

THINKING ABOUT CASES

BY SHELLY KAGAN

I. THE PRIORITY OF CASE SPECIFIC INTUITIONS

Anyone who reflects on the way we go about arguing for or against moral claims is likely to be struck by the central importance we give to thinking about cases. Intuitive reactions to cases—real or imagined—are carefully noted, and then appealed to as providing reason to accept (or reject) various claims. When trying on a general moral theory for size, for example, we typically get a feel for its overall plausibility by considering its implications in a range of cases. Similarly, when we try to refine the statement of a principle meant to cover a fairly specific part of morality, we guide ourselves by testing the various possible revisions against a carefully constructed set of cases (often differing only in rather subtle ways). And when arguing against a claim, we take ourselves to have shown something significant if we can find an intuitively compelling counterexample, and such counterexamples almost always take the form of a description of one or another case where the implications of the claim in question seem implausible. Even when we find ourselves faced with a case where we have no immediate and clear reaction, or where we have such a reaction, but others don't share it and we need to persuade them, in what is probably the most common way of trying to make progress we consider various analogies and disanalogies; that is to say, we appeal to still other cases, and by seeing what we want to say there, we discover (or confirm) what it is plausible to say in the original case. In these and other ways, then, the appeal to cases plays a central and ubiquitous role in our moral thinking.

Admittedly, some moral philosophers officially disavow the legitimacy of such appeals to our intuitions about particular cases. They attempt to make do without them, arguing that moral claims are better justified by appeal to something else, perhaps general principles that are themselves intuitively attractive or that can be shown to cohere well with other philosophical (or empirical) claims we find ourselves inclined to accept. But whatever the official pronouncements, I suspect that in practice the deft appeal to intuitions about cases is never actually eliminated. Like everyone else, moral philosophers—even those moral philosophers who think they know better—tend to be suspicious of moral claims that yield counterintuitive implications in particular cases. And like everyone else, moral philosophers—even those who insist that no legitimate comfort is

to be had in this way—are reassured when intuitions about particular cases support the particular moral claims they are putting forward. In short, whatever it is that some of us may say, what all of us actually do is appeal to, and give considerable weight to, our intuitive judgments about cases.

This is not to say, of course, that we all take such intuitions as fixed points, judgments that must be endorsed by any adequate moral theory. Our intuitions about cases provide us with *evidence* for and against rival moral claims—and it is difficult to imagine giving them no weight whatsoever. But that is not to say that the evidence must always be taken to be decisive, overriding any considerations at all that might lead us to reject our intuitive judgment about some particular case. On the contrary, most of us are prepared to dismiss some intuitions as ill-considered, or the result of mere bias or prejudice, or perhaps even moral illusion.¹ Still, the fact of the matter is that none of us is genuinely prepared to write off all of our intuitions in this way.

What seems more open to genuine debate is the question of just how *much* weight should be given to our intuitions about particular cases. Absent compelling reason to dismiss some particular intuition, most of us are inclined to give our intuitions about cases considerable weight. We trust them to a remarkable extent, using them, as I have already indicated, as the touchstones against which our various moral claims are to be judged. We take our intuitions about cases to constitute not only evidence, but compelling evidence indeed. I think it fair to say that almost all of us trust intuitions about particular cases over general theories, so that given a conflict between a theory—even one that seems otherwise attractive—and an intuitive judgment about a particular case that conflicts with that theory, we will almost always give priority to the intuition.

It would be tempting to describe this priority by saying that we trust intuitions over theories, but that wouldn't be quite right. For the fact is that we can have intuitions about theories and general principles them-

¹ In ethics, then, as elsewhere, we need to distinguish between intuition and belief, since one need not believe one's intuitions. At best, intuition involves something more like a disposition to believe. But of course not all dispositions to believe are intuitions. While it would be useful to have a general characterization of intuitions, this is a complicated subject and I will have to restrict myself to two further remarks. First, intuitions are normally taken to be "immediate" or "spontaneous," and while this apparently rules out dependence upon current conscious inference or reflection, it seems to leave open the possibility that prior reflection (or current unconscious inference) may have played a role in generating the present "immediate" intuition (and so, among other things, intuitions need not be unlearned). Second, not all "immediate" and "spontaneous" dispositions to believe qualify as intuitions. There is, I think, a further characteristic quality—one that I, at least, find difficult to describe—that is required as well: roughly, its simply "appearing" to one that something is the case. (Thus, although I am immediately disposed to believe that Washington D.C. is the nation's capital, it doesn't seem to me that I have any *intuition* to this effect.) It is not clear which of these features (or others) are relevant to justifying our reliance upon intuition, in ethics or elsewhere.

selves. After all, even a general principle can strike us as intuitively plausible, and thus garner support from that very fact. And yet it seems to me that even intuitively plausible principles can come into conflict with intuitions about particular cases—thus giving us a conflict between intuition and intuition—and when this happens it remains true that we will almost always be inclined to have greater trust in the intuition about the particular case. (Of course, once again, various considerations might ultimately lead us, on reflection, to endorse the general principle rather than the particular judgment about the particular case; but insofar as we focus solely upon the evidence provided by intuition itself, we tend to trust the intuition about the case far more than the intuition about the principle.) Thus, what is striking is not only our reliance upon intuition but, more particularly, our reliance upon intuitions about particular cases.

