NOÛS

NOÛS 00:0 (2019) 1–29
doi: 10.1111/nous.12291

Trusting Moral Intuitions

JOHN BENGSON
University of Wisconsin—Madison

TERENCE CUNEO
University of Vermont

RUSS SHAFER-LANDAU
University of Wisconsin—Madison

We develop an argument for a novel version of moral intuitionism centered on the claim that moral intuitions are trustworthy. Our argument employs an epistemic principle that we call the *Trustworthiness Criterion*, a distinctive feature of which is its emphasis on oft-neglected social dimensions of cognitive states, including non-doxastic attitudes such as intuition. Thus our argument is not that moral intuitions are trustworthy because they are regress-stoppers, or because they are innocent until proven guilty, or because denying their epistemic contribution would be self-defeating, or because they are presupposed in rational inference, or because they are analogous to perceptions, or because they are based on understanding—individualistic claims that have elsewhere been used (controversially) in defense of the thesis that moral intuitions are in good epistemic standing. Rather, our argument appeals to the idea that moral intuitions are trustworthy because they are the outputs of a cognitive practice, which has epistemically-fecund social elements, that is in good working order. This means that, as John Rawls writes in another context,

> if a person engaged in the practice is asked ... to defend what he does, then his ... defense lies in referring the questioner to the practice.[1]

On our approach, the trustworthiness of moral intuitions is to be accounted for not in terms of features of the intuitions themselves, but by reference to the broader practice in which the agents who possess those intuitions are engaged.

Our argument thus encourages a reorientation in philosophical work on moral intuitionism, which has largely operated within an individualistic framework. But first we must elucidate our notion of a moral intuition and get clear on what it is to be trustworthy. Once these clarifications are in place, we will be in a position to develop our argument that moral intuitions are trustworthy.

1. Moral Intuitions

We understand an intuition to be a conscious non-sensory mental state or event in which it *strikes* one that things are a certain way when one reflects on the matter.

For example, when addressing the question whether identity is transitive, it may consciously strike one that the answer is affirmative. Similarly for the intuitions that nothing could be colored without being a particular shade of color, that parallel lines must not converge, or that believing against all evidence is irrational. One or more of these claims may turn out to be false, and we may fail to believe or judge that one or more are true. But that is compatible with our intuiting all of them, its striking us that they are true—an experience in which it appears, from the inside, as if their truth is manifest, when we consider them. That experience is typically absent when considering the claims that the square root of 4489 is 67, that language is a precondition for thought, or that there is a fundamentally non-disjunctive type of entity that is simultaneously the semantic value of a full indicative sentence and the primary bearer of truth and falsity. These claims may all be true, and each may be believed, but they are not typically intuited.[2]

By 'moral intuitions,' we mean intuitions about moral matters, such as the moral status of a type of behavior (e.g., that it is *wrong* to recreationally slaughter fellow persons) or particular behaviors in particular circumstances (e.g., that it was *admirable* for you to come to the aid of a friend last week). Moral intuitions may also concern moral reasons, virtues, duties, or obligations, whether in the form of general principles (e.g., that beneficence is a duty), or verdicts about specific actual or hypothetical cases (e.g., concerning what moral reasons you would have in counterfactual scenarios).

Moral intuitions are sometimes the psychological bases of moral beliefs and judgments (and perhaps other mental states).[3] For example, you might have recently formed the belief that you should visit an ailing relative. You might have formed this belief because, when reflecting on the question of whether to visit her, it struck you that you should do so. This is not to say that moral beliefs are always formed in the light of moral intuitions; other types of mental states, events, or processes can serve as the full psychological bases of moral beliefs. (We return to this point later.)

These claims are not by themselves highly controversial or of deep philosophical significance. But the related claim that moral intuitions are trustworthy (a positive epistemic status) is both. For this claim entails the rejection of a prominent strand of skepticism about moral intuitions that many philosophers have found attractive. We will explain this connection in the next section, after clarifying the relevant notion of trustworthiness. But let us first make a comment about the position that our argument aims to establish.

We will defend a version of moral intuitionism according to which moral intuitions

> (α) have positive epistemic status, namely, trustworthiness;
> (β) do not have that status fully in virtue of facts about coherence, inferential role, or other such intra- or inter-attitudinal connections; and
> (γ) are in normal circumstances[4] sufficient to yield trustworthy moral beliefs.

It is worth clarifying that the version of intuitionism for which we will argue is non-standard in several respects.

First, moral intuitionists have often identified moral intuitions with non-inferential moral beliefs or judgments.[5] As will be evident, we understand intuitions to be not beliefs or judgments but rather non-doxastic epistemic bases of such doxastic states.

Second, intuitionists are sometimes reputed to hold that moral intuitions are infallible. Even if some intuitionists have held such a view, we want to go on record as rejecting it.[6]

Third, some intuitionists believe that intuitions originate in conceptual understanding[7] or concern truths for which there is no explanation or proof.[8] As indicated by our clarification of what we mean by 'moral intuitions,' the version of intuitionism we defend is more capacious, being free of such commitments.

Fourth, intuitionists have often focused on evidence or knowledge. By contrast, our version of intuitionism showcases a more basic epistemic status that we call 'trustworthiness,' which (as explained in the next section) subserves epistemically conscientious inquiry.

Fifth, as foreshadowed at the outset, while standard defenses of moral intuitionism focus on properties of the moral intuitional states of individual thinkers or agents, the version of moral intuitionism for which we'll argue emphasizes a range of features, with social dimensions, of the practice of intuiting.

Sixth, philosophers have often used the terms 'moral intuitionism' and 'moral nonnaturalism' nearly interchangeably. We do not use the terms in this way, holding instead that nonnaturalism is a metaphysical position regarding the nature of moral reality. Nonnaturalists have typically embraced one or another version of intuitionism, and vice versa. This is probably no accident, since the views are arguably natural companions. Still, neither entails the other; each can be endorsed independently of the other. What we argue here, then, does not rely on or presuppose the truth of nonnaturalism. While nonnaturalists will have a special interest in our project, a wide array of metaethical positions could endorse the version of moral intuitionism we defend here.

2. Trustworthiness

We now turn to the notion of trustworthiness that plays a central role in our argument. As we'll understand it, for an agent to *trust* φ is for her to 'cognitively rely' on φ—to employ φ to form, maintain, or revise other of her attitudes with the aim of increasing, improving, or otherwise positively contributing to her stock of knowledge or understanding regarding φ's contents.[9] For φ to be *worthy* of trust by an agent is for φ to be such that the agent is warranted in cognitively relying on it (or, equivalently, that φ warrants cognitive reliance). Thus, trustworthiness is warranted cognitive reliance. The notion of trustworthiness plays a central role in our discussion because it designates a property such that, when a state has that property, an agent is thereby epistemically conscientious (i.e., responsible to a high degree) in relying on that state in order to obtain understanding or knowledge.[10]

Warrant is a property that applies to a variety of agential states and events, both doxastic (such as beliefs) and non-doxastic (such as plans of inquiry). We will focus

on events, since our primary concern in what follows is reliance (an event). We will not here attempt a reductive analysis of warrant. Nor will we rely solely on ordinary usage or concrete cases (although we'll consider several examples below, in §4.1). The reason is that we find it more illuminating to highlight the ways in which warrant connects to a range of other important normative properties, which are familiar enough to provide a stable anchor for our explication and constrain our argument below (as explained in §4.2).

> First, warrant entails being favored; correlatively, warranting is a species of favoring.
> Second, warrant is gradable, coming in various qualities or degrees.
> Third, warrant entails permission: if an event is warranted for an agent, she is permitted to partake in that event.
> Fourth, warrant does not entail obligatoriness: that an event is warranted for an agent does not imply that the agent is required to partake in that event.
> Fifth, warrant entails blamelessness: if an event is warranted for an agent, she does not merit blame for partaking in that event.
> Sixth, warrant entails appropriateness: if an event is warranted for an agent, it is appropriate for her to partake in that event.
> Seventh, warrant entails reasonableness: if an event is warranted for an agent, it is reasonable for her to partake in that event.
> Eighth, warrant entails goodness: if an event is warranted for an agent, then it is good for her to partake, or there is value in her partaking, in that event.[11]
> Ninth, warrant is defeasible: if an event is warranted for an agent, then, at least in conditions of uncertainty, the relevant permission, blamelessness, appropriateness, reasonableness, and goodness may be defeated (i.e., fail to obtain).
> Tenth, warrant is an epistemic, rather than a merely practical, property in the sense that the above features do not hold merely in virtue of a particular event's connections to practical agency. The relevant permission, blamelessness, appropriateness, reasonableness, or goodness obtains not simply because (say) the event in question is useful or backed by practical reasons. Instead, they are at least partly grounded in various not merely practical facts or considerations (such as, for example, the fact that the state promotes accurate representation of reality, or that being in such a state contributes to the realization of intellectual virtues).

In what follows, we leave these elements of warrant at an intuitive level, in order to avoid prejudging various controversies whose resolution is independent of our project (e.g., whether goodness is always kind-relative). We also leave open the question of how all of the elements hang together, allowing that some might explain or mutually support others.[12] For our purposes here, what matters is that the foregoing yields a substantive and informative characterization of warrant: it is the gradable and defeasible property, not reducible to consistency, not identical to knowledge, and not requiring antecedent attitudes with that property, which an agential state or event possesses when and only when it is favored, permissible

though not necessarily obligatory, blameless, appropriate, reasonable, and good, and its having these features is not due merely to its connections to practical agency.[13] Thus a state φ is *trustworthy* for an agent when cognitive reliance on φ by that agent instantiates that property.

As we've indicated, trustworthiness is of considerable interest and importance because it is arguably a precondition for epistemically conscientious inquiry. After all, when agents lack any states that are trustworthy, they suffer a type of epistemic paralysis: they are *not* warranted in relying on *any* of their states to expand their stock of knowledge or understanding. But if they are not warranted in relying on any of their states to expand their stock of knowledge or understanding, then it is not epistemically favored, permissible, blameless, appropriate, reasonable, and good for them to rely on any of their states with that aspiration in view. In such a case, it clearly is not epistemically conscientious of them to engage in an activity, such as inquiry, that aims to increase, improve, or otherwise positively contribute to knowledge or understanding.[14] It follows that epistemically conscientious inquiry requires having trustworthy states.

