

Gibbons, John. *The Norm of Belief*. Oxford: Oxford University Press, 2013.

“From Outside of Ethics” Review, *Ethics*, Forthcoming.

Daniel Star

I.

John Gibbons’s impressive monograph attempts to resolve a disagreement between ethicists and epistemologists when it comes to the topic of normative reasons. Epistemologists tend to think of reasons (for belief) as perspective-dependent, and ethicists tend to think of reasons (for action) as perspective-independent. Some people, across both camps, don’t think of this as a source of disagreement at all: reasons for belief are simply fundamentally different than reasons for action; it is sometimes said that the first are “subjective” and the second “objective”. Gibbons contends, rather plausibly, that this is an unsatisfactory place to end up. We want to know what unifies reasons for action and reasons for belief, which is to say that we want a general explanation of why many things that we are inclined to put under each heading really do count as reasons. While Gibbons demonstrates a keen appreciation of some of the considerations that draw philosophers to objectivism about reasons, he comes down firmly on the side of subjectivism. He thinks it is a mistake to overlook or downplay the connection between reasons and reasonableness, and that this is where the objectivist goes wrong. At the same time, Gibbons aims to defend a version of subjectivism that explains some of the key intuitions that make objectivism tempting, and that adopts some of the virtues of that position.

This way of describing his project might lead one to think that the title of Gibbons’s book is misleading. His concerns certainly are broader than it suggests. Throughout the book, he moves back and forth between the general disagreement concerning reasons just described, and the disagreement between subjectivists and objectivists as it plays out in arguments regarding theoretical reason, in particular. With respect to the norm of belief, the issue is whether we should believe what is true (or only things that we will thereby know), or whether we should instead simply believe whatever it’s reasonable to believe. Gibbons spends the last three chapters of his book masterfully articulating and defending his specific version of the view that epistemic *justification*, rather than *truth* or *knowledge*, is the norm of belief. These chapters constitute an important contribution to epistemology, and they may be read independently of the rest of the book; they should be read by everyone with a serious interest in theoretical reason. My aim in this review is to explore what ethicists might make of this book, so, following a short summary of each chapter, I will focus on what Gibbons has to say about normative reasons and requirements in general, rather than the norm of belief *per se*.

The book consists of a preface and eleven chapters. In the Preface, Gibbons provides a brief introduction to the problem of reconciling the view of reasons common amongst ethicists with the view of reasons common amongst epistemologists. He begins by describing some plausible motivations for thinking of reasons in general as non-perspectival, and some plausible motivations for thinking of reasons in general as perspectival. The ethicists are right to reject theories that would entail that reasons can supervene on merely introspectively accessible properties, for since my brain-in-a-vat twin can not be required to take care of his child, it would follow that I am not required to take care of my child (viii). But the epistemologists are right to be wary of the “equally insane” idea that facts that are epistemically inaccessible to agents could make a difference to what they ought to do: “if a code of

conduct tells you to do irrational things on a daily basis, you shouldn't take that code of conduct seriously, and calling it an ethical theory simply does not get around the problem" (viii). The key for reconciling these convictions lies with taking seriously the possibility that an agent's reasons in general depend on what the agent is in a position to know, where what an agent is in a position to know is not to be limited to what is internally accessible, although we might still think of the relevant facts as being part of our mental landscape, broadly speaking: "On one way of looking at things... reasons have one foot in the mind and another in the world... But on another way of looking at things... if what you can see from where you are doesn't supervene on the internal, then neither does your mind." (xi).

Chapter 1, "The Puzzle", provides an extended discussion of the manner in which the subjectivist and objectivist options arise in debates about the norm of belief, in particular. The puzzle is that all of the following claims – each providing a different view as to what the norm of belief is – are independently attractive, but they direct us to believe different things: (T) You ought to believe p only if p is true; (K) You ought to believe p only if you would thereby know that p ; and (J) You ought to believe p if and only if you are justified in believing p . Objectivists accept (T) or (K), and subjectivists accept (J).

