessed reasons for agents to do various acts (such as, with respect to ordinary agents, the fundamental reasons that concern ethical theorists), all facts that the agent is in a position to know, including, crucially, all *a priori* accessible fundamental ethical truths, will be relevant.

Given what I just said about the need to transpose Littlejohn's arguments into a different key, his main argument against R can be restated as follows, without weakening it.

1. All evidence is constrained by the relevant agent's epistemic position.
2. If so, the Commonality Thesis implies that all reasons are constrained by the relevant agent's epistemic position.
3. Not all reasons are constrained by the relevant agent's epistemic position.
4. Thus, the Commonality Thesis is mistaken.
5. If the Commonality Thesis is mistaken, so is R.
6. Thus, R is mistaken.

Assuming, as I now do, that 1 is true, I deny 3. Hence it is clear that I am endorsing perspectivalism. I could, alternatively, deny 1, and have sometimes been inclined to do so in the past, but I worry that doing so must either commit one to a highly controversial account of evidence (where subject-independent propositions stand in subject-independent evidential relations to each other), or broaden the meaning of the word "evidence" to such an extent that R is no longer an interesting, substantive claim.¹⁵ It also appears to radically divorce possessed from unpossessed evidence/reasons.¹⁶

Perspectivalism, as I understand it, must be committed to what Littlejohn calls Epistemic Containment, when it is combined with *reasons as evidence*: "A fact is a normative reason only if it is a fact that belongs to some body of evidence by virtue of some subject bearing an interesting epistemic relation to it (e.g. knowledge, possible knowledge, etc.)".¹⁷ Littlejohn presents two arguments against Epistemic Containment, which I will

¹⁵ Ralph Wedgwood has expressed the latter concern to me.

¹⁶ Hawthorne and Magidor (Forthcoming) object to *reasons as evidence* that it seems highly counterintuitive to say that the fact that a reliable book says that an apple I am now holding is poisonous is an unpossessed reason for me not to eat it, in a case where that book is hidden in a cave on the other side of the world. A key advantage of going perspectivalist is that this objection no longer has any bite, as they rightly do not think there is anything odd about saying that a fact that a reliable book says an apple is poisonous is a reason not to eat the apple when one has just read the book (nor, I would contend, when the book is open on one's desk and the relevant sentence is staring one in the face). A more theoretical point also worth making is that the non-perspectivalist version of *reasons as evidence* will need to deny that possessed reasons can be factored into (unpossessed) reasons and possession conditions (given certain reasonable assumptions).

¹⁷ Littlejohn's inclusion of possible knowledge here points to a difference between perspectivalism and Epistemic Containment, since possible knowledge is a much broader notion than that of being in a position to know. Errol Lord (2015) defends a version of perspectivalism that isn't committed to epistemic containment. Like me, Lord wants to rely on a distinction between unpossessed and possessed reasons, but, unlike me, he thinks that unpossessed reasons need not stand in any interesting epistemic relation to agents for whom they are reasons, so he need not accept Epistemic Containment (also, unlike me, he thinks that what one ought to do depends merely on the reasons one possesses). One of the selling points of perspectivalism for Lord is that it delivers the intuitively right verdict about the well-known Mine Shaft case. The perspectivalism I am defending does this as well, since one knows all the relevant facts one is in a position to know in the Mine Shaft case.

now consider in turn. The first argument appeals to our seemingly shared endorsement of a principle regarding practical deliberation, the knowledge norm. The first thing I should say here is that I do *not* think it's always the case that one *ought* to only treat known facts as premises in practical deliberation. It's true, as Littlejohn says, that I endorse one interpretation of the knowledge norm in my book, but I was careful to state the principle I stand by in this way: "It is *rationally appropriate* for *A* to treat *p* as a reason to ϕ only if *A* knows *p*." (p. 116, some italics added). I claim, in other words, that it is only *fitting* or *appropriate*, by the rational standards that govern or are constitutive of (good) reasoning, to treat a proposition as a premise in practical reasoning when one knows that premise, just as it is only fitting or appropriate to believe propositions that one knows or will thereby know, by the standards that govern or are constitutive of (good) beliefs.¹⁸

Let's focus on the details of Littlejohn's first objection. Fred treats *p* as a reason, and *p* is unknowably unknowable. It follows from my interpretation of the knowledge norm that it is not rationally appropriate to treat *p* as a reason. And it follows, on Williamson's (2000) knowledge-based account of evidence (which, in broad terms at least, I endorse), that *p* is not a reason. It does not follow that Fred ought not treat *p* as a reason, nor, hence, that there is a reason not to treat *p* as a reason (if it did, we would have discovered that there are some reasons that are not evidence, as the objection would have it). Is it rational (although not appropriate) for Fred to treat *p* as a reason? Ought he treat it as a reason? There will be cases where the answer to both of these questions is positive. It will be rational for Fred to treat *p* as a reason if it is rational for him to believe that *p* is a reason (or that it is evidence that he ought to do or believe something).¹⁹ And Fred ought to treat it as a reason if, were he to know everything he is in a position to know, it would be rational for him to believe that *p* is a reason.