Establishing that moral intuitions are trustworthy would therefore have a very significant payoff: it would show that we are warranted in relying on moral intuitions in our pursuit of moral knowledge or moral understanding, contrary to the skeptical thesis that moral intuitions can't serve as the basis for any epistemically meritorious doxastic states.[15] A full rebuttal of such skepticism must vindicate all three of the moral intuitionist theses identified above: (α), (β), and (χ). We will devote the bulk of our discussion to the first two, arguing for the trustworthiness of moral *intuitions*; in §6, we show how to extend our argument for these claims to ratify claim (χ), regarding the trustworthiness of intuition-based moral *beliefs*.

3. The Trustworthiness Criterion

We turn now to the project of formulating and defending the key principle in our argument: the Trustworthiness Criterion. The principle employs the notion of a *cognitive practice*. Perhaps the best way to gain a grip on this notion is to have a paradigm example of a cognitive practice before us. Consider what we'll call the *perceptual practice*. At a young age, in the course of undergoing a variety of sensory experiences, most of us are introduced to concepts such as TRIANGLE or TULIP and instructed in their proper application or employment. Along the way, we develop various abilities, such as the ability to discriminate some types of shapes or plants from others: we develop the ability to distinguish triangular things from rectangular things, tulips from roses, and so on. Moreover, we learn how to pay attention to our physical environment so that we can notice triangles and tulips and various properties they might have, and we come to adopt strategies for double-checking our experiences and for evaluating our application of the relevant concepts in various conditions. As a result, most of us have perceptual states of (or as of) triangles and tulips and their properties. It is on the basis of these that we form perceptual beliefs.

This example helps us to state what a cognitive practice is: in the paradigm case, a cognitive practice is a set of activities and events, with social dimensions, into which we are inducted—often as small children—and participation in which, in suitable conditions, yields a certain range of representational states as outputs. These outputs include what we earlier called psychological bases, such as perceptual states, intuitions, and the doxastic states formed thereon (which may serve as psychological bases for further states, and in some cases ground or otherwise promote knowledge or understanding). The activities and events that constitute a cognitive practice typically include: undergoing certain types of experiences, or having certain types of attitudes, and acquiring techniques for prompting, forming, or maintaining them in certain types of circumstances; mastering concepts of a particular kind, after having been instructed in their proper employment by other participants in the practice; learning conventional patterns and strategies of attention, discrimination, and inquiry, as well as means for individual and collective innovations; being introduced to and trained in the use of methods for revising and evaluating its outputs; and employing these concepts, patterns, strategies, means, and methods to form, prompt, maintain, revise, or evaluate the outputs of that practice.[16]

As noted, the outputs of a cognitive practice are a certain range of representational states. By calling them 'outputs' we do not mean to imply they never serve as inputs to any of the activities and events constituting the practice. They may. For example, when a visual experience is hazy or obscure, a second look may be deemed appropriate. In such a case, the perceptual experience is treated as the input to a method for evaluation, whose application may result in another output, a second perceptual experience, against which the first can be checked. Such double-checking is made possible by the methods for evaluation that partly constitute a given practice; such methods are often collaborative, and they can (as we explain below) be more or less critical, and more or less sophisticated.

There are many cognitive practices other than the perceptual one, including the mnemonic practice (whose outputs include memory states), the introspective practice (whose outputs include introspective states), the arithmetic practice (whose outputs include judgments regarding sums), and the scientific practice (whose outputs include experimental observations, explanatory hypotheses, statistical analyses, and more).[17] Of special interest for our purposes is a cognitive practice that we'll call the *moral intuition practice*. Like other cognitive practices, the moral intuition practice is constituted by activities and events, with social dimensions, into which we are inducted—often as small children—and participation in which, in suitable conditions, yields a variety of outputs. The outputs include intuitions concerning the moral status of acts, character traits, policies, and the like. The activities and events include—but are not limited to—having experiences of various sorts (e.g., mnemonic, imaginative, reflective, and dialogical) regarding matters that are assessable morally; mastering moral concepts to which one has been introduced, and being instructed in their proper employment by members of one's community; learning patterns and strategies of attention, discrimination, and inquiry into moral matters; and employing these concepts, patterns, and strategies to form, prompt, maintain, revise, or evaluate the outputs of the moral intuition practice. While the

outputs of this practice may arise in diverse ways, it is worth emphasizing that they are often occasioned by quite ordinary experiences, such as that of describing your promise to provide a meal for a colleague who has just had a child. When speaking with friends who are aware that you'd prefer to get away for the weekend, it may just strike you that you should fulfill your promise.

As we've indicated, cognitive practices have various social elements and are cultivated by social learning: they are social practices.[18] With any cognitive practice, it is possible for participants to engage in that practice in a more or less serious and competent way. Such engagement is a matter of participating in the practice so as to occupy good conditions for the production of the relevant outputs. The good conditions are largely determined by characteristic features of the practice and its outputs, specified by reference to their success. For example, a characteristic feature of the perceptual practice is that an important range of its outputs require normal lighting conditions for their success; consequently, such lighting will be among the good conditions in the case of the perceptual practice. Similarly, given that a characteristic feature of the moral intuition practice is that the success of its outputs requires (*inter alia*) attentiveness to ambiguities and pertinent distinctions, familiarity with relevant non-moral facts, and a capacity for empathy or sympathy, such will be among the good conditions in the case of the moral intuition practice.

Imagine that a particular cognitive practice meets the following conditions. It

(i) is socially well-established: there is broad participation by members of the community over time;[19]
(ii) is deeply entrenched: participation in the practice is to some extent "second nature," where in the limit case recurrent participation is practically inescapable or indispensable;
(iii) makes available sophisticated methods of critically evaluating its outputs (including procedures for detecting and to some extent correcting errors or infelicities among its outputs, standards for legitimate appeal to a community of putative experts, and acknowledged criteria for reliability or infelicity), which are readily employed by those engaging in the practice in a serious and competent way;
(iv) engenders achievement: it enables apparently successful engagement in substantive practical and theoretical projects among those engaging in the practice in a serious and competent way (by, *inter alia*, helping them to make accurate predictions, to enforce some control over their environment, and to pursue socially meaningful ends);[20]
(v) is internally harmonious: for those engaging in the practice in a serious and competent way, it yields outputs that (a) are not massively and systemically inconsistent with each other,[21] and (b) positively cohere (e.g., bear explanatory and inferential relations) with one another.

When a cognitive practice satisfies all five of these conditions, then we will say that it is *in working order*. Suppose that a cognitive practice that is in working order is also such that it

(vi) is externally harmonious: its outputs (a) are not massively and systemically inconsistent with, and (b) in fact positively cohere well with, the outputs of a sufficiently wide range of other cognitive practices in working order.

When a cognitive practice meets all six conditions, we will say that it is *in good working order*.

We are now in a position to present

> **The Trustworthiness Criterion (initial statement)**: If a cognitive practice P is in good working order, then P's outputs are trustworthy for participants in P.

In other words, participants in P are warranted in cognitively relying on P's outputs—to employ them to form, maintain, or revise other of their attitudes in order to increase, improve, or otherwise positively contribute to their stock of knowledge or understanding.

Note that this statement of the Trustworthiness Criterion identifies just a sufficient condition for warranted trust. (So, for all we say here, there may be other criteria whose satisfaction indicates when an agent is warranted in relying on a given state.) We do not regard the criterion itself as stating a necessary condition for trustworthiness. Still, we do think it highly plausible to regard conditions (v.a) and (vi.a) as separately necessary conditions for the trustworthiness of the outputs of a cognitive practice.[22] That a practice generates massive and systemic internal or external inconsistencies is the classic way to show that (in our terminology) the outputs of a cognitive practice are untrustworthy. Agents are not well-positioned to increase their stock of knowledge or understanding when their cognitive practices render verdicts that massively and systemically conflict with each other, or with those of cognitive practices in working order. In what follows, we will take the Trustworthiness Criterion to incorporate this friendly amendment, such that it provides a sufficient condition for warranted cognitive reliance on the outputs of a cognitive practice, while also taking the absence of massive and systemic inconsistency to be a necessary condition for such trustworthiness. In the interest of both explicitness and precision, let us state the amended version of the criterion:

> **The Trustworthiness Criterion**: If a cognitive practice P is in good working order, then P's outputs are trustworthy for participants in P, provided conditions (v.a) and (vi.a) are satisfied.

4. The Status of the Trustworthiness Criterion

We will offer two defenses of the Trustworthiness Criterion. First, we argue that it yields the right results in a wide range of cases. Second, we argue directly that satisfaction of conditions (i) – (vi) is sufficient for a cognitive practice to be such that cognitive reliance on its outputs possesses all ten of the features of warrant enumerated in §2.

4.1 Four examples of cognitive practices

Consider four examples of cognitive practices, beginning with what we earlier called the perceptual practice. Does this cognitive practice satisfy the Trustworthiness Criterion? It would appear so. The perceptual practice is in working order, as it satisfies conditions (i) – (v): it is socially well-established and deeply entrenched, it engenders practical and theoretical achievements, makes available sophisticated methods for critical evaluation, and yields a largely harmonious set of outputs. This is not to deny that the practice sometimes yields incompatible outputs. After all, serious and competent perceivers sometimes disagree about the way things look or what color they are. Still, such disagreement is not massive and systemic, but is limited in its scope and in many cases can be resolved—sometimes straightforwardly—through further participation in the practice (e.g., by looking again, more attentively, in better lighting, etc.). In a similar fashion, while the perceptual practice sometimes yields outputs that are, or seem to be, inconsistent with the outputs of various other cognitive practices (such as the mnemonic practice) that appear also to be in working order, the conflict is neither massive nor systemic. If this is correct, then the perceptual practice satisfies the antecedent of the Trustworthiness Criterion: it is in good working order. It follows, according to the criterion, that its outputs are trustworthy. Independently of the Trustworthiness Criterion, of course, that is the right thing to say about the perceptual practice: we are warranted in relying on perceptual states regarding our immediate environment. The Trustworthiness Criterion delivers the right result.