In Chapters 2 and 3 (Part I), Gibbons focuses on the response to the idea that we must choose between subjectivism and objectivism that would have it that there is no need for us to make a choice at all. Perhaps "ought" is lexically ambiguous (Ch. 2), or context-sensitive (Ch. 3). Gibbons's sensible strategy here is to studiously avoid getting bogged down in arguments in the philosophy of language. The debate between subjectivists and objectivists isn't simply one about which sentences are true – either lexical ambiguity or context-sensitivity might deliver up plenty of true sentences that fit (T) to a glove *and* plenty of true sentences that fit (J) to a glove. The problem is that the availability of such true sentences would do nothing to undermine the sense that one would be in a seriously conflicted state if one tried to wholeheartedly follow both subjective and objective norms.

The focus of Chapters 4 and 5 (Part II) is on objectivism, and strategies available to the objectivist who wants to hold on to one of (T) or (K), as well as (J), for explaining the relationship between the objective norm and the subjective norm. In Ch. 4, Gibbons considers two strategies that explicitly take it that (J) is a genuine, but secondary norm, whose role is derived from the more fundamental (T) or (K). He then turns, in Ch. 5, to the teleological view of belief, which has it that either truth or knowledge is a *constitutive* aim of belief. He argues, rather convincingly, that these strategies for holding on to both a subjective and an objective norm are unsuccessful.

Chapters 6 and 7 (Part III) are concerned with explaining what an intuitive reaction to certain claims about reasons that Gibbons calls "the natural reaction" commits us to, and they constitute a valiant effort to undermine objectivism about reasons and requirements. Suppose I am in a building that is on fire and I could get to safety by jumping out of the window into the river below, but I have no evidence whatsoever that the building is on fire. Someone looking at me from afar might say that there is a reason for me to jump out of the building, and even that I ought to jump out of the building. But it would be *unreasonable* for me to get up and jump out of the window, and there is no rational way I could follow any norm concerning reasons that might be said to apply to me at this time that would direct me to jump out of the window. Or suppose all my evidence suggest that turning left will take me to the restaurant, but actually turning right will take me to the restaurant. It would be *unreasonable* for me to nonetheless believe that turning right will take me to the restaurant, as (T)

would say I should. The *natural reaction* is the response that what makes appeals to norms of this kind inappropriate is that they can not get a grip on us.

The natural reaction is said to involve two notions: *guidance* and *accessibility* (the first is explored in Ch. 6 and the second in Ch. 7). Reasons and requirements must be able to guide us in belief formation and action, and they can only guide us when they are accessible in the right way. Guidance does not require that our reasons are always transparent to us (the relevant mental states need not be luminous), and Gibbons is keen to argue that many supposed problems for subjectivism are really just problems for internalism about epistemic justification. He here argues for an account of accessibility that extends to facts that we are in a position to know, an idea I will say more about below.

The last three chapters of the book (Part IV) are concerned to show that the subjectivist can give us everything we need in order to understand certain intuitions about truth and knowledge that appear to favor objectivism. The most important ingredient here involves the notion of *commitment*. We are told (in Ch. 8) that rather than thinking of justification as a means for accumulating true beliefs or knowledge, we should think that whenever we believe a particular proposition we always *ipso facto* commit ourselves to knowing that proposition (similarly, when we intend an end we thereby commit ourselves to taking the necessary means to that end). An extended examination of Moore-paradoxical propositions reveals that the problem that arises for any agent that would attempt to believe them is not that some objective condition thereby fails to be met (as others have argued), but is instead that the agent must, in so doing, view herself as irrational (Ch. 9). Furthermore, any attempt to view knowledge or truth as merely optional (as a naïve subjectivist might) will lead to Moore-paradoxical claims, hence irrationality (Ch. 10). This is what is said to explain why one is always committed to knowing *p* when one believes *p*.

2.