It is not at all clear to me what to make of this fact. Perhaps our pervasive and deep-seated reliance on intuitions about particular cases—what we might call “case specific intuitions”—is misguided. It is puzzling, at any rate, for it seems to me that although the extent to which we rely upon intuitions about cases is widely recognized, we don’t yet have anything like an adequate account of our practice—that is, a careful description of the various ways in which we appeal to, and give priority to, our case specific intuitions. Nor, I think, do we have anything like an adequate justification of our practice. While it is obvious that we constantly appeal to our intuitions about cases, it is far from clear what, if anything, makes it legitimate for us to give these intuitions the kind of priority we typically give them.

One (“deflationary”) possibility, of course, is that our reliance upon intuitions about particular cases is simply a reflection of a more general epistemic policy of relying on *all* of our various beliefs—and inclinations to believe—to the extent that we are confident about them. On such an account, all we could say is that we just happen to be especially confident about our various case specific intuitions; and while this might be a fact that would call for some sort of explanation (perhaps along evolutionary grounds), it would need no further *justification*. But the more ambitious epistemological alternative is to think that there is indeed some special justification for our reliance on case specific moral intuitions, something that warrants our particular confidence in them and our giving them the kind of priority that we do. I take it that most of us are actually drawn to this second view, and so the question remains whether there is in fact a plausible way to defend this idea, a way to justify our particular confidence in and reliance upon case specific intuitions.

II. THE ANALOGY TO EMPIRICAL OBSERVATION

The closest we typically come, I think, to justifying this reliance on moral intuition is to appeal to a certain analogy. It is often suggested (and

it is, at any rate, a natural suggestion to make) that we should think of case specific intuitions as playing a role in moral theory similar to that of *observation* in empirical theory. The suggestion, I presume, is sufficiently familiar that a bare sketch of the analogy should suffice.

Let's start with the role of observation. When arguing for or against empirical theories, we give unique weight to accommodating our observations of the world. We can simply see—immediately, and typically without further ado—that the liquid in the test tube has turned red,² or that the needle on the meter is pointing to 3, and an adequate empirical theory must take account of these facts. We appeal to such observations to provide support for a given theory, and we are very strongly inclined to reject any theory that runs afoul of them. Even a theory that seems otherwise attractive, and that strikes us as intuitively plausible in its own right, will be rejected if it contradicts the evidence provided by our empirical observations. To be sure, any given observation can itself be rejected (we might discover, for example, that we had unwittingly observed the test tube in red light), but for all that, no one seriously proposes that we should give no weight to our observations at all; and typically we give far greater priority to preserving the judgments of our observations than we do to maintaining our allegiance to any particular general empirical theory.

Similarly, then, when arguing for or against a moral theory we should think of our case specific intuitions as akin to observations. When thinking about particular cases we can simply see—immediately, and typically without further ado—whether, say, a given act would be right or wrong, or that it is morally relevant whether or not you have made a promise. An adequate moral theory must take account of these facts, it must accommodate these intuitions. To be sure, any given intuition can be challenged or rejected (we might, for example, realize that we made some judgment while inappropriately angry or embarrassed), but it would be quite implausible to suggest that we should give no weight to our moral intuitions at all. Indeed, even an otherwise plausible moral theory should be rejected if it contradicts the evidence provided by these intuitions; and so typically we appropriately give far greater priority to endorsing the judgments of intuition than we do to maintaining our allegiance to any particular general moral principle.

The analogy is indeed an appealing one, and it would be silly to dismiss it out of hand. But if we try to take it seriously certain points of disanalogy immediately suggest themselves. The most obvious worry—also familiar, and a natural one to think about—is this: in the case of empirical observation we have a tolerably good idea of how it is that the

² This is similar to an example of Judith Jarvis Thomson's, offered while making a similar point; see Thomson, *Rights, Restitution, and Risk* (Cambridge, MA: Harvard University Press, 1986), 257.

observations are produced. Visual observations depend upon the eyes, auditory observations depend upon the ears, and so forth. More generally, empirical observations depend upon the presence of well-functioning sense organs. In contrast, in the moral case, it is not at all obvious how it is that the corresponding “observations”—the moral intuitions—are produced. Is there a corresponding organ, a “moral sense,” that is at work here? If so, it must be admitted that we know precious little about it.

Now this complaint must not be misunderstood. The main complaint about an appeal to a moral sense had better not be that we don’t know how it *works*. For if that were the complaint it might not be especially worrisome. I take it, after all, that for most of human history we knew next to nothing about how the various sense organs worked either. But despite our ignorance, what was never in question was the existence of the various sense organs themselves (or that they were, indeed, sense organs). It was always fairly obvious, for example, that eyes were tied to visual observation, ears to auditory observation, and so on. In contrast, talk of a “moral sense” is nothing more than a place holder, a name for a supposed organ of moral intuition, something whose existence we may be led to infer (so as to have an account of the generation of moral intuitions), but concerning which we know virtually nothing else. And it is this, I take it, that gives us ground for skepticism, leaving us worried that there may be no such organ at all. Yet without a moral sense to correspond to the sense organs, the analogy to empirical observation is threatened.

Just how serious is the threat? Actually, this isn’t at all obvious. Even if there were no moral sense, no organ generating moral intuitions, the rest of the analogy might still go through. We could still regard moral intuitions as “input” for our moral theories, in roughly the way that we let empirical observations function as input for empirical theories. Perhaps there is no single moral organ (or set of organs) corresponding to the sense organs; still, the fact of the matter is that we have the various intuitions and we can treat them as input, accommodating them and giving them priority in the way that empirical observations are accommodated and given priority.