Consider, as a second example of a cognitive practice, what we'll call the *astrological practice*. This practice yields predictions about the future, such as those that forecast imminent romance, based on information concerning the positions of stars and other celestial bodies. While this practice might satisfy some of the six conditions we have identified, such as being socially well-established over time, it nonetheless fails to satisfy them all. It neither enables us to make accurate predictions nor helps us to control the course of perceivable events. Moreover, it lacks sophisticated methods of critically evaluating (notably, for detecting and correcting errors in) its outputs, which are, moreover, massively and systemically inconsistent with one another, even among serious and competent participants in the astrological practice; they are also massively and systemically inconsistent with the deliverances of other cognitive practices in working order, such as those of physics, astronomy, biology, and psychology. In this way, the astrological practice fails to satisfy the two necessary conditions on trustworthiness, viz., conditions (v.a) and (vi.a). In contrast to the perceptual practice, the astrological practice is a paradigm case of a family of activities whose deliverances do not warrant trust. Once again, the Trustworthiness Criterion delivers the right result.

Take a more difficult case as our third example, one focused on what we'll call the *gustatory practice*. This practice yields outputs that include evaluative gustatory states—those in which one finds food or drink delicious, revolting, mildly appealing, or the like. At first glance, the gustatory practice seems to satisfy the criteria of being socially well-established over time, deeply entrenched, and being endowed with sophisticated methods of critically evaluating its outputs (e.g., it has its own array

of experts). It also engenders various sorts of achievement (e.g., the improvement of recipes, the development of better-tasting fruit and vegetable hybrids). That said, there is considerable intra- and cross-cultural disagreement about evaluative gustatory states among serious and competent participants in the gustatory practice. It is not our aim to determine whether this disagreement is massive and systemic. Even leaving that question unresolved, the gustatory practice illustrates two instructive points about the Trustworthiness Criterion.

The first is that the Trustworthiness Criterion can be satisfied to different degrees or extents (we use these terms interchangeably). A cognitive practice might be on the borderline of being in good working order, or might be in good working order only with respect to a central class of outputs. In such cases, the cognitive practice satisfies the Trustworthiness Criterion *to some extent*. Suppose it turned out that, among serious and competent tasters, there is disagreement regarding evaluative gustatory states (e.g., whether Yirgacheffe coffees taste better than Sidamo coffees) that is on the borderline of being massive and systemic. But suppose, too, that there is a central class of evaluative gustatory states that is not subject to such disagreement (e.g., that ripe peaches and strawberries are tasty). In that case, the criterion would deliver the result that the outputs of the gustatory practice are trustworthy to some extent, albeit not to the same extent as (say) the outputs of the perceptual practice.

Second, while the Trustworthiness Criterion addresses the trustworthiness of the outputs of a given cognitive practice, it is neutral with respect to other features of these outputs. For example, if it turned out that the gustatory practice is in good working order, it wouldn't follow that its outputs represent objective features of reality, at least in any strong sense of "objective." For it might be that had our preferences been somewhat different, different evaluative gustatory states would be correct. The fact that satisfaction of the Trustworthiness Criterion is compatible with such subjectivity is important, insofar as it makes clear that the criterion is not theoretically biased in the following sense: it does not in and of itself stack the deck in favor of non-subjectivist theories, such as moral realism or other forms of moral objectivism.

Turning now to a fourth and final example, consider another less than clear-cut case, which we'll call the *religious practice*. The outputs of the religious practice include experiential states as of God, the gods, or divine activity or presence more generally. Like the gustatory practice, this practice might also satisfy the criteria of being socially well-established over time, deeply entrenched, and endowed with sophisticated methods of critically evaluating its outputs (e.g., it may have its own array of experts who employ rigorous tests to assess its outputs for veracity). It may also engender apparent practical and theoretical successes (e.g., it may foster compassion for the weak and make us less judgmental of others). Still, the religious practice is not a clear-cut case of a cognitive practice in working order, much less good working order, since its outputs appear to be subject to large-scale disagreement even among serious and competent participants. Such conflict appears present even when we focus on central outputs of the religious practice, such as the judgments that express the doctrines essential to the world's most

prominent theistic religions. Arguably, the conflict at least approximates, even if it does not fully realize, massive and systemic inconsistency. Consequently, although the Trustworthiness Criterion allows that the outputs of the religious practice may be on better footing than those of the astrological practice, at a minimum, the criterion implies a serious downgrade in the trustworthiness of the outputs of the religious practice relative to those of (say) the perceptual practice. We take this to be the right result.

That the Trustworthiness Criterion yields the correct verdicts in the four cases we've considered should bolster confidence in its plausibility. We have also noted that the criterion is neutral with respect to many features of the outputs of a given cognitive practice (e.g., whether they are subjective or objective), as it should be. Together, we believe that these considerations provide prima facie reason to think that the Trustworthiness Criterion is a plausible, theoretically unbiased criterion for determining whether the outputs of that practice are trustworthy.

4.2 The epistemic significance of the six conditions

Our second defense of the Trustworthiness Criterion addresses the question of why satisfaction of conditions (i) – (vi) is sufficient for a cognitive practice to be such that its outputs are trustworthy. Our answer is that those outputs have a certain pedigree—namely, that of having been produced by a cognitive practice that is epistemically excellent for a community of inquirers (or, simply 'epistemically excellent').[23] We proceed in two steps. First, we argue that the satisfaction of the six conditions, taken together, contributes to the epistemic excellence of a cognitive practice and its outputs. Second, we review the elements of warrant, elucidated above (in §2), and reveal how the epistemic excellence of a cognitive practice and its outputs explains each of these elements.

We begin with a concessive point: satisfaction of either or both of conditions (i) and (ii) obviously does not entail the presence of epistemic excellence. We allow, further, that it does not even imply the *potential* to possess such excellence. (After all, there are familiar examples of socially well-established and deeply entrenched cognitive practices that wholly lack epistemic excellence and seem incapable of being retooled so as to someday earn such excellence—recall the example of the astrological practice, or the practice of wishful thinking.) By contrast, if a cognitive practice satisfies conditions (iii) – (vi), then it possesses substantial potential for epistemic excellence—potential that is realized if conditions (i) and (ii) are also satisfied. Or so we will argue.[24] We begin by contending that satisfaction of each of conditions (iii) – (vi) by a cognitive practice implies that the practice harbors potential for being epistemically excellent. We then argue that simultaneous satisfaction of conditions (i) and (ii) ensures that this potential is realized.

Condition (iii) states that a cognitive practice must be endowed with sophisticated methods of critical evaluation. The availability of such methods and their ready application by serious and competent participants imply that the outputs of the practice are subject to a kind of "quality control," thus incorporating safeguards against potential threats such as conflict, disorder, and arbitrariness. If such methods were absent, or could not be readily applied by even the most serious and

competent participants, that would be strong reason to be suspicious of the epistemic excellence of that practice. The existence and competent application of these methods does not, of course, guarantee that all is well with a cognitive practice. Rather, the point is that when such methods are available and can be competently applied, a cognitive practice is to that extent the bearer of greater, rather than lesser, potential for epistemic excellence.

The same is true with regard to condition (iv), which requires of a cognitive practice that it facilitate apparent success in theoretical and practical endeavors. What counts as such success is both extremely wide-ranging and contested in some cases. Still, a cognitive practice that consistently frustrates, rather than enables, such success is to that extent epistemically suspect. Here we may recall our example of the astrological practice, which neither enables us to make accurate predictions nor helps us to control the course of perceivable events. Other practices (e.g., one using powerful hallucinogens) may fail to satisfy condition (iv) for different reasons, for example, because they preclude psychologically stable or broad access to their target domain. Once again, the ability of a cognitive practice to facilitate apparent practical and theoretical successes does not, in and of itself, establish that the practice is epistemically excellent. Rather, the point is that to the extent that a practice affords such success, its potential for epistemic excellence is thereby improved.

Turn next to condition (v). The absence of massive and systemic internal inconsistency is no guarantee that all is well with a given cognitive practice. Neither is the presence of substantial coherence relations among its outputs. Still, it would be a bad sign indeed if a cognitive practice were to yield massive and systemically incompatible outputs among its serious and competent participants. It would be better, though not much better, were its outputs *merely* consistent, avoiding conflict with one another while failing to hang together coherently. That would provide strong evidence that the practice is epistemically suspect. By contrast, when the outputs of a cognitive practice are free of massive and systemic inconsistency among its serious and competent participants, while also yielding outputs that stand in important explanatory and inferential relations with one another, then a cognitive practice is to that extent better positioned to realize epistemic excellence.

Now consider condition (vi). If satisfaction of condition (v) implies the potential for epistemic excellence, as we have just argued, then satisfaction of condition (vi) likewise implies the potential for such excellence, since it simply requires consistency and coherence—the key criteria in condition (v)—with the outputs of any cognitive practice that fulfills conditions (i) – (v).

We have argued that the fulfillment of each of the conditions (iii) – (vi) implies that a cognitive practice harbors potential for epistemic excellence. Since the conditions are complementary and not in any way in tension, it is safe to say that they jointly entail that this potential is substantial. We now turn to our claim that when a cognitive practice also fulfills conditions (i) and (ii), this potential is realized: a cognitive practice in good working order is epistemically excellent.

Condition (i) stipulates that a cognitive practice must be socially well-established over time, while condition (ii) requires that a cognitive practice be deeply

entrenched. Together, these conditions entail that the practice enjoys broad participation over a significant period, and that such participation has become second nature to the practice's many participants. If conditions (iii) – (vi) are fulfilled by the practice during this period in the wake of its employment by these participants, then the practice has delivered numerous outputs that have been subjected to considerable quality controls, engendered apparent success, and achieved internal and external harmony. Given that these accomplishments contain the potential for epistemic excellence, as we argued above, it follows that a practice that realizes such accomplishments thereby realizes epistemic excellence.