The arguments of *The Norm of Belief* are complex, multifaceted, and highly ambitious. The case for subjectivism about reasons in general is, in large part, developed by focusing on the case for thinking the norm of belief is not best construed as an objective norm, and that properly understanding it requires a particular form of subjectivism. For the ethicist drawn to objectivism about practical reasons, a simple response is tempting: even if one were to grant that the author has established that the norm of belief is justification, and that epistemic reasons are perspective-dependent, one would not need to grant that *all* reasons are perspective-dependent, given that the case for thinking they are largely rests on arguments that concern the norm of belief, in particular. It might even be granted that it would be nice to be able to be either a subjectivist or an objectivist about all reasons, rather than have a mixed view. But what is nice is not always possible.

As the detailed description of the book provided above indicates, Gibbons does have some resources for dealing with the mixed view of reasons that do not spring merely from considerations concerning the norm of belief. In particular, a reader might be persuaded to accept the view that all genuine normativity involves guidance and accessibility, as he understands these two conditions. Nonetheless, an odd feature of his book is that he says close to nothing at all about one of the key motivations for objectivism about *practical* reasons, so far as many ethicists are concerned: *moral* objectivity. It is here that I think he leaves himself most vulnerable to attack on the basis of intuitions that ethicists take seriously. This is *not* because there is no good way to square perspective-dependence

with moral objectivity – there is a way to do so – but because it is far from clear that this option is available to Gibbons, and there are some reasons to think it is not.

I don't want to give the reader the wrong impression. Gibbons does focus on some people's intuitions concerning moral objectivity, inasmuch as he focuses, repeatedly, on the idea that somehow we ought to do what is actually best. This is an idea he forcefully rejects, on the grounds that what is actually good is generally beyond our ken. However, the philosopher concerned with moral objectivity need not accept that what we ought to do is what is actually best. She might, for instance, think that what we ought to do is whatever the balance of *pro tanto* duties that presently applies to us directs us to do. Even if she does adopt a consequentialist framework, she may opt for an expected value theory (Gibbons does raise some particular objections to expected value theory [130-31], but they are concerned with guidance, rather than moral objectivity). Fundamental moral principles can be formulated either in such a way that what we ought to do depends on worldly facts (what is actually best, or perspective-independent *pro tanto* duties, etc.), or in such a way that what we ought to do depends on features of our epistemic position. The idea that some moral principles of either kind are true, and that the beliefs of rational people converge on them, is an idea that Gibbons does not discuss. Even with respect to the good, he does not discuss essentially ethical claims about what types of things are fundamentally good (claims that expected value consequentialists, as much as actual value consequentialists, must commit themselves to).

What additional claim does the Gibbons-style subjectivist need, then, in order to be true to the idea that we should all accept one set of fundamental ethical principles? Very plausibly, the claim that is needed here is that rational agents are all in a position to know the same set of ethical principles *a priori*, where these ethical principles direct us to focus on our evidence, or what we know, rather than on perspective-independent facts.

Most of the time Gibbons is careful to stay true to his unified subjectivist picture of practical and theoretical reason, but there are moments where one might be forgiven for thinking he is having second thoughts: "...I think that epistemic justification supervenes on what you're in a position to know. But when it comes to practical and affective reasons, it may well matter what you want, or value, or care about. Or it may depend on what you ought to want, or value, or care about. And it may well be that these things are not determined by the evidence. This is a big question, and we just have to leave it open." (177) We are not told in what sense it may be the case that we *ought* to value anything in particular, but Gibbons must be either committed to the idea that the ought in question here is the "subjective" ought, or that there is nothing one ought to value. Given this, why not contend that practical justification with respect to what we ought to value (if there is anything we ought to value), will also turn on what we are in a position to know?

It might seem unfair to complain about this lacuna regarding our access to fundamental ethical truths. Gibbons has enough on his plate trying to get clear about the foundations of theoretical reason. However, we are told to pursue unity repeatedly throughout the book. It would certainly be too much to ask for a substantive account of essentially ethical knowledge, but some suggestions about how we might hope to find a unity to reasons, across the practical and theoretical divide, when we think about our moral commitments, might have been helpful. And by leaving it open that we might instead think that practical reason is simply concerned with what we happen to want, Gibbons leaves himself open to the charge that his subjectivism can not, after all, give us what many actual objectivists want.