In any event, given the undeniable fact that we do have our various moral intuitions, it is not clear what harm there is in simply going ahead and positing a moral sense in the first place. Presumably, *something* generates the intuitions—they do not arise out of thin air!—and if we want to talk of the mechanism (or mechanisms) responsible for generating them as a “moral sense” or a “moral faculty” it is not clear what objection there can be to doing so, so long as we don’t thereby presuppose anything further about the structure or inner workings of that faculty.

The important question, rather, is whether we have special reason to *trust* our moral intuitions. Whether or not we posit a moral sense, the question remains whether there is good reason to take our intuitive judg-

ments as *evidence* in anything like the way we do. Even if there is a moral sense, an organ capable of generating moral intuitions, we still need to know whether it is more or less reliable.

It is precisely at this point, of course, that the analogy to empirical observation seems to beg the crucial question. After all, we all come to the discussion already convinced of the general reliability of the sense organs. (That is, we come to *this* discussion convinced of it; skepticism about the senses is not a worry we normally embrace when doing moral philosophy.³) Roughly speaking, then, we take the sense organs to be generally reliable, which is to say that empirical observations are generally reliable as well: *that* is why empirical theories must accommodate them. Similarly, then, once we make the assumption that our moral intuitions are generally reliable—that our moral sense, whatever it is, is generally reliable—then of course it will follow that our moral theories must accommodate our intuitions as well. But what justifies our assumption that our moral intuitions are reliable? Insofar as the analogy to empirical observation *presupposes* the reliability of our moral intuitions, it is not obvious how it can provide us with any reason to accept the claim that they are indeed reliable.

It is possible, however, that the analogy to empirical observation might still be found helpful, even here. For it might be suggested that our reasons for trusting our moral intuitions are analogous to our reasons for trusting the evidence of our senses.

Very well, then, what exactly *is* it that justifies us in thinking our empirical observations generally reliable in the first place? This is, of course, a complicated and much contested question, but at least one attractive answer begins by emphasizing the fact that we find ourselves strongly inclined to believe these observations—immediately, and without further ado—and so in the absence of a good reason to reject them, it is reasonable to (continue to) accept them. What's more, we are able to incorporate these observations into an overall attractive theory of the empirical world, one which admittedly rejects some of the observations as erroneous, but which for the most part endorses the claims of observation as correct. These two facts—the lack of reason for wholesale skepticism concerning

³ In point of fact, I don't think it altogether obvious to what extent empirical observations are indeed reliable. Consider shapes and sizes. Many people, I suppose, would be inclined to say, for example, that the stick in the water *appears* to me to be bent, even if I know better (and so believe it to be straight). But is it also the case—more controversially—that the building in the distance appears small (though I correctly infer that it is large, given its distance and apparent size), or that the penny seen from an angle appears to be oval (though I correctly infer that it is circular)? What, exactly, is it that I *observe* in such cases? What is it that I "just see," without further ado? Pursuing these questions would illuminate the nature of empirical observation, and thus might illuminate the nature of moral intuitions as well. But they seem to me quite complex—and they would certainly take us rather far afield—so I am going to put them aside, and assume in what follows that empirical observations are, indeed, generally accurate.

our senses, and our ability to construct an overall theory that in the main endorses our observations—together go a considerable way toward justifying us in taking our senses to be reliable.

Of course, to say that our senses are reliable is to say more than that they *happen* to be accurate, that empirical observations happen to be true. It is to claim that this level of accuracy is nonaccidental, that there is a connection between the truth of the relevant claims and the fact that they are given by empirical observation. (Very roughly, the presumed connection is this: it is because of the fact that P is true that we make the empirical observation that P; and were it not the case that P, we would not “observe” that P.) As we normally put it, our sense organs *respond* to the underlying empirical realities (and do so accurately, of course).⁴

This nonaccidental connection between observation and empirical reality is, obviously enough, an important part of what justifies our practice of actually *relying* upon our observations. After all, it is not as though we first construct a complete theory of the empirical world, and only then decide that our observations are, in the main, accurate. Rather, we construct enough of an account of the empirical world to justify us in taking our observations to be generally reliable, and then we use *further* observations to give us evidence concerning those aspects of the empirical world whose character we have not yet determined (as well as providing further confirmation for those aspects already known). We can *rely* on our observations only because we take it to be nonaccidental that our observations are accurate; we assume, that is, that our sense organs are *responding* to the world.

But what justifies us in taking our sense organs to be not just (accidentally) accurate, but reliably responsive in this way? I suspect it is primarily the very two facts already noted: we are strongly and immediately inclined to believe our empirical observations, and we can offer an (admittedly incomplete) overall theory of the empirical world that largely endorses the claims of observation as correct.

Given these two facts, we are justified in believing that ultimately—even if not initially—an account will be forthcoming which will display the inner mechanics of the sense organs in such a way as to explain just how this responsiveness is accomplished (that is, how it is that the nonaccidental connection between observation and fact is maintained). Of

⁴ Not surprisingly, talk of “responding” to the empirical world suggests that the world exists independently of—that is, metaphysically prior to—our empirical sensations and observations. But I take it that even “anti-realists” (who take the world to be somehow metaphysically constituted by our sensations or our reports of their contents) want a way to express the thought that our empirical observations are appropriately connected to the empirical facts, and for present purposes that is the only point at issue. Throughout this essay—both with regard to the empirical world, and the moral domain—I make use of familiar “realist” locutions. But I believe that roughly similar issues arise (concerning the reliability of both moral intuition and empirical observation) for both realists and anti-realists.

course, to believe that such an account can be produced is not yet to produce it. And eventually, no doubt, that promissory note must be made good: the account must indeed be produced. But I take it that our belief in the possibility of such an account can justifiably remain a mere promissory note for a good long time, since, as I have already noted, for much of human history we couldn't actually produce even the basic outlines of the relevant accounts. Still, given that we *were* able to produce an attractive overall theory of the empirical world that largely accommodated our empirical observations, it was nonetheless reasonable to conclude (albeit provisionally) that empirical observations are, indeed, not only accurate, but reliably so.