Having argued that fulfillment of conditions (i) – (vi) makes for a cognitive practice that is epistemically excellent, our next task is to explain why the outputs of such an epistemically excellent cognitive practice are *trustworthy*. Here the question that needs answering concerns the battery of positive elements that partly constitute trustworthiness, or warranted cognitive reliance. Why is cognitive reliance on the outputs of a practice in good working order *favored*, *permitted*, *appropriate*, *blameless*, *reasonable*, and *good*? The answer is straightforward: such reliance exhibits these features because the cognitive practice that generated these outputs is itself epistemically excellent. If we ask, for instance, why we are warranted in relying on our perceptual seemings when seeking to expand our knowledge or gain understanding of our surroundings, it makes good sense to cite the fact that the perceptual practice is socially well-established, deeply entrenched, possessed of sophisticated methods of critically evaluating its outputs, able to facilitate theoretical and practical achievements, and both internally and externally harmonious. That the perceptual practice exhibits these features is what makes it, and its outputs, epistemically excellent. And that they are epistemically excellent explains why reliance on them is favored, permitted, appropriate, blameless, reasonable, and good.

Three additional features of warrant pertain to ways in which it is nonetheless a qualified or modest achievement. First, warrant is *gradable*. This is readily explained by the fact that cognitive practices in good working order fulfill conditions (i) – (vi) to varying extents. Reliance on their outputs is, correlatively, warranted to varying extents. Second, while relying on the outputs of a cognitive practice in good working order is warranted, it is *not necessarily obligatory*. Third, such warrant is *defeasible*. In order to explain why warrant is qualified or modest in these ways, we can point to the fact that a cognitive practice can be epistemically excellent without being complete, invincible, or perfect. Serious and competent participants in exemplary cognitive practices may, for instance, encounter some countervailing evidence or learn that in a given case they are poorly situated to cognize things aright. And while an exemplary cognitive practice may point its participants in a promising direction, it need not thereby impose an obligation on them to follow its guidance.

Finally, warrant is an *epistemic* notion. That is because reliance on these outputs puts one in a better position to gain knowledge and understanding—not (or not merely) a promotion, popularity, or untold riches. When one relies on the outputs of a practice in good working order, one is to that extent epistemically better off, even if there are strong practical considerations that oppose such reliance.

This completes our case for the claim that conditions (i) – (vi) explain why the Trustworthiness Criterion is true. Our argument for this claim entails that these conditions are also sufficient for its truth, since the same considerations that help to explain *why* agents are warranted in relying on the outputs of a cognitive practice in good working order also go to show *that* such reliance is warranted, just as the Trustworthiness Criterion says.

5. The Status of Moral Intuitions

We are now in a position to consider whether the moral intuition practice (as characterized in §3) satisfies the Trustworthiness Criterion. We defend the claim that this practice satisfies all six of the conditions invoked by this criterion to a fairly high degree.

5.1 Conditions (i) – (ii)

The moral intuition practice is clearly socially well-established. Having intuitions about moral matters isn't a recent innovation, but a longstanding element of humanity's way of experiencing and assessing the world. The moral intuition practice is also deeply entrenched—so much so that, for all practical purposes, it is nearly inescapable.[25] As a result, the moral intuition practice easily satisfies the first two conditions.

5.2 Condition (iii)

Condition (iii) requires that the moral intuition practice incorporate sophisticated methods of critically evaluating its outputs—moral intuitions. We believe that it does, although establishing the success of the moral intuition practice on this front will require that we be specific about the methods available to competent participants in the practice. Here we describe five such methods, all of which have close analogues in other cognitive practices.

A first method is the straightforward one of testing for *consistency*, establishing *corroboration*, and pursuing *confirmation*.[26] In many cases, it is easy to check whether moral intuitions are inconsistent, whether the moral intuitions of one individual are corroborated by the moral intuitions of other individuals, or whether moral intuitions are confirmed by experience or reflection. By engaging in moral dialogue and deliberation, we can ascertain whether our moral intuitions are consistent. If they are not, we can attempt to resolve these inconsistencies by revisiting cases and clarifying our understanding of them, reflecting on them anew with the aim of finding whether our intuitions shift.[27] We can seek corroboration by consulting the views of other participants in the moral intuition practice whom we trust. And we can pursue confirmation by checking whether our moral intuitions imply absurdities or repugnancies. Just as we can assess solipsism by noting its unpalatable consequences, we can also assess a set of moral intuitions by noting their outrageous implications (e.g., that we are morally required to kill an unsuspecting person in order to save five patients requiring organ transplants).[28] Such assessment constitutes a familiar kind of critical reflection, which can be undertaken at both

individual and community-wide levels, and whose function is to root out intuitions that do not merit belief on their basis.

To the "three Cs" of consistency, corroboration, and confirmation, we should add a fourth, which is establishing that our moral intuitions are *coherent*—that is, that they are somehow mutually supporting.[29] When we systemize a set of moral intuitions in such a way that they cohere with other moral and non-moral intuitions, we thereby help to bring these intuitions into what John Rawls dubbed "wide reflective equilibrium."[30] As Rawls noted in his discussions of this notion, the fact that a set of moral states—Rawls had his eye not on intuitions but "considered judgments"—are situated in wide reflective equilibrium is no guarantee that all is well with them. In particular, it does not guarantee that they are trustworthy. But our task here is not to locate guarantees; it is, rather, to identify widely acknowledged criteria for epistemic success and improvement. And whatever limitations wide reflective equilibrium may possess, it nevertheless has the virtue of operating according to a widely acknowledged criterion for this, namely, coherence. The fact that the moral intuition practice readily incorporates this criterion reflects well on its critical resources.

A second method of critical evaluation consists in certifying that participants in the moral intuition practice are in good conditions for the production of the relevant outputs, moral intuitions, and taking steps to remedy flaws in those conditions when possible. Such conditions include:

- being sensitive to the clarity and strength of participants' intuitions;
- being alert to potentially significant (e.g., diverting) features of the target example, principle, or question;
- being careful about the content (e.g., metalinguistic or comparative) possessed by participants' intuitions;
- marking relevant distinctions (e.g., epistemic possibility versus metaphysical possibility);
- disambiguating (scope ambiguity, lexical ambiguity, etc.);[31]
- noting context-sensitivity;
- flagging pragmatic effects;
- securing familiarity with relevant non-moral facts;
- remaining attentive throughout the process of reflection;
- manifesting awareness of participants' affective situation (e.g., as a potential source of either distortion or enhancement);
- fostering and properly exercising the capacity for empathy or sympathy;
- considering a variety of examples;
- being wary of esoteric or outlandish scenarios (e.g., seeking replication when considering such scenarios); and
- avoiding alcohol and other drugs.

Although this list is incomplete, it brings to light a wide range of factors that bear directly on the success of moral intuitions. These factors are manipulable, at least to some extent, by serious and competent participants in the moral intuition practice. Being part of a community whose members are also active participants in the practice introduces opportunities for collaboration, division of critical labor,

systems of checks and balances, and social pressures to execute the above factors well rather than poorly. In effect, participants in the moral intuition practice are in a position to individually and collectively structure and facilitate their efforts to use moral intuitions responsibly, with the aim of increasing, improving, or otherwise positively contributing to their stock of moral knowledge or moral understanding.

A third method of critical evaluation consists in contrast and comparison with *exemplars*, who are widely acknowledged by the community to be models in a particular domain. Think, for example, of how we make progress in the domains of musical and athletic performances: we look to those who do them well. The same is true in the moral domain. To make progress in the moral life, we look to moral exemplars, whose intuitions, judgments, and behavior set the bar. Having done so, we can then detect errors in our own moral intuitions and judgments, and identify ways in which they can be improved. This is another place where the social dimensions of cognitive practices are important: simply put, we learn from others. For example, we can notice and reflect on the concepts that moral exemplars employ, as well as their patterns and strategies of attention, discrimination, and inquiry. In some cases, we might realize that we need to employ, or strive to master, a wider range of moral concepts, which enable us to make distinctions that are relevant to appropriately responding to situations that we have previously failed to classify properly and in which we have, as a result, morally faltered. In other cases, we might determine that we need to develop better habits of attention, focusing more carefully on aspects of situations that we may have otherwise overlooked. In still other cases, we might recognize that we need to develop better habits of inquiry, such as stepping back from fraught or explosive situations to gain increased perspective or distance. In these and other ways, contrast and comparison with moral exemplars serves as a potent source of critical evaluation, especially when embedded in a rich social context with conventions of mentorship and the like.

A fourth method of critical evaluation consists in what J. S. Mill called *experiments of living*, which involve trying out several different ways of being and acting in order to discern what it is like, for oneself and others, to be and act in those ways. This could take the form of an apprenticeship or internship, provisionally undertaken with the goal of determining whether a given trade or career is the right fit. In the case of morality, as Mill observed, the point may be to discern whether and how certain patterns of conduct are amenable or inimical to the well-being of oneself and others. For example, some have taken it upon themselves to "go undercover" and live in historically persecuted or neglected groups; others have taken to the woods, desert, or grove (think: Thoreau, Jesus, Buddha), permanently or temporarily abandoning familiar routines and social settings. Such experiences often result in profound alterations and improvements, particularly when undertaken as a community-wide experiment. There are of course limits on what experiments may be done here: the infamous Stanford prison experiments are no longer permitted, and for good reason. Fortunately, when real-life enactments would be deeply morally problematic, thought experiments may function as comparatively safer ways to explore the implications of various ways of being and acting. Engaging in experiments in living, whether hypothetical or actual, individually or collectively,

is a fertile and versatile method for critical evaluation of the outputs of the moral intuition practice.

The fifth and final method of critical evaluation that we will mention consists in *reflection on concepts*. To cite this method is not to commit to substantive theses about conceptual analysis, but to summon the observation, familiar from pedagogy, that improvement in a given domain is often facilitated by improved grasp of the concepts involved in that domain. To give a simple example, teachers of mathematics convey the subject's basic concepts—EQUALITY, ADDITION, MULTIPLICATION, PRIMALITY, POSITIVITY, NEGATIVITY, and so forth—and then encourage their students to study and master these concepts by thinking carefully about, and with, them (where such thinking includes, but is not limited to, reflecting on these concepts' conditions of application, their purposes, their relations to other concepts, specific instances in which they are either applicable or inapplicable, and so on). In the normal case, the result is a range of more accurate or felicitous mathematical intuitions and judgments by the group.