I just suggested that the Gibbons-style subjectivist might appeal to the availability of *a priori* ethical knowledge when thinking about whether we are in a position to know ethical principles. But there are, of course, things that we might worry about here. Most relevantly, it seems that much essentially ethical knowledge (assuming there is such knowledge to be had) is quite difficult to come by. Although it may not require much time and effort to come to know that lying is normally wrong, it may require a great deal of time and effort to come to know the correct fundamental normative ethical principles (the principles of a specific ethical theory). This is what makes normative ethics a research program. If some such difficult-to-arrive-at principles are *a priori* knowable, should we say that *everyone* is in a position to know them, or not? This question is particularly worth raising in the present context, because of the central role that being-in-a-position-to-know plays in Gibbons's theory of reasons and requirements.

We are actually given surprisingly little guidance as to what is to be included when we are thinking about what we are in a position to know. Here is Gibbons's main example: "I'm in the dining room searching frantically for my keys. And they are staring me in the face on the otherwise empty dining room table. If this were an actual situation, and you were standing there watching me, you wouldn't be able to keep that smile from your face. When I give up and start searching the kitchen, I believe my keys are not in the dining room. But I should have known better." (179). The focus here is not on it seeming to one that the keys are in the dining room (after one heads to the kitchen, perhaps). There can be cases where I am able to come to know that the keys are in the dining room through introspection when in the kitchen, but Gibbons rejects the notion that what one ought to know supervenes on what is available through introspection. So, the paradigm case is simply one where evidence sufficient for knowledge is right in front of me in the external world, detectable by my senses. It seems quite reasonable to require me to know that the keys are on the table when I am in this situation. *Obviousness* is the crucial thing here: "What we need here is not a precise metric for measuring obviousness and a specific number on that metric above which knowledge is required. All we need is the idea that at least sometimes people ought to know the really obvious facts that are staring them in the face whether they believe those things or not" (179).

Many ethical truths are far from obvious, and some – the principles we aim to discover when doing normative ethics – may be anything but. Perhaps the subjectivist should say that we are *not* in a position to know them, in any sense that would connect far from obvious ethical principles we happen to be ignorant of to our normative reasons, even if they are knowable *a priori*. This would be a natural thing to conclude, given the emphasis on obviousness and reasonableness in many places in the book. However, there is significant room for doubt here, because of the particular way Gibbons is committed to relying on a distinction between good and bad reasons, and his rejection of epistemic internalism. These things take us away from, and are in tension with, the appeal, here and elsewhere, to obviousness.

3.

Ethicists generally distinguish between motivating reasons and normative reasons. If all the evidence I possess suggests there is gin in the glass in front of me, when there is actually petrol in it, I can have a motivating reason to drink from the glass, but no normative reason to drink from it. If the building I am in is on fire, but I possess no evidence that it is on fire, I will have no motivating reason to jump out the window into the river below, yet will have a normative reason to jump out of it.

Interestingly, these are verdicts most internalists and externalists about practical reason agree about. Gibbons thinks this is all a big mistake. Like pretty much everyone else, he thinks motivating reasons are the kind of reasons we cite when providing rationalizing explanations of actions. But normative reasons are just a subset of motivating reasons, on his view.

Crucial to Gibbons's subjectivism is a distinction between *good* reasons and *bad* reasons: "...to say that these [considerations] are motivating reasons is only to say that the transition from them to [a belief or action] is an instance of a general type of transition. And we can say something good about the transition type without saying anything about what it's a transition from on a particular occasion. So calling them reasons doesn't mean that they're good reasons. ... Motivating reasons are the kinds of things that make things reasonable. But that doesn't mean that they always succeed in making things reasonable. If there's something rationally criticizable about the inputs to the transition, there's something rationally criticizable about the outputs as well" (155). This distinction is crucial in at least two ways. First, when we ask what we *ought* to believe we are *not* asking what we have most reason to believe; Gibbons claims (on 177, for example) that we are asking what we have most *good* reason to believe (and *mutatis mutandis* for practical reasons and requirements). Second, if we were simply asking what we have reason to believe (not focusing on what we have good reason to believe), it would not follow that what we have reason to believe supervenes on what we are in a position to know, which is another of his crucial claims.