Analogously, then, it might be argued that we are also justified in taking our moral intuitions to be reliable. We certainly find ourselves strongly inclined to believe our moral intuitions—immediately, and without further ado—and so, in the absence of good reason to reject them, it is reasonable to (continue to) accept them. And if, going beyond this, we are also able to incorporate our intuitions into an overall attractive theory of morality, one which for the most part endorses these intuitions as correct, then even if the theory rejects some of the intuitions as erroneous, we will still be justified in taking our moral intuitions to be generally reliable.

Here too, of course, we will still find ourselves with a further explanatory obligation. If we are to justify our *reliance* on moral intuition it won't suffice if moral intuitions merely happen to be accurate: there must be, instead, a nonaccidental connection between moral intuition and the underlying moral realities. Thus, we must believe that ultimately an account will be forthcoming that will display the inner mechanics of the moral sense in such a way as to reveal how it succeeds in being responsive to the moral "facts." Eventually, no doubt, we will need to make good on this promissory note, and produce the requisite account. But just as we were justified in taking sense organs to be reliably responsive, even though we lacked (for most of human history) an account of how it is that this responsiveness was accomplished, we may still be justified (for the time being) in taking our moral sense to be reliably responsive as well, even if we still lack an account of how *that* responsiveness is accomplished. In short, given the compelling nature of our immediate moral intuitions, and given the existence of an overall moral theory that largely accommodates those intuitions, we are justified in believing that the requisite account of the moral sense may yet be forthcoming. Which is to say: we are justified in taking intuition to be reliable.

If an answer along these lines is to be accepted, however, it is important to give due weight to the claim that our various moral intuitions can indeed be incorporated into an overall attractive theory of morality. For it is only if we are truly able to construct such a theory that we are entitled to take our moral intuitions to be reliable.

To see this, consider the case of empirical observation again, and imagine that we were unable to construct a theory of the empirical world which largely endorsed our empirical observations. We would then dismiss the evidence of our senses as unreliable—illusory, not to be trusted. After all, our sense organs can hardly be reliable if empirical observations are not generally accurate, but we are only justified in taking empirical observation to be accurate given our ability to construct a plausible theory of the empirical world that largely endorses the observations. Thus, if we were unable to construct such a theory, we would be forced to dismiss the evidence of our senses as inaccurate and unreliable.

The point can perhaps be put this way: the fact that we find ourselves immediately and unreflectively inclined to accept our empirical “observations” only gives us reason to accept these observations as reliable *given* that we have no reason to be skeptical of their accuracy. It provides only a presumptive argument for accepting them. But if we find that we cannot construct an overall theory of the empirical world that (in the main) endorses the observations, then this very failure provides us with good reason to be skeptical. The presumptive argument provided by the intuitive force of the observations is overcome. Similarly, then, in and of itself the mere fact that we find ourselves immediately and unreflectively inclined to accept our case specific moral intuitions provides us with only a presumptive argument for accepting them. If we were to discover that we could not actually construct an attractive overall moral theory that (in the main) endorses these intuitions, then this presumptive argument would be overcome, and we would have reason to be skeptical about our moral intuitions. So the question we must ask ourselves is this: can we indeed produce a moral theory that appropriately accommodates our moral intuitions, incorporating them into an overall theory of morality that is itself plausible and attractive?

I don’t think the answer to this question is obvious, especially once we bear in mind that the requisite theory presumably must go beyond merely *organizing* the various “appearances,” but must itself be sufficiently explanatory so as to provide at least the beginnings of an account of the relevant phenomena. Consider the empirical case, yet again: we are satisfied that the requisite theory of the empirical world can indeed be produced, but we would not be satisfied if all we could do was organize our various empirical observations into systematic patterns. Instead, what we want, and what we take ourselves to be able to produce, is a theory that goes below the surface and provides something of an explanation of the empirical phenomena that are the subject matter of our empirical observations. We offer, that is, a theory of objects in space and time, interacting with one another and with ourselves, a theory that begins to explain how it is that the empirical world can have the particular features reported in our observations.

Similarly, then, in looking for a moral theory that will accommodate our case specific moral intuitions, it won't suffice if all we can do is organize these intuitions into systematic patterns. Instead, what we need to find is a moral theory that goes below the surface and provides at least the beginnings of an explanation of the moral phenomena that are the subject matter of our moral intuitions. That is to say: we need a theory that offers at least the outlines of an explanation of how the moral domain can indeed have the particular features ascribed by our various intuitions. What I take to be far from obvious is whether we can in fact produce an overall moral theory that is sufficiently explanatory in this way, while still accommodating the bulk of our moral intuitions.

Of course, the difficulty of this task will depend on at least two further issues: first, the precise content of the moral intuitions we are trying to accommodate, and second, the standards we impose concerning what will constitute an explanatorily adequate moral theory. Unfortunately, pursuing either of these issues here would take us too far afield. But let me register the following skeptical note. I have argued elsewhere⁵ that, in point of fact, certain widely accepted views—views central to common-sense morality and supported by the case specific intuitions of a great many individuals—cannot be provided with the kind of theoretical underpinnings we are here calling for. If I am right about this, then despite the immediate appeal of the relevant intuitions, they cannot be incorporated into an adequate overall moral theory, and in this regard, at least, our moral intuitions are unreliable.