The point we are making is not limited to pedagogical contexts. A Chicagoan curious about her city's architecture, for example, may reflect seriously on the concept FORM and realize—perhaps prompted by dialogue with her tour guide from the city's architectural foundation—that it encompasses the space surrounding a structure just as much as it does the positive space occupied by the structure itself. Subsequently, she may find her architectural intuitions and judgments shifting accordingly, so as to render them more accurate or felicitous. Participants in the moral intuition practice may engage in analogous reflection: moral agents may reflect seriously on the concept FAIRNESS and realize that it encompasses both descriptive and evaluative components. Subsequently, they may find their moral intuitions and judgments shifting accordingly, so as to render them more accurate or felicitous.[32] As these examples indicate, the method of reflection on concepts is not restricted to the classroom or technical disciplines, but is available to mature reflective agents engaged in diverse theoretical and practical projects seeking to critically evaluate their intuitions with the goal of better executing those projects.

As we have noted, these five methods for critically evaluating the outputs of the moral intuition practice are not unique to this practice; they have close counterparts in many other cognitive practices.[33] That said, we recognize the existence of some important differences between the moral intuition practice and other cognitive practices. To confirm whether another agent's vision is accurate, you might look at the very same object as that agent from a different angle. No such method is available in the moral intuition practice. You could not ascertain the accuracy of another's intuition that an inequitable distribution of goods can be fair by attempting to look at fairness, let alone from a different angle. To ascertain whether a given moral intuition is accurate, you would have to engage in a different activity, such as contemplating whether it coheres with other principles you accept, imagining hypothetical situations in which there is an inequitable distribution of goods, or reflecting carefully on the concept FAIRNESS with the aim of discerning whether there are reasons for holding that inequitable distributions of goods could satisfy it. To which it should be added that no such set of methods has any hope

of success in the domain of perception. There is no way to determine whether your visual perceptions are accurate simply by contemplating whether they cohere with various principles you accept, imagining hypothetical situations in which things are one way rather than another, or reflecting on concepts such as OBJECT or COLOR. In this respect, at least, the moral intuition practice may have resources for critical evaluation that are unavailable to various other cognitive practices, such as the perceptual practice.

One implication of these remarks is that the methods of critical evaluation we have canvassed expose the limitations of the view that productive normative inquiry is limited to pure procedural reasoning. The five methods we've enumerated (and very likely others) offer means of expanding our epistemic horizons in ways that extend well beyond the confines of such reasoning. Importantly, this shows that intuition-based moral inquiry need not be unduly conservative; moral intuitionists are not committed to a picture of moral reasoning that takes moral intuitions as incorrigible (or only minimally corrigible) starting points for reflection, somehow beholden to commonsense, ordinary language, or current orthodoxy.[34]

Were our project to engage in a fine-grained comparison of the moral intuition practice to other cognitive practices, there would be much more to say about these matters. But that is not our project. It is, instead, to argue for the claim that the moral intuition practice satisfies the antecedent of the Trustworthiness Criterion, including condition (iii). And that claim is compatible with (*inter alia*) there being interesting points of both similarity and difference between the moral intuition practice and other cognitive practices that satisfy that antecedent. We would only counsel that, given the variety of methods of critical evaluation available to different cognitive practices, we be careful not to identify some particular set of methods of critical evaluation—such as confirmation by simple sense perception or repeatable empirical testing—as the *sine qua non* of a practice in good working order. When it comes to methods of critical evaluation that bear on the trustworthiness of the outputs of cognitive practices, pluralism—per condition (iii)—must be our watchword.

5.3 Condition (iv)

Condition (iv) requires a cognitive practice to enable apparently successful engagement in substantive practical and theoretical projects among serious and competent practitioners. The moral intuition practice does this. A host of practical benefits accrue to serious and competent moral agents who rely on their moral intuitions for guidance. Though the moral intuitions of serious and competent agents will not help fix a flat tire or repair a toaster, they will help one navigate specific situations in which there is insufficient time or information for explicit moral reasoning. They will also enable the development of a sound moral character, which, in fortunate circumstances, will in turn help (*inter alia*) to sustain loving relationships, to earn the trust of others, and to cultivate the capacity for sympathy and empathy.[35]

The moral intuition practice also enables us to engage in apparently successful theoretical inquiry. For example, it is difficult to imagine making headway in normative ethics without relying on the moral intuitions of serious and competent

moral agents. Indeed, according to W.D. Ross, such intuitions constitute or deliver "the data of ethics."[36] Though we ourselves find this view congenial, one needn't go quite so far as this in order to accept the claim that a great deal of apparently fruitful moral inquiry is helped along by reliance on our moral intuitions. For even those who flatly reject the evidential status of moral intuitions in the "context of justification" can and should allow that such intuitions play an important role in the "context of discovery," in this way potentially serving as engines of theoretical progress.

5.4 Condition (v)

We turn now to condition (v) in the Trustworthiness Criterion, which invokes internal harmony. There are two questions to address here. The first is whether the outputs of the moral intuition practice are free from massive and systemic internal inconsistency, as required by (v.a); the second is whether these outputs exhibit robust coherence relations, as required by (v.b). It is sometimes said that moral disagreement—endorsement by different agents of opposing moral claims—is pervasive, which might be taken to imply that neither of these requirements is met. Here we argue that such an inference would be illegitimate, and that the moral intuition practice satisfies both requirements. We will focus on (v.a), making several points that (as we'll explain) are also relevant to (v.b).

Let us begin by recognizing that there is substantial disagreement about many moral matters, such as whether it is permissible to eat animals or to perform elective abortions, whether we have duties to the poor or to oppose policies that would impose life-sentences on those who commit non-violent crimes, and whether lying or stealing and the imposition of various gender and sex norms is morally justified. We do not wish to downplay the subtlety and interest of such conflicts, and we intend to say more about some of them below. But when it comes to the assessment of (v.a), there are two points to emphasize.

The first is that despite widespread disagreement on a wide variety of moral issues, among serious and competent participants in the moral intuition practice, one finds nearly complete agreement among moral intuitions[37] that are expressed by moral platitudes such as:

> It is right to protect one's children from lethal danger.
> It is wrong to recreationally slaughter fellow persons.
> An agent has a right against others that they not torture her for mere entertainment.
> One ought to act justly, if acting justly is costless.

Any genuine intuitional conflict about such platitudes and their direct applications, such as one might find by surveying the Charles Mansons of the world, is either minimal (hence not widespread) or traceable to one or another flaw or idiosyncrasy (hence not systemic).[38]

The second point is that there are a great many other moral intuitions that are overwhelmingly congruent, although they are not obviously platitudinous. For instance, the intuition that it is impermissible to harvest the organs of one healthy

person, chosen randomly, to distribute to five sickly patients is widely shared. We could say the same for cases in which agents, upon being made aware of the psychological effects of waterboarding, or learning the details of a botched execution, have the intuitions that these actions are morally problematic. These join countless other cases in which our moral intuitions are largely uniform across place, time, class, gender, and education. While there are also many moral questions that elicit differing moral intuitions, the extent of intuitional convergence is easily overlooked, as it is usually the difficult, contested cases that gain the headlines and attract the attention of theorists in this area.

We've observed that there is, in fact, little conflict in intuition regarding the moral platitudes and many other moral matters. Is the remaining moral disagreement sufficiently massive or systemic to threaten the moral intuition practice's standing? We contend that we presently lack reason to think that it is. For while there is undoubtedly widespread disagreement on many moral matters, only some of it reflects conflict among moral intuitions. A significant portion can be traced instead to:

(a) different, perhaps incompatible factual beliefs;
(b) different, perhaps incompatible beliefs about the principles governing the relative weight (epistemic, prudential, or moral) of various sorts of considerations; or
(c) different, perhaps incompatible ancillary moral beliefs or judgments.[39]

The beliefs in (a) are never fully based on moral intuitions, and the beliefs and judgments in (b) and (c) are only sometimes fully based on moral intuitions. Indeed, in a great many cases, the mental states in (b) and (c) will be the psychological results of complex patterns of association, reasoning, or non-conscious processing whose inputs are not solely, or even predominantly, moral intuitions.

To illustrate, consider the question of whether it is permissible (for agents situated as we are, with our typical purchasing habits) to eat animals. Here, the moral disagreements that exist do not seem solely, or even predominantly, traceable to a conflict among moral intuitions. After all, there is very little clash of intuitions about, for example, whether there is something morally amiss about current factory farming practices. Instead, the moral disagreement on this topic appears substantially due to other factors. Some is owing to the non-intuitional, non-moral differences referred to in (a): for instance, some people are largely *ignorant* of how poorly factory-farmed animals are treated, and of the environmental effects of factory farming—obviously, not the fault of moral intuitions. Among those who are equally well-acquainted with such facts, much of the persisting disagreement about the permissibility of eating animals stems from either (b) or (c). Regarding the former, agents may draw upon conflicting *non-moral beliefs*, regarding the weight of the various prudential and epistemic considerations in play, to reach incompatible conclusions—once again, not the fault of moral intuitions.[40] Finally, there is the disagreement that arises from the non-intuitional moral differences referred to in (c): for example, agents possess disparate background moral views—

arising not from moral intuitions but from *desires*, *emotions*, *hypotheses*, *complex reasoning*, and sundry other non-intuitional sources—about the comparative importance of human interests over those of non-human animals.