I think this takes care of Littlejohn's first objection, but one might worry it does so at too high a cost. I would emphasize that I think run of the mill cases will be cases where one rationally and appropriately treats a reason as a reason, and not just rationally but inappropriately takes some proposition to be a reason. Importantly, it isn't a feature of my view that one can knowingly and rationally take something to be a reason when it isn't a reason (Fred falsely believes *p* is a reason)—the whole point of Littlejohn's objection is that there are cases where ignorance runs deep, and it's only when ignorance runs deep that rationality and accuracy with respect to what counts as a reason can come apart. Such ignorance offers a good excuse. Ideals or standards can be important even

¹⁸ I use the phrase "rationally appropriate" in the book, because I wish to indicate that I have in mind appropriateness with respect to rationality, and I was thinking that what is rational and what is appropriate might come apart, with respect to beliefs and reasoning. It will sometimes be rational, but not appropriate, to believe unknown propositions (in Gettier cases, notably). And in some cases one may be in a position to know a proposition, hence it would be appropriate to believe it, yet it may not be rational to do so, either because one must first go through some reasoning or first attend to something presently perceptible (transitions from being in a position to know to knowing involve the relevant beliefs being formed on the right basis or using the right method), yet the closest possibilities involve one not doing that, or because one possesses misleading higher-order evidence (I'm assuming it's not always the case that the possession of misleading higher-order evidence prevents agents from being in a position to know).

¹⁹ This is a rough characterization of what is required. Rather than requiring Fred to be in a position where it is rational to believe *p* is a reason, one might want to merely say he needs to have a rational true belief that *p* (as in Gettier cases). "Rational" here corresponds to what many epistemologists mean by "justified" (like Littlejohn, I'm inclined to think it is best to reserve "justified" for cases involving more than the kind of justification that is available in Gettier cases).

when one can't always rationally live up to them. Furthermore, given the role that essentially ethical knowledge ends up playing in my book in explaining the weight of practical reasons, and that some philosophers (e.g. ethical nihilists) think that such knowledge is never available to us, it is highly desirable for me to leave room to be able to assert that "even if moral realism... is false, we may still have been acting rationally, in this narrower sense, all along" (135). Still, some perspectivalists, in particular, see it as a virtue of perspectivalism that it's not possible to be rationally mistaken as to what one's reasons are.²⁰ And some fans of the knowledge norm might have particular objections to the interpretation of this norm that I favor, and flesh out in the book (I claim there that *reasons as evidence* provides a much needed supplement to the norm; it cannot by itself explain which bits of knowledge it is appropriate to act on, nor distinguish between when it's appropriate to treat known facts as *pro tanto* reasons and when it's appropriate to act on them).

Rather than discuss Littlejohn's second objection in as much detail, let me simply indicate that I'm far from convinced that Agnes, given what she knows, ought to think that she has strong enough evidence, with respect to each room, considered individually, to believe that she ought to illuminate it. It seems she had strong enough evidence to conclude this prior to her encounter with the reliable guru who tells her she ought not illuminate all the rooms, but I'd say she now has evidence that blocks her from rationally concluding that she ought to turn the light on in room 1 (and so on).²¹ What would it be rational for Agnes to do and believe? Here, an important difference between belief and action is relevant. We always need to do something or other (when we're not asleep, etc.), but we don't always need to believe propositions that we are considering. So, Agnes does not need to believe she ought to illuminate room 1 (although she would be right to assign quite a high credence to the proposition that she ought to), but she does need to decide whether or not to illuminate room 1. Assuming the guru hasn't revealed anything about the stakes of wrongly turning the light on being particularly high, I think it is rationally appropriate for her to turn the lights on in room 1 (and so on, for each room). This follows from the more general principle that it's rationally appropriate to act on one's strongest reasons (see 115, where a precise statement of this principle is provided). In actual fact, with respect to all but one of the rooms, Agnes ought to turn the lights on (and she knows this much). Still, we won't, on this approach, find a deontic conflict is generated, but simply a conflict between what it is rationally appropriate for her to do, and what she ought not do (she ought not turn on one of the particular lights, and she ought not turn them all on). This doesn't seem odd or worrisome to me, given that it's a feature of cases like this one that there is a kind of ignorance at play regarding how one could go about satisfying all of the legitimate ought claims that one knows apply to one in these circumstances.