I realize, of course, that many people would reject the particular arguments I've previously offered concerning the impossibility of providing an appealing and coherent moral theory that endorses these common moral intuitions. It is important to note, however, that in at least some cases the rejection of these arguments would simply take the form of pointing out how counterintuitive the implications of these arguments are, and in the present context, at least, such an appeal to intuition would constitute begging the question. For insofar as we are trying to establish whether our case specific moral intuitions are to be trusted or not, a simple appeal to the *force* of these intuitions shows nothing. We are only justified in trusting our intuitions if we can indeed construct a moral theory that adequately explains and incorporates them, and this, of course, is precisely what I am saying we cannot do. Thus, the mere fact that the conclusions for which I have argued are incompatible with many forceful and widely held intuitions does nothing to show that the requisite moral theory *can* be constructed. Indeed, as I have already noted, I think there are good reasons to conclude that we cannot, in fact, produce the requisite moral theory.

⁵ See, especially, Shelly Kagan, *The Limits of Morality* (Oxford: Oxford University Press, 1989).

III. ERROR THEORIES

Let's recap. We have been taking seriously the analogy between moral intuitions and empirical observations, so as to see what might justify our practice of giving our case specific intuitions the kind of priority that we do. I have been suggesting, of course, that if we are to be justified in trusting our intuitions in this way, there must be an explanatorily adequate moral theory that endorses (not all, but most of) our case specific intuitions, just as we take ourselves to be justified in trusting our empirical observations by virtue of having an explanatorily adequate empirical theory that endorses (most of) our empirical observations. And as I have already noted, my own opinion is that once we take seriously the need to construct a general moral theory that would endorse our case specific intuitions as being for the most part accurate, we will find it difficult, indeed impossible, to produce the requisite theory. Theories that attempt to accommodate the bulk of our various case specific intuitions fail, I believe, at one or another explanatory task, and fall short in overall plausibility. What we are led to, instead, is a general moral theory according to which many of our specific moral intuitions are simply mistaken.

If I am right about this, then at a minimum we will have reason to be skeptical about these particular common moral intuitions. More generally, however, and for our current purposes more importantly, we will have reason to conclude as well that moral intuition is not, on the whole, reliable. Instead, the appropriate stance to take toward our moral intuitions will involve accepting an *error theory*, according to which at least many of our case specific moral intuitions are mistaken.⁶

Of course, there are various kinds of error theories—some more radical than others—and we've not yet addressed the question of whether our moral intuitions need to be discounted altogether. At one extreme lies just such wholesale skepticism concerning our case specific moral intuitions. But more modest versions of error theories are possible as well, and it might be that our best overall moral theory still endorses some specified range of moral intuitions, while nonetheless writing off other classes of intuitions as mistaken.

However, even such moderate error theories will seem unattractive to many. They will hold, correctly, that to accept an error theory—even a modest one—is to retreat significantly from our current practice, where appeals to intuition are generally taken across the board to be a particularly important source of evidence concerning the moral domain.

And so, despite my own skepticism, many will insist on remaining optimistic about the prospects for constructing a moral theory that actually succeeds quite generally in accommodating our case specific intu-

⁶ Strictly, of course, everyone accepts an error theory of at least a rather *minimal* sort, since no one thinks that moral intuitions can never be mistaken. But I have in mind more ambitious theories of this type.

itions. They will want to reject any error theoretic approach to moral intuition at all. They will claim that our moral intuitions are, in point of fact, typically accurate, and that we are justified in thinking that it is nonaccidental that this is so. Thus, they will insist that we are justified in taking moral intuition to be reliable.

There are, however, still further grounds for skepticism about the overall reliability of our moral intuitions that we have not yet considered. What I have in mind is the surprising—and typically overlooked—extent to which people's intuitions actually differ with regard to specific cases. The extent of the disagreement is overlooked for the simple reason that we normally don't *look* for such disagreement. We barely entertain the possibility that others may not agree with us, and so we typically don't look around very carefully to see just how widely shared our particular intuitions actually are. And when we do stumble upon such cases of intuitive disagreement, it surprises us. Our own intuitions are sufficiently compelling and powerful that the relevant judgments strike us as virtually self-evident, and we are, accordingly, shocked if other, apparently reasonable individuals don't share them.

I do not mean to suggest, of course, that intuitive disagreements arise with regard to every case, though it does seem to me—based on years of discussing such cases with students and others—that even the most compelling examples typically fall short of garnering complete agreement. And in many cases, I think, once one probes a bit one finds that there is actually a considerable amount of disagreement. Consider, for example, “trolley problems” of the kind frequently used to determine the precise nature of the prohibition against harming others.⁷ In my own classes I generally find that only about three fourths of the students share the majority intuition (say, that it is permissible to turn the trolley), while up to a fourth disagree; and even the apparent agreement of the three fourths majority dissolves when one asks further questions (for example, whether one is required, or only permitted, to turn the trolley).