There is nothing unique about the case we've just considered. In many prominent cases of moral disagreement, we see similar dynamics at work, and the disagreements are best explained by reference to the beliefs described in (a) – (c), rather than conflict among moral intuitions.[41] Given this, together with the fact that there is a lot of convergence on the moral platitudes and their direct applications, as well as a great many other cases, we cannot legitimately infer massive and systemic inconsistency among moral intuitions from the observation that there is a lot of actual moral disagreement.[42]

Some philosophers will maintain that the real problem lies elsewhere, in deeply entrenched disagreement among moral *theorists* (including theoretically-inclined laypeople) about, for instance, what makes actions right or wrong, or what grounds the instantiation of a specific instance of rightness or wrongness (or the factors on which such instantiations depend), or whether there are any right or wrong actions in the first place.[43] However, such disagreement has fairly limited implications. One reason is that much, if not most, of it can be traced to conflict among the sorts of beliefs described in (a) – (c). It is a familiar point that disagreement between realists and error theorists, for example, is due not to divergent moral intuitions, but to opposing ancillary non-moral views (e.g., beliefs not fully based on intuition regarding the status of metaphysical naturalism, empiricism, the relation between God and morality, or the criteria of action explanation). We do not deny that moral intuitional conflicts play some role in theoretical debates like this one. But that gives us no basis for extrapolating to the conclusion that conflict among moral intuitions is *massive and systemic*.[44]

We take the foregoing points to justify our contention that the moral intuition practice satisfies (v.a), at least to some extent. Several of these points can be applied, *mutatis mutandis*, to show, per (v.b), that the moral intuition practice generates outputs that cohere well with one another. Though there is doubtless a greater degree of consistency than of coherence among the moral intuitions of participants in the moral intuition practice, the same could be said about the outputs of any cognitive practice in good working order, as coherence is more demanding than mere consistency. Further, the moral intuitions that are the object of widespread convergence do exhibit substantial degrees of coherence. Fully substantiating this claim is a project for another day, but consider this important piece of evidence: most intuitions can be gathered under a set of increasingly general moral principles that culminate in requirements to display respect and to protect or promote welfare. In these ways, there is a lot of mutual support among the moral intuitions whose contents are concrete verdicts or moral principles at varying levels of generality.

The emerging fit, of course, is far from perfect: our actual moral intuitions are not *fully* coherent—but then neither are the actual outputs of any other actual practice in good working order. In our estimation, moral intuitions are probably less coherent than the outputs of the cognitive practices in the very best working order, such as the perceptual practice or the arithmetic practice. But it is no

part of our brief to argue for full parity between the moral intuition practice and these other exemplary cognitive practices. We claim only that there are a wide range of moral intuitions, most notably (but not restricted to) those enjoyed by moral agents participating in the moral intuition practice in a serious and competent way, that are sufficiently internally coherent so as to imply that (v.b) is fulfilled by this practice to at least a moderate degree.

5.5 Condition (vi)

The argument thus far supports the conclusion that the moral intuition practice is in working order, since the practice satisfies conditions (i) – (v). To establish that the moral intuition practice is in *good* working order we need to furnish reasons that its outputs (vi.a) are not massively and systemically inconsistent with, and (vi.b) in fact cohere well with, the outputs of a sufficiently wide range of other cognitive practices in working order. To fix attention, consider the wide range of cognitive practices in working order that involve familiar, widely respected methods for prompting and evaluating its outputs—for example, in addition to the perceptual practice, the mnemonic practice, and the introspective practice, there are those cognitive practices associated with the disciplines of logic, mathematics, medicine, electrical engineering, geology, botany, horticulture, oceanography, immunology, metallurgy, chemistry, psychology, and astronomy. Call such practices *paradigm practices*. We submit that there is no reason to believe that the outputs of the moral intuition practice are massively and systemically inconsistent with the outputs of any plausible candidates for paradigm practices.[45]

Consider once again moral intuitions whose contents are moral platitudes. Moral platitudes are not massively and systemically incompatible with the deliverances of the perceptual practice. Nor are they massively and systemically incompatible with the deliverances of the mnemonic and introspective practices. The same is true of other paradigm practices. One reason that they are compatible is that moral intuitions (regarding the platitudes and other moral matters) are morally prescriptive or evaluative, unlike the outputs of these paradigm practices. Indeed, none of those latter outputs constitutes a moral prescription or a moral evaluation; as a result, there are *no* direct conflicts, much less any conflicts that are massive and systemic, between the outputs of paradigm practices and those of the moral intuition practice.[46]

The feature that secures the satisfaction of condition (vi.a) may seem to threaten the possibility of fulfilling (vi.b): since paradigm practices issue no morally prescriptive or evaluative outputs, it may seem puzzling how moral intuitions could cohere well with the outputs of these other cognitive practices. But this puzzle is easily solved. For all concrete moral intuitions (i.e., those that do not represent moral principles but rather assessments of particular cases) can cohere more or less well with the outputs of paradigm practices, by virtue of the empirical presuppositions that underlie those moral intuitions. When it strikes moral agents that an action is immoral, or a person's character evil, or a situation morally regrettable, such a striking implies certain propositions—many of which involve empirical assumptions—about the ways things are non-morally (e.g., that factory farming

causes animals unnecessary suffering; that a powerful man is serially assaulting women; that human actions are adversely affecting the climate). When moral intuitions are formed in good conditions, by agents engaging in the practice in a serious and competent way, there is every reason to suppose that those intuitions cohere well with the outputs of diverse paradigm practices, whose contents include those propositions.

We can say something similar about moral intuitions concerning moral principles (including the moral platitudes). For example, when moral agents intuit that it is wrong to torture another simply because he has inconvenienced them, their intuition implies certain non-moral propositions about torture—propositions that may focus on such things as our physical vulnerabilities, psychological susceptibilities, and the likely erosions of bureaucratic safeguards in a regime that permits torture. Such propositions regularly cohere with the outputs of other cognitive practices in working order. Indeed, when serious and competent moral agents are reflecting in good conditions, they will often be well-placed to have moral intuitions whose non-moral implications fully cohere with the outputs of many such practices.

6. From Moral Intuitions to Moral Beliefs

The moral intuition practice provides a center and a station from which moral agents, and the community to which they belong, may engage in epistemically responsible moral inquiry. We have argued that such agents are warranted in relying on moral intuitions to expand their stock of knowledge or understanding regarding the moral status of acts, character traits, policies, and the like. While our success on this front is tantamount to vindicating two core tenets—theses (α) and (β) in §1—of our brand of moral intuitionism, which is a traditional view in moral epistemology, the argument itself is not traditional. For it does not privilege features of any mental states or events *per se* (such as self-evidence or a particular type of phenomenology), but rather assigns a central role to cognitive practices in good working order—socially established practices of inquiry and evaluation that generate a variety of epistemically excellent outputs for participants in the practice. We have argued that, and explained why it is the case that, when a cognitive practice satisfies the antecedent of the Trustworthiness Criterion, its outputs are trustworthy. We have also argued that the moral intuition practice manages this feat. It follows that moral intuitions are trustworthy, in virtue of being the outputs of a cognitive practice that is in good working order.

Although our argument has focused on moral intuitions, it also bears on the epistemic status of moral beliefs that are based on such intuitions. To appreciate this connection, consider the following principle: *if a given mental state or event is trustworthy, then, in normal circumstances, so too are the doxastic states to which it gives rise*.[47] Given this principle, it follows from our argument that agents are warranted in relying on those moral beliefs that are formed on the basis of moral intuitions to expand their stock of knowledge or understanding: per the moral intuitionist thesis (γ), intuition-based moral beliefs, too, are trustworthy.[48]

Notes

[1] Rawls (1955, 27). We enter the caveats that Rawls's topic differs in several ways from our own, and we do not intend to endorse the whole of his view.

[2] These remarks are not intended to offer a *theory* of the nature of intuition. While we endorse the theory presented in Bengson (2015a and 2015b), we will not presuppose its details here.

[3] Hereafter we use 'belief,' although we take our claims to hold for judgments and many other states and events.

[4] This caveat is explained when defending (χ) in §6 (see note 47).

[5] See Ross (1930) and Audi (1997, 1999, 2004). Stratton-Lake (2016) provides a fine analysis of why moral intuitionists should prefer a non-doxastic conception of intuition to the more traditional doxastic one.

[6] The attribution is controversial. As Hurka (2015, 112) observes, "many in the [intuitionist] school emphasized the difficulty of achieving moral knowledge and the consequent fallibility of our intuitions."

[7] See, e.g., Audi (2015).

[8] See, e.g., Moore (1903, x and §86).

[9] Some philosophers, such as Baier (1986), Holton (1994), and Jones (1996), have focused on *trusting a person* (to do something); our discussion, like the discussions of trust by Wright (2004, §V) and Elgin (2008), is not restricted to interpersonal trust.

[10] We return to this point below. Throughout we focus on cognitive reliance; for brevity, we sometimes elide 'cognitive.'

[11] As indicated just below, the relevant value is epistemic. For a recent collection of essays on epistemic value, discussing the centrality of value to epistemology, see Haddock, Millar, and Pritchard (2009). We remain neutral on the correct theory of epistemic value here (e.g., monistic versus pluralistic accounts).

[12] Warrant—hence, by implication, trustworthiness—is distinct from several other epistemic merits, such as possessing knowledge, being rational, or enjoying inferential support (at least on certain common construals of these notions). Unlike knowledge, warrant does not require non-accidentally true belief. Unlike rationality, warrant does not consist, even partially, in consistency or coherence (though it may of course be undermined by certain types of inconsistency or incoherence). And unlike inferential support, the notion of warrant does not, in and of itself, require that one be warranted in any of one's antecedent attitudes. We remain neutral on the relation between warrant and justification, whose proper analysis (or even pretheoretical characterization) is contentious. What about various alethic properties, such as reliability? In one respect, warrant is *weaker*: it does not guarantee truth-conduciveness. But in other respects, warrant is *stronger*: for example, it entails the features enumerated above, which alethic properties do not (as shown by Bonjour's (1980) clairvoyance case and Zagzebski's (2003) espresso machine case, among many others).

[13] This property, which we have termed 'warrant,' bears important similarities to the property designated by Wright's (2004) use of the same term, although his silence on several of the features we've listed and his emphasis on warrant as what makes a thinker "beyond rational reproach" (*ibid.*, 175) suggests that his property is not identical to ours in all respects. Our use of 'warrant' differs markedly from others' use of the same term to designate what makes the difference between true belief (or justified true belief) and knowledge.