Littlejohn ends his paper on a helpfully forthright note. I agree with him that, at base, our disagreements hinge on a question as to what one should expect philosophical talk of

²⁰ Gibbons (2013), for instance, appears to see things this way (see Star 2016 for a relevant discussion of Gibbons). He defends the principle that if you justifiably believe that you ought to ϕ , then you ought to ϕ . It's worth noting, however, that he thinks what you justifiably ought to believe depends on a subset of all of your epistemic reasons, viz your good epistemic reasons, and he thinks of knowledge as the gold standard with respect to good epistemic reasons.

²¹ Is Littlejohn assuming that the ought in question is *not* evidence relative? I'll assume he isn't begging the question, and that he doesn't need to in order to make his objection.

reasons to do. Are we to understand reasons merely as tools in theories, such that it doesn't matter whether reasons are accessible to agents, as long as they play an explanatory role, or, alternatively, as guides to deliberation that must be accessible to agents? I think I made it clear in my book how I believe we should answer this question. Here I would add merely that, in this respect at least, I am part of a venerable tradition in ethics, that is, I hope, on the right side of history. In rejecting something that is also often associated with this tradition, as it runs through Bernard Williams's work, in particular, and that is the suspicion that highly general, principle-based normative ethical theory is essentially a waste of time, I think it's not unreasonable of me to hope I am likewise on the right side of history.²²

References

- Arpaly, Nomy (2003). *Unprincipled Virtue*. Oxford: Oxford University Press.
- and Timothy Schroeder (2014). *In Praise of Desire*. Oxford: Oxford University Press.
- Bommarito, Nicolas (Forthcoming). “Knowing Better: Virtue, Deliberation, and Normative Ethics,” by Daniel Star.” *Australasian Journal of Philosophy*.
- Driver, Julia (2001). *Uneasy Virtue*. Cambridge: Cambridge University Press.
- Finlay, Stephen (2014). *Confusion of Tongues: A Theory of Normative Language*. Oxford: Oxford University Press.
- Gibbons, John (2013). *The Norm of Belief*. Oxford: Oxford University Press.
- Hawthorne, John and Magidor, Ofra (Forthcoming). “Reflections on the Ideology of Reasons,” in Daniel Star (ed.), *The Oxford Handbook of Reasons and Normativity*. Oxford: Oxford University Press.
- Kearns, Stephen and Star, Daniel (2008). “Reasons: Explanations or Evidence?” *Ethics* 119 (1): 31–56.
- (2009). “Reason as Evidence.” *Oxford Studies in Metaethics* 4: 215–42.
- (2011). “On Good Advice: A Reply to McNaughton and Rawling.” *Analysis* 71 (3): 506–508.
- (2013a). “Reasons, Facts-About-Evidence, and Indirect Evidence.” *Analytic Philosophy* 54 (2): 237–43.
- (2013b). “Weighing Reasons.” *Journal of Moral Philosophy* 10 (1): 70–86.
- (2015). “Weighing Explanations,” in Iwao Hirose & Andrew Reisner (eds.), *Weighing and Reasoning: Themes from the Philosophy of John Broome*. Oxford: Oxford University Press, 232–51.
- Kolodny, Niko (Forthcoming). “Instrumental Reasons,” in Daniel Star (ed.), *The Oxford Handbook of Reasons and Normativity*. Oxford: Oxford University Press.
- Lord, Errol (2015). “Acting for the Right Reasons, Abilities, and Obligation.” *Oxford Studies in Metaethics* 10: 26–52.

²² I especially have in mind Williams's *Ethics and the Limits of Philosophy* (1985). I agree with Williams that philosophers should be interested in ethics in general, and not just morality narrowly construed, and I take it that it is a virtue of *reasons as evidence* that it allows us to understand how agents can weigh prudential reasons against moral reasons without needing to label reasons in one way or the other, and without themselves possessing any theory about the relationship between prudence and morality (Star 2015, 20–21).

- Star, Daniel (2015). *Knowing Better: Virtue, Deliberation, and Normative Ethics*, Oxford: Oxford University Press.
- (2016). “John Gibbons. *The Norm of Belief*. Oxford: Oxford University Press, 2013.” *Ethics* 126: 1139–48.
- (Forthcoming). “Reasoning with Reasons,” in Conor McHugh, Jonathan Way, and Daniel Whiting (eds.), *Normativity: Epistemic and Practical*. Oxford: Oxford University Press.
- Williams, Bernard (1985). *Ethics and the Limits of Philosophy*. London: Fontana Press.
- Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Worsnip, Alex (2016). “Daniel Star, *Knowing Better: Virtue, Deliberation, and Normative Ethics*, Oxford University Press, 2015.” *Notre Dame Philosophical Reviews*.