A "good reason" is simply what many of us call a "normative reason", with the caveat that there are no normative reasons of the kind that the objectivist accepts (no reasons to jump out of burning buildings when your evidence is impoverished, etc.). In ordinary language, we often pick out normative reasons by using "good reason". In a classic discussion of normative reasons, T. M. Scanlon writes, "...when I say that something is or is not a reason I will not be concerned with whether it is or could be someone's operative reason but with whether it is a *good* reason – a consideration that really counts in favor of the thing in question." (*What We Owe to Each Other* [Cambridge, MA: Harvard University Press, 1998], 19). Here we see both the identification with good reasons that Gibbons accepts, as well as the objectivism he rejects. Gibbons thinks even good reasons are motivating reasons, whereas objectivists, like Scanlon, think they are not. Nonetheless, they *all* agree that the good reasons are the ones we really want to be following when forming beliefs, and acting. We should not conclude that this debate is merely terminological. Gibbons contends the objectivist will have no need to speak of *bad* reasons (155). This is because they don't think motivating reasons are really reasons to do or believe anything at all, whereas Gibbons thinks bad motivating reasons are still reasons to do or believe things.

It is a significant feature of the particular way in which Gibbons develops his case for subjectivism that he appeals to common judgments that, with respect to various cases, we can only be normatively required to do or believe what is reasonable from the subject's own perspective. He takes "Be reasonable!" to be a categorical imperative, in the sense that it applies to all of us, regardless of our contingent desires (5). However, it turns out that, on his own account, the way the world appears to subjects can come apart from their *good* reasons and what they are required to believe and do (since such requirements concern good reasons only). Considering brains in vats, Gibbons writes: "If genuine reasons are things that make things reasonable, then the regular 'ought' expresses the requirements of rationality. If what you are required to do is determined by your nonfactive mental states, then if you duplicate the inner life, then you duplicate the requirements. So you'll be required to do the same things

that your recently envatted twins are required to do. And since they can't get out of their vats and go to the store for milk, you can't be required to get out of your chair and go to the store for milk." (175)

For Gibbons, good reasons for belief are epistemically justified. The notion of justification at work here is externalist, in that justification supervenes on more than the inner mental life. Justification is not, in general, perspicuous. We ought to believe we have hands, but our envatted twins may not be required to believe they have hands (and we, but not our envatted twins, are required to know we have hands). They are obviously not required to know that they have hands. It may appear to them that they are justified in believing they have hands, but it may be the case that they are not. If so, they may think they are justified in believing they ought to believe they have hands, but they are not. This is because, in a very interesting discussion, Gibbons defends the claim that if you are justified in believing you ought to believe a proposition then you ought to believe it (164-70); it would follow from it not being the case that our envatted twins ought to believe that they have hands that they are not justified in believing they have hands.

There is a complication here that points to a significant difference between reasons for action and reasons for belief. Gibbons notes that, despite what I just said, it is far from obvious that our envatted twins are not justified in believing they have hands: "None of your twins in vats can see that they have hands. So they believe for different reasons. This doesn't automatically entail that they shouldn't believe that they have hands. It all depends on what they have most reason to believe and what they're in a position to know" (187). Here Gibbons probably has in mind that recently envatted twins have many relevant justified beliefs in the form of relevant memories, and that all envatted twins may be in a position to know that it *appears* to them that they have hands, hence may justifiably believe they have hands on this basis.