[14] That inquiry is an activity with this, rather than mere truth or reliable belief, as its aim is a common point in recent literature. See especially Pritchard (2016) and Friedman (forthcoming).

[15] Such skepticism is the natural conclusion of critics' allegations that it would be mysterious if intuitions had epistemic power; that intuitions are mere feelings or hunches, which are epistemically indigent; that intuitions could have epistemic power only if foundationalism were true, though it is not; that intuitions are unacceptably conservative; that intuitions generate disagreements that are not in principle rationally resolvable; that intuitions cannot be "calibrated." (See, e.g., Wittgenstein 1937 (417–19), Ayer 1946 (73), Singer (1974, 2005), Mackie (1977, 39-40), Brandt (1979, 16-23), McDowell (1985, 111), Brink (1989, ch. 5), Hintikka 1999 (130–3), Kitcher 2000 (75ff), Haidt (2001), Kagan (2001), Weinberg, Nichols and Stich (2001), Sinnott-Armstrong (2002, 2007, 2008), Weinberg (2007), Williamson (2007, ch. 7), Gibbard (2008, 20-1), McGrath (2008), Della Rocca (2011), Fanselow (2011), and Greene

(2013).) As should be clear, we are not here seeking to show that the epistemic status of moral intuitions is *never* defeated. Critics do not regularly distinguish worries about defeat from skepticism proper, though many of the critics just cited are for various reasons committed to the latter. We believe that an adequate reply to skepticism—*viz*., our argument below for the trustworthiness of moral intuitions—introduces resources that enable intuitionists to develop a compelling reply to concerns about defeat as well. We sketch this reply in note 33, acknowledging that its full development lies beyond the scope of the present paper (see Bengson, Cuneo and Shafer-Landau ms.).

[16] Those familiar with Alston's (1991, ch. 4) epistemology will note that cognitive practices resemble what he calls "doxastic practices." However, there are several differences: for example, whereas Alston's doxastic practices yield only beliefs as outputs, some cognitive practices yield non-doxastic outputs such as perceptual experiences or intuitions. Alston's larger aim is also importantly different from ours. We will argue that cognitive practices that are in what we call 'good working order' yield outputs that are *trustworthy* (in the specific sense defined in §2). This is not a claim that Alston defends; in contrast, his central thesis is that doxastic practices whose outputs are subject to what he calls "assessment procedures" (158) yield beliefs that there is prima facie *practical* reason to *take to be reliable*.

[17] We remain neutral on how to individuate cognitive practices. For present purposes, we allow a host of diverse cognitive practices, which perhaps overlap or stand in various compositional or hierarchical relations.

[18] Inspired by MacIntyre (1984, ch. 14), Wolterstorff (2015, ch. 8) maintains that in a social practice, participants (a) "hand on the know-how" of engaging in the practice to others, (b) sometimes act jointly, and (c) are aware of others' participation. Thompson (2008, Part III) adds that in a social practice (d) convergence of action is not accidental, in the sense that agents' "coinciding actions exhibit a common source or account." Our characterization emphasizes (a-b) and readily accommodates (c-d). Railton's (2014, §XIV) discussion of moral intuition also gives a role to social factors. But whereas Railton focuses on intuitions' reliability and treats social conditions as playing a merely causal role in their development, we emphasize trustworthiness and argue below that social factors figure into the very grounds of their trustworthiness.

[19] We intend to construe the notion of community liberally; it might include (e.g.) various subcultures or suitably integrated cadres of like-minded individuals. This first condition, like several of the others, has been discussed within the field of social epistemology (e.g., Longino 1990; Solomon 2007; Wagenknecht 2016; Elgin 2017, chs. 5-7). However, these conditions' connection to what we are calling 'trustworthiness' (warranted cognitive reliance) has not to our knowledge been explicitly considered.

[20] We include the 'apparently' qualifier in order to recognize that we cannot—and need not—*guarantee* that exemplary cognitive practices, such as the perceptual practice, are in fact facilitating successful practical and theoretical projects. (We might, after all, be brains in a vat or subject to some other radical skeptical threat.)

[21] This is not identical to the condition that conflict among such outputs always be rationally resolvable. Nor is rationally unresolvable conflict evidence of massive and systemic inconsistency. Rational resolvability is orthogonal to condition (v.a); if the arguments in §4 are sound, it is not necessary for trustworthiness.

[22] Weinberg (2007, §§2-3) can be read as contending that condition (iii)—which is of a piece with a criterion he labels 'hopefulness'—is a necessary condition for trustworthiness. Although it is compatible with our arguments, we do not find this contention compelling and will not assume its truth in what follows.

[23] We use 'epistemic excellence' to designate a determinable epistemic quality, with positive valence, above a sufficiently high threshold. We remain neutral on the correct theory of this quality, but assume that it is intuitive, and that paradigms are familiar (e.g., the scientific method, canonical patterns of logical reasoning, mathematical induction).

[24] A less cautious stance is that a cognitive practice that satisfies only conditions (iii) – (vi) is epistemically excellent for a community of inquirers. If so, this would only strengthen our case in what follows.

[25] These points are compatible with the idea that the moral intuition practice as we know it today (replete with contemporary moral concepts, patterns and strategies of attention, etc.) represents a substantial development of the practice as it was initially employed. These points are also typically

recognized by metaethicists and moral psychologists. For example, nearly all prominent antirealist positions are fully compatible with, and indeed explicitly endorse, the idea that the moral intuition practice is pervasive and has a rich history.

[26] Bealer (1992) labels these the "three C's," which he emphasizes in another context. We understand each "C" to include its negative corollary: one can check for inconsistency, discorroboration, and disconfirmation, respectively.

[27] See, e.g., Bealer (2008, §7) on the role of "local misunderstanding" in the handling of intuitional error.

[28] It is also possible to pursue confirmation of moral intuitions by checking their *non-moral* implications. Consider, for example, the outlandish claim that it is *metaphysically impossible* for there to be an infinite number of happy and sad people. This absurdity is arguably implied by (the complex set of) moral intuitions that support aggregation consequentialism, which are thereby rendered suspect. *Cp*. Bostrom (2011).

[29] See, e.g., Ewing's (1941, 34 ff.) discussion of "coherence tests." See also Bonjour (1998, 117–18).

[30] Rawls (1971). DePaul (1993) offers a comprehensive defense of wide reflective equilibrium in the case of morality. See also, among others, Brink (1989, §§5.8-9), Sayre-McCord (1996, 141 ff.), and DePaul (2011).

[31] Zwicky and Sadock (1979) provide a useful guide.

[32] This type of moral reflection—reflection on moral concepts in the service of more accurate or felicitous moral intuitions and judgments—has been emphasized by Murdoch (1970) and Nussbaum (1990), among others.

[33] We observed in §3 that the methods of evaluation that partly constitute a cognitive practice may be applied to the practice's initial outputs, generating further outputs that are improved in various respects. The latter outputs may be used to evade or overcome defeaters afflicting the initial ones, when the methods are critical and sophisticated to a very high degree. That the moral intuition practice incorporates the five methods of critical evaluation we've canvassed reveals that the practice meets this bar. We believe that this is an important step on the way to an intellectually responsible reaction to the prospect of defeaters for moral intuitions.

[34] *Cf*. Singer (1973), Brandt (1990), and Della Rocca (2011).

[35] Such successes do not entail epistemic excellence, but they also are not irrelevant to it. For example, the first and second further opportunities for corroboration, while the third facilitates the occupation of good conditions.

[36] Ross (1930, 41).

[37] We focus on intuitions here and below because, as explained in §1, not all moral beliefs are based on moral intuitions; consequently, disagreement among such beliefs does not reflect conflict among moral intuitions.

[38] After all, Charles Manson and his ilk are not serious and competent participants in the moral intuition practice. The point we are making here marks a contrast with paradigmatically problematic cognitive practices such as extra-sensory perception ("ESP") and aura reading, which fail to generate a set of outputs that are immune to conflict: to our knowledge, there is no analogue in these practices to the moral platitudes and their direct applications.

[39] Recall that, throughout, our use of 'belief' is not meant to exclude judgments and other mental states. (See note 3.)

[40] For example, some hold that "Morgan's Canon," which does not obviously imply that many animals experience tremendous suffering, has substantial epistemic weight, while others privilege "Newton's Rule," which implies that they do (see Klein 2017). In this case, disagreement is due not to moral intuitions but divergent *epistemic attitudes*.

[41] The full explanation may be complex: for example, some moral disagreements may be due to methodological leanings that stem from privileging the simplicity of theories or a preference for metaphysical naturalism, which are not plausibly viewed as products of specifically moral intuitions. *Cp*. the explanatory options described in Bengson (2013, §5).

[42] *Cp*. Bealer's (1992, 110-14) analysis of intuitional conflict. He goes further: "In the case of intuition, no one yet knows how far the elimination of apparent conflict goes. At this point we cannot rule out with certainty that it does not go all the way" (131, n.18).

[43] *Cp*. Leiter (2014). *Cf*. Parfit (2011) on convergence among moral theorists.

[44] *Cp*. Williamson (2007, 191). We can bolster the points in the main text by considering the case of logical disagreement. There is substantial disagreement among theorists of logic (including theoretically-inclined laypeople) regarding the accessibility relation and the nature of implication. But it would be incredible to infer that logical intuitions are in general massively and systemically in conflict on the basis of such disagreement. The same is true of disagreement among moral theorists.

[45] Some might concede that moral intuitions satisfy condition (vi) but argue that this is explained by epistemically invidious factors, such as those cited in so-called *adaptive link* accounts (Street 2006). The idea here is that our moral intuitions positively cohere with, and are not massively and systemically inconsistent with, the outputs of paradigm practices not because moral intuitions are converging on the truth, but rather because some alethically insensitive causal factors have prompted us to have the moral intuitions we do. As we acknowledge, the trustworthiness that attaches to the outputs of practices that fulfill conditions (i) – (vi) is defeasible. If some such adaptive link account were correct, then we concede that a defeater would indeed be in place. That said, we are highly skeptical of the merits of any such account, for the reasons given in Bengson, Cuneo and Shafer-Landau ms.