To be sure, it is difficult to be confident that the opinions being reported in such informal polls truly state the immediate moral intuitions of the students in my classes. As we have already noted, we need to distinguish between the immediate pronouncements of our case specific intuitions and the various beliefs about a case one might have instead (for example, as a result of conscious reflection). In short, when students vote in such polls, are they reporting moral intuitions, or simply stating their own tentative beliefs about the cases? It might well be that despite the exis-

⁷ In the basic case, a runaway trolley will hit and kill five children, unless you throw a switch which will divert the trolley onto a side track, saving the five, but killing a sixth child trapped on that side track (who would otherwise be safe). A large number of variants of this basic case have been discussed. See, e.g., Judith Jarvis Thomson, “Killing, Letting Die, and the Trolley Problem,” in Thomson, *Rights, Restitution, and Risk*, 78–93; and Frances Kamm, “Harming Some to Save Others,” *Philosophical Studies* 57, no. 3 (1989): 227–60.

tence of widespread disagreement in opinions about the relevant cases, there is actually far greater agreement with regard to the immediate intuitions themselves.

This is certainly a possibility, and I don't mean to suggest that I conduct my polls with sufficient care to rule it out. (It would be useful to have some careful empirical studies of these matters.) Still, it seems to me likely that intuitive disagreement is indeed a fairly widespread phenomenon.

What's more, I suspect that such disagreement is far from a random affair. It is not that any given individual almost always agrees with the majority, but sporadically finds himself faced with an idiosyncratic intuition, one as much at odds with the rest of his own intuitions (at other times, or in other cases) as it is at odds with the majority. If this were the nature of intuitive disagreement, we might well feel free to write off the occasional, quirky intuition as a mere aberration—a random misfiring in an otherwise reliable moral sense. In fact, however, it seems to me that moral disagreement is systematic and patterned. A given individual is likely to be regularly responsive to certain features that cases might display, while other individuals are routinely indifferent to the presence (or absence) of those features, or react to them in quite different ways. In short, intuitive disagreement doesn't take the form of norm and aberration. Rather, it is as though moral senses fall into distinct types, each with its own regular pattern of intuitive responses.

If I am right about this, obviously enough, it greatly complicates the position of anyone who hopes to endorse moral intuitions as largely correct. For if people actually differ considerably as to the content of those intuitions, even when thinking about the very same cases, then clearly not everyone's intuitions can be largely reliable. So what should we say?

One possibility, I suppose, would be to hold that everyone's intuition is indeed reliable, but only in those areas where there is complete agreement (assuming that such an area of complete agreement is to be found at all). But if we do say this, then we face the difficult task of explaining why intuition is indeed reliable in exactly those areas. What is it about the areas of agreement that makes intuition there function properly, and what is it about the other areas that causes intuition to break down and malfunction? Apparently, even those who hope to endorse moral intuition to this limited extent require an error theory, and an error theory of a fairly subtle sort, for they need to explain why intuition malfunctions in certain areas while working reliably in others. Absent a story about the mechanics of moral intuitions—the workings of the moral sense—any confidence that intuition is indeed to be trusted at all, even where there is agreement, may seem strained or premature.

More ambitiously still, some might hold out the hope of justifying reliance upon moral intuition even in those cases (considerable, as I believe) where there is intuitive disagreement. Clearly, however, this requires dismissing as flawed the moral senses of all those who stand in

intuitive disagreement with the intuitions being endorsed. At best, the moral intuitions of only certain individuals can be held to be generally reliable. For the rest, then, we will inevitably need to embrace an error theory of a different sort: we will require an account which explains how most (or at least many) people end up with unreliable moral intuitions, while the moral sense of others nonetheless ends up functioning properly and reliably. And we will need an epistemological account as well, so as to justify us in our position concerning just whose intuition is to be trusted as reliable. (Obviously, it won't do to simply assume without further ado that it is *mine* that functions properly.)

This is not to say that these various explanatory burdens could not possibly be met. Once again, empirical observation provides a helpful analogy, for we do find ourselves, in the case of color blindness, arguing for something at least roughly comparable. Certain individuals are said to have damaged or flawed visual senses—leading to inaccurate visual observations, in at least a specifiable range of cases—while the rest of us are held to have properly functioning and reliable visual senses nonetheless. If something like this can be plausibly held to occur in the case of empirical observation, why not in the case of moral intuition as well? Is it so implausible to think that certain individuals are “morally blind”—cursed with inaccurate moral intuitions, in at least a specifiable range of cases?

The analogy to color blindness certainly suggests that something similar might arise in the case of moral intuition as well. But it is one thing to admit the mere possibility of something like this, it is quite another to make good on the claim that “moral blindness” actually occurs, and still another thing to warrant applying this label to some particular individual. In the case of color blindness, after all, we are able to demonstrate, even to the satisfaction of the color blind themselves, that their visual apparatus is indeed impaired, and that they fail to respond accurately to genuine features of the empirical world, features that the rest of us are able to detect through our own unimpaired visual senses. It is far from clear whether anything analogous can be done in the case of disagreement of moral intuitions, or even how one would go about trying to make out a comparable case. Instead, the charge of moral blindness more typically seems little more than name calling, where we blithely dismiss the intuitions of those who disagree with us, assuming without any further evidence than the mere fact of the disagreement itself that it is they who are blind, rather than us.⁸

I have been arguing that given the nature of intuitive moral disagreement, no one, not even those who hope to endorse moral intuition as

⁸ The situation is further complicated by the fact that *each* side may fail to respond to features that the other side's intuitions mark out as morally significant. Thus, unlike the normal case of color blindness, moral disagreement may actually be closer to a situation in which many groups claim to see one or more colors that some other groups do not, and yet each group still fails to see some of the colors that other groups claim to see.

generally reliable, can escape the need to accept some kind of error theory with regard to at least many moral intuitions. And I have suggested as well that until we produce at least the beginnings of a story about the mechanics of moral intuition it is difficult to be confident that the requisite error theory can be produced. Attempts to limit the error theory—so that it impugns only a certain range of intuitions, or a certain group of moral senses—may easily fail, so that we are left with no good reason to believe our moral intuitions to be especially reliable at all.

But I do not mean to suggest that matters are particularly easier for those who hope to embrace far more radical error theories, dismissing most, or all, of our moral intuitions as suspect. For the fact is, producing a plausible error theory even of this radical sort is extremely difficult as well.

Consider, for example, the suggestion that is sometimes made that our case specific intuitions can be dismissed out of hand, as the mere historical by product of outdated religious views or neuroses about sex, or that they are merely the results of internalizing dubious moral teachings received in childhood.⁹ Were this the case, there might well be little reason to give any weight at all to our case specific moral intuitions, and the wiser course of action would be to attempt to elaborate moral theories simply without appeal to them, however difficult that might prove to be.

But although accounts along these lines may well rightly cast doubt upon certain case specific intuitions (say, about sex), they seem rather inadequate as general explanations of the origins of our moral intuitions. Consider again the appeal to trolley problems as a means of determining the precise content of the prohibition against harming. Such cases are highly stylized, and unlike anything most of us have ever faced in real life, read about, or even imagined before being introduced to them for the first time as adults. Yet once the given case is described, we typically find ourselves with a moral intuition about it. I think it highly implausible, accordingly, to suggest that what happens here is that some vestige of a (perhaps forgotten) religious teaching now comes into play. No one is taught about trolley problems in childhood—nor even anything remotely similar to them—and yet we still find ourselves with intuitive reactions to the cases once they are described. Thus, whatever the actual origins of these case specific intuitions, we cannot dismiss them as artifacts of outmoded or unjustified teachings and accidental historical influences. For the simple fact of the matter is that most of our case specific intuitions cannot be plausibly explained in this way.

We may do somewhat better if we appeal, instead, to some of the primitive beliefs about physics or the nature of agency that we may well inherit as a result of our evolutionary history, as well as to certain innate

⁹ See, for example, Peter Singer, "Sidgwick and Reflective Equilibrium," *The Monist* 58, no. 3 (1974): 516.

psychological biases in terms of how to group people and events.¹⁰ An error theory that dismisses (many of) our case specific intuitions on the ground that they are implicitly based on inherited but dubious physical theories may well have an easier time of it explaining how we can have immediate and intuitive reactions to trolley cases, say, despite never having considered such cases previously. We may, for example, react to a given case as we do because we are innately disposed to view it in terms of mistaken concepts of causation and agency.

Here, too, such an account may rightly cast doubt upon certain of our case specific moral intuitions. But even an account of this sort seems inadequate, in large part because of the very universality of the inherited biases and beliefs that it presupposes. If our case specific intuitions are to be explained in terms of innate (though false) views about physics, say, then we would expect that people's intuitions would be fairly uniform—all reflecting the same set of inherited, though dubious, physical beliefs or psychological dispositions. In fact, however, as I have already suggested, it seems to me that we differ from one another in terms of our moral intuitions, in ways that this sort of account cannot easily accommodate. Intuitive disagreement is widespread and systematic, and it is implausible to dismiss our case specific intuitions on the ground that they are based on shared, inherited—and false!—views about the world, if in point of fact many of the relevant intuitions are not universally shared at all.

An error theory adequate to the facts about our moral intuitions would apparently have to be a rather subtle affair. It would need to accommodate the simple fact that we readily have intuitive reactions to cases quite unlike anything that we have faced or been taught about previously, and yet at the same time it would need to accommodate the fact that when we think about such cases our intuitive reactions are not all the same: people's intuitions differ, in systematic and patterned ways. It is not at all obvious what such an error theory would look like.¹¹

I don't mean to suggest, however, that it will be impossible to produce an error theory adequate to the facts. Indeed, if I am right that everyone needs an error theory of some sort—both those on the whole trusting of moral intuition, and those on the whole skeptical of it—then it seems inevitable that *some* sort of error theory must be right, and I see no particular reason to assume that we cannot eventually articulate and

¹⁰ See, for example, Peter Unger's discussion of protophysics and psychological grouping principles in his *Living High and Letting Die* (New York: Oxford University Press, 1996).

¹¹ It might seem that an emotivist or expressivist account of moral claims would have an easy time accommodating these facts, since there is nothing especially surprising in the suggestion that people's emotional (and other) attitudes vary, and that they can be readily generated in response to never before considered cases. But even accounts of this kind, it seems to me, should be troubled by the ease and force with which intuitions can be generated in response to trolley problems (and the like) since it is not at all obvious why these should so readily engage our emotions or other pro-attitudes, nor why minor changes in the cases should elicit such drastically altered reactions.

defend this theory, whatever it is. But for the time being, at any rate, it seems to me that we are rather far from having an adequate account of what this theory looks like, and so, lacking it, we are rather far from knowing to what extent our moral intuitions can be trusted.

IV. PARTICULAR CASES AND GENERAL CLAIMS

Let me close by noting one further complication. Recall the fact, previously noted, that our moral intuition is capable of responding not only to particular cases but also to general moral principles and moral theories. Consider how different this is from the case of empirical observation, where all we can directly observe are the features of particular cases. I can simply see that the meter is pointing to 3, but I cannot simply see the truth of Ohm's law or other principles of physics at all. General empirical claims must be inferred from the evidence; one cannot simply observe their truth. Apparently, our sense organs are incapable of responding directly to general empirical truths in this way.

In itself, this may be no more than a striking disanalogy between the case of moral intuition and the case of empirical observation. But it points to a deeper problem. For we have also already noted the fact that we do not give the same kind of priority to our intuitions about general moral claims. What we particularly trust, rather, are our case specific intuitions, so that given a conflict between an intuition about a particular case and an intuition about a general moral claim, we are almost always inclined to endorse the intuition about the particular case (at least, insofar as what we are attending to is the evidential force of the intuitions themselves). We give priority not to intuition in general, but, more particularly, to our case specific intuitions.

Yet how is this fact to be explained? If the situation were like that of empirical observation—with the relevant sense only capable of responding directly to particular cases rather than to general principles as well—there would, of course, be nothing further to explain (although, no doubt, we would ultimately want to explain just why it is that the given sense can respond only to particulars). But given that moral intuition is capable of reacting both to particular cases and to general principles, we do need a further explanation: we need to understand just why it should be the case that intuition is particularly reliable only with regard to specific cases. What makes our intuition more reliable for the one sort of object rather than the other?

Once we put the question this way, however, it may seem that the answer won't be particularly hard to come by. Even if moral intuition (unlike empirical observation) is capable of reacting both to particular cases and to general claims, there is no particular reason to assume that it will be equally adept at handling both *kinds* of objects. Although, no doubt, the details of the explanation will need to await a theory of the

inner mechanics of moral sense, there is nothing particularly perplexing in the claim that intuition reacts more reliably when directed to one particular kind of object.

But this reassuring answer is itself threatened by the realization that this very distinction between two *kinds* of objects for intuition may well be misguided. For the fact of the matter, I believe, is that when we react to particular cases we are actually reacting to things of the very same type as when we react to general moral claims. It is easy to lose sight of this, given our common practice—one that I have followed in this paper as well—of saying that we are reacting to *particular* cases. But what we are actually reacting to, I think, are *types* of cases.

This is easiest to see in the situation where the kind of case we are thinking about is purely imaginary. What we are presented with, then, is only a description—and typically, all things considered, a fairly thin description at that. There is no actual, particular, concrete case that we are confronted with. So when our intuition tells us, say, that some particular act would be the right thing to do in that particular case, what we are actually intuiting, it seems, is that a certain *kind* of act would be the right thing to do in a certain *kind* of case. And this, of course, is a general moral claim.

The same thing is true, I think, even when the particular case being judged is an actual one. Again, this is easiest to see if the case, despite being real, is not one that we actually observe. We might only be *told* about the case, which means, of course, that we are again presented with a mere description. But this means, I take it, that we are not actually reacting to a particular, concrete case, but rather to a *type* of case. So here, too, when we react to the case what we are actually intuitively responding to is, it seems, something general: we are intuitively seeing that, say, this kind of act would be the right thing to do in this kind of case.

Although the point is controversial, I think the same is probably true even in those situations where we are literally faced with an actual, concrete case. Even in cases like this, I suspect that what we are actually responding to is its being a case with various salient features. By virtue of being literally faced with the case—able to observe it for ourselves—we better come to see that it has certain features, and we then intuit that the right thing to do, given a case with *these* features, is such and such. But if that is right, then here, too, we are reacting to something general: we are seeing that such and such an act is the right thing to do in this *kind* of case.

This is not to deny that being actually presented with a concrete case may elicit a different intuitive reaction than merely being presented with a description of the case. (When we literally *see* the needs of others we may intuitively see the importance of helping them, in a way that no mere description of their needs would elicit.) But even if it is true that in such cases there can be something special about intuition in the face of genuinely concrete particulars, the fact would remain that typically when we

think about cases, we are only thinking about *kinds* of cases. Which is to say, typically when we think about cases we are intuitively reacting to something general.

This makes it harder to explain the priority we want to give to our intuitive reactions to “particular” cases. If all, or at least most, case specific intuitions are not actually reactions to something concrete and particular at all, then we cannot readily claim that what makes intuition more reliable here is that it is directed at a different kind of object than when we intuitively respond to a general moral claim. In both cases, it seems, what we see is something general.

Of course, there will still be differences in degrees of generality, and it might be that what we should give priority to are our intuitive reactions to the less general rather than to the more general. But this, too, calls out for explanation, and it is not clear what could be said in its defense.

For when we face the fact that typically (at least) when we think about a case, we are indeed only *thinking* about it, we are reminded of the fact that intuitive reactions are, in some suitably broad sense of the term, *a priori*. Typically, at least, we don’t need to actually see the case; we only need to think about it. But it is not, as far as I can see, a general feature of the *a priori* that such thoughts are more reliable when they are directed to the less general rather than the more general. So it remains unclear why moral intuition should be thought particularly reliable in just such cases.

V. CONCLUSION

I have been arguing that our reliance upon case specific moral intuition is problematic, and in need of a justification that we do not yet possess. Most importantly, of course, anyone who is going to rely on intuition at all—and that, I think, means all of us—needs to explain exactly why we are justified in taking intuition to be particularly reliable in the first place. This is a justificatory burden that has not, I think, been satisfactorily discharged. In particular, despite the obvious appeal of an analogy to the case of empirical observation, there are, it seems, sufficient disanalogies here, so that at a minimum considerably more needs to be said. Furthermore, if, as I think, we must all accept some sort of error theory (whether modest or radical) with regard to moral intuition, then we must face the further fact that providing an adequate error theory is itself a surprisingly difficult task. Apparently, our reliance upon intuition must be tempered; but how, or in what ways, is not yet clear