When it comes to reasons for action, the gap between perception and reality will, it turns out, be much wider than this. One might be tempted to reason as follows. When we find we are out of milk, we ought to believe we should go to the shop to get milk – in fact, describe the situation in the right way, and we are required to go to the shop to get milk. But, we have seen, Gibbons is adamant that our envatted twins are not required to go to the shop to get milk. They may think they are justified in believing they ought to go to the shop, but, in fact, they are not justified in believing this (given the principle that if you are justified in believing you ought to do something then you ought to do it, it would appear to follow from it not being the case that our envatted twins ought to believe that they should go to the shop that they are not justified in believing they should go to the shop.) However, there is a problem with this line of reasoning. Gibbons wants to restrict the range of actions where we are allowed to apply the principle that says that if you are justified in believing you ought to do something then you ought to do it. The restriction in question is to cases where one is in a position to do the relevant act *for* a reason (188). Since my envatted twin can not go to the shop, it can not go to the shop for a reason. On the other hand, it may be able to believe that it has hands for good reasons.

Despite how things will *appear* to my envatted twin, he has no reasons to do all the very many acts it appears to him that he has reason to do. And he may not even have reasons to *try* to do these things, given what Gibbons has to say about reasons to try to do things being dependent on reasons to do things (175). "Be reasonable!" appears not to be an imperative he will be able to follow in relation to his actions, even if he can follow it in relation to his beliefs.

The basic point here is that it seems like normative requirements on this picture ask too much of some possible agents, and that this is a source of tension in Gibbons's book, given the appeal to *obviousness* and pretheoretical intuitions concerning the rule "Be reasonable!" There will be cases, according to Gibbons's subjectivism, where what is apparent to agents will come apart from their good reasons. Gibbons might be entitled to think that this is a feature of his view, and not a bug. However that may be, Gibbons's appeal to *reasonableness* is not, despite initial appearances, theory neutral, given that what his subjectivism would ask of us will come apart from some of our intuitions about what it is reasonable to expect of us.

Gibbons tells us that the objectivist can be criticized for claiming that there are normative reasons out there that we are unable to respond to. His subjectivist, on the other hand, says that there are *good* reasons that it will *seem* wrong to restrict one's responses to in lieu of responding to other (bad) reasons. We can respond to the good reasons, but if, in some situations, we can not internally distinguish between them and the bad reasons, are we not being asked to do something intuitively *unreasonable* when we are told we are required to do so? The problem with our envatted twins is not that they must somehow leap out of their vats and respond to reasons they don't have access too, but rather that they must somehow stop responding to a lot of things they naturally take to be good reasons, and instead respond to a seemingly impoverished subset of their reasons (e.g. they must now form many beliefs on the basis of appearances, rather than on the basis of facts about the world). The problem is one of too few good reasons, rather than too many, yet there is something oddly similar between being in this predicament and the predicament that we find ourselves in if the objectivist take on reasons is correct.

Finally, let me return to difficult-to-arrive-at essentially ethical knowledge. Since, as we have seen, Gibbons actually distances himself from pretheoretical intuitions regarding which requirements are 'reasonable' to impose on agents, it is *not* at all clear that we should say that it is unreasonable to hold agents to the content of difficult-to-arrive-at *a priori* ethical knowledge (i.e. fundamental principles concerning what types of things are good, and what we ought to want, value, or do). Gibbons rejects the idea that we should hold them to a standard of accepting *all* knowledge that is *a priori* accessible (he rejects the idea that we ought to believe everything that logically follows from our justified beliefs [44]), but knowledge with non-contingent ethical content seems like a crucially important subset here. After all, Gibbons wants to avoid proliferating *oughts*: he would think it quite unattractive to say that sometimes agents rationally ought to do what they ethically ought to do (when they possess the requisite ethical knowledge) while, at other times, agents rationally ought to do something that comes apart from what they ethically ought to do (when they lack ethical knowledge). This means he needs to either conclude that there is no fundamental ethical knowledge that is difficult to arrive at, or accept that it's reasonable to hold people to the standards provided by the content that such knowledge has when people possess it.

Gibbons's subjectivist account of reasons is very much worth taking seriously, and how it ultimately compares with the best forms of objectivism is a question that deserves to be considered carefully. Thinking more about the norm of belief, and the many fascinating things that Gibbons has to say about it, should prove helpful when answering this question, especially given the attractive goal Gibbons shares with a number of us of trying to provide a unified account of reasons across ethics and epistemology. But it will also be essential to think much more about the nature of practical reasons.