[46] Of course there may be *indirect* conflicts, once we supplement the outputs of paradigm practices with moral claims. But this is irrelevant to assessing the present threat, which is the charge that moral intuitions—the outputs of the moral intuition practice—are *themselves* massively and systemically inconsistent with the outputs of other practices in good working order. They are not.

[47] The qualification 'in normal circumstances' is meant to mark the fact that (for example) things can sometimes go awry, or a skeptical threat may obstruct the move from a good epistemic ground to the doxastic state to which it gives rise. Normal circumstances are just those in which there are no such obstacles—in such cases, the doxastic state is properly based on the epistemic ground, thereby acquiring the epistemic merit provided by its ground. This explains the truth of the principle articulated in the text.

[48] We're grateful to audiences at Antwerp, Dartmouth, the Princeton-Brown Moral Epistemology Workshop, Purdue, Vanderbilt, Virginia Tech, Uppsala, and Yale, as well as members of seminars at Harvard and the University of Wisconsin-Madison, for their feedback on earlier versions of this essay. We'd also like to extend our thanks to David Plunkett, Jonathan Shaheen, and Teemu Toppinen for their comments.

References

Alston, William P. 1991. *Perceiving God*. Cornell University Press.

Audi, Robert. 1997. *Moral Knowledge and Ethical Character*. Oxford University Press.

Audi, Robert. 1999. "Self-Evidence." *Philosophical Perspectives* 13: 205–28.

Audi, Robert. 2004. *The Good in the Right*. Princeton University Press.

Audi, Robert. 2015. "Intuition and its Place in Ethics." *Journal of the American Philosophical Association* 1: 57–77.

Ayer, A.J. 1946. *Language, Truth, and Logic*, second edition. Harmondsworth: Penguin.

Baier, Annette. 1986. "Trust and Antitrust." *Ethics* 96: 231–60.

Bealer, George. 1992. "The Incoherence of Empiricism." *Proceedings of the Aristotelian Society* 66: 99–138.

Bealer, George. 2008. "Intuition and Modal Error." In Quentin Smith, ed., *Epistemology: New Essays*. Oxford University Press.

Bengson, John. 2013. "Experimental Attacks on Intuitions and Answers." *Philosophy and Phenomenological Research* 86: 495–532.

Bengson, John. 2015a. "The Intellectual Given." *Mind* 124: 707–60.

Bengson, John. 2015b. "Grasping the Third Realm." *Oxford Studies in Epistemology* 5: 1–38.

Bengson, John, Terence Cuneo and Russ Shafer-Landau. ms.

BonJour, Laurence. 1980. "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy* 5: 53–73.

Bonjour, Laurence. 1998. *In Defense of Pure Reason*. Oxford University Press.

Bonjour, Laurence and Ernest Sosa. 2003. *Epistemic Justification*. Blackwell.

Bostrom, Nicholas. 2011. "Infinite Ethics." *Analysis and Metaphysics* 10: 9–59.
Brandt, Richard. 1979. *A Theory of the Good and the Right*. Clarendon Press.
Brink, David. 1989. *Moral Realism and the Foundations of Ethics*. Cambridge University Press.
Della Rocca, Michael. 2011. "The Taming of Philosophy." In Mogens Lærke, Justin E. H. Smith, and Eric Schliesser, eds., *Philosophy and Its History*. Oxford University Press.
DePaul, Michael. 1993. *Balance and Refinement*. Routledge.
DePaul, Michael. 2011. "Reflective Equilibrium." In Christian Miller, ed., *The Continuum Companion to Ethics*. Continuum.
Elgin, Catherine. 2008. "Trustworthiness." *Philosophical Papers* 37: 371–87.
Elgin, Catherine. 2017. *True Enough*. MIT Press.
Ewing, A.C. 1941. *Reason and Intuition*. H. Milford.
Fanselow, Ryan. 2011. "Self-Evidence and Disagreement in Ethics." *Journal of Ethics and Social Philosophy* 5: 1–17.
Friedman, Jane. Forthcoming. "Inquiry and Belief." *Noûs*.
Gibbard, Allan. 2008. *Reconciling our Aims: In Search of Bases for Ethics*. Oxford University Press.
Greene, Joshua. 2003. "From Neural 'Is' to Moral 'Ought': What are the Moral Implications of Neuroscientific Moral Psychology?" *Nature Reviews Neuroscience* 4: 847–50.
Haddock, Adrian, Alan Millar, and Duncan Pritchard. 2009. *Epistemic Value*. Oxford University Press.
Hintikka, Jaakko. 1999. "The Emperor's New Intuitions." *Journal of Philosophy* 96: 127–47.
Holton, Richard. 1994. "Deciding to Trust, Coming to Believe." *Australasian Journal of Philosophy* 72: 63–76.
Hurka, Thomas. 2015. *British Ethical Theorists from Sidgwick to Ewing*. Oxford University Press.
Jones, Karen. 1996. "Trust as an Affective Attitude," *Ethics* 107: 4–25.
Kagan, Shelly. 2001. "Thinking about Cases." *Social Philosophy and Policy* 18: 44–63.
Kitcher, Paul. 2000. "A Priori Knowledge Revisited." In Paul Boghossian and Christopher Peacocke, eds., *New Essays on the A Priori*. Oxford University Press.
Klein, Colin. 2017. Review of *Tense Bees and Shell-Shocked Crabs: Are Animals Conscious*? by Michael Tye. *Notre Dame Philosophical Reviews*. Available online at: http://ndpr.nd.edu/news/tense-bees-and-shell-shocked-crabs-are-animals-conscious/.
Leiter, Brian. 2014. "Moral Skepticism and Moral Disagreement in Nietzsche." In Russ Shafer-Landau, ed., *Oxford Studies in Metaethics*, vol. 9. Oxford University Press.
Longino, Helen. 1990. *Science as Social Knowledge*. Princeton University Press.
Lynch, Michael. 2006. "Trusting Intuition." In Peter Greenough and Michael Lynch, eds., *Truth and Realism*. Oxford University Press.
Mackie, J.L. 1977. *Ethics*. Penguin.
MacIntyre, Alasdair. 1984. *After Virtue*, second edition. University of Notre Dame Press.
McDowell, John. 1985. "Values and Secondary Qualities." In Ted Honderich, ed., *Morality and Objectivity*. London: Routledge & Kegan Paul.
McGrath, Sarah. 2008. "Moral Disagreement and Moral Expertise." In Russ Shafer-Landau, ed., *Oxford Studies in Metaethics*, vol. 3. Oxford University Press.
Moore, G. E. 1903. *Principia Ethica*. Cambridge University Press.
Murdoch, Iris. 1970. *The Sovereignty of Good*. Routledge.
Nussbaum, Martha. 1990. "The Discernment of Perception: An Aristotelian Conception of Private and Public Rationality." In *Love's Knowledge*. Oxford University Press.
Parfit, Derek. 2011. *On What Matters*, volumes I and II. Oxford University Press.
Pritchard, Duncan. 2014. "Seeing It for Oneself: Perceptual Knowledge, Understanding, and Intellectual Autonomy." *Episteme* 13: 29–42.
Railton, Peter. 2014. "The Affective Dog and Its Rational Tale: Intuition and Attunement." *Ethics* 124: 813–59.
Rawls, John. 1955. "Two Concepts of Rules." *Philosophical Review* 64: 3–32.
Rawls, John. 1971. *A Theory of Justice*. Harvard University Press.
Ross, W.D. 1930. *The Right and the Good*. Oxford University Press.
Sayre-McCord, Geoffrey. 1996. "Coherentist Epistemology and Moral Theory." In Walter Sinnott-Armstrong and Mark Timmons, eds., *Moral Knowledge*. Oxford University Press: 137–89.

Singer, Peter. 1974. "Sidgwick and Reflective Equilibrium." *The Monist* 58: 490–517.
Singer, Peter. 2005. "Ethics and Intuitions." *Journal of Ethics* 9: 331–52.
Sinnott-Armstrong, Walter. 2002. "Moral Relativity and Intuitionism." *Philosophical Topics* 12: 305–28.
Sinnott-Armstrong, Walter. 2007. *Moral Skepticisms*. Oxford University Press.
Sinnott-Armstrong, Walter. 2008. "Framing Moral Intuitions." In Sinnott-Armstrong, ed., *Moral Psychology*. MIT Press.
Solomon, Miriam. 2007. *Social Empiricism*. MIT Press.
Stratton-Lake, Philip. 2016. "Intuition, Self-evidence and Understanding." In Russ Shafer-Landau, ed., *Oxford Studies in Metaethics*, vol. 11. Oxford University Press.
Thompson, Michael. 2008. *Life and Action*. Harvard University Press.
Wagenknecht, Susann. 2016. *A Social Epistemology of Research Groups*. Palgrave Macmillan.
Weinberg, Jonathan. 2007. "How to Challenge Intuitions Empirically without Risking Skepticism." *Midwest Studies in Philosophy* 31: 318–43.
Weinberg, Jonathan, Shaun Nichols, and Stephen Stich. 2001. "Normativity and Epistemic Intuitions." *Philosophical Topics* 29: 429–60.
Williamson, Timothy. 2007. *The Philosophy of Philosophy*. Blackwell.
Wittgenstein, Ludwig. 1937. "Cause and Effect: Intuitive Awareness." In Peter Winch, trans., 1976: *Philosophia* 6: 409–25.
Wolterstorff, Nicholas. 2015. *Art Rethought: The Social Practices of Art*. Oxford University Press.
Wright, Crispin. 2004. "Warrant for Nothing (and Foundations for Free)?" *Proceedings of the Aristotelian Society* 78: 167–212.
Zagzebski, Linda. 2003. "The Search for the Source of Epistemic Good." *Metaphilosophy* 34: 12–28.
Zwicky, Arnold M. and Jerold M. Sadock. 1975. "Ambiguity Tests and How to Fail Them." In J.P. Kimball, ed., *Syntax and Semantics IV*. Academic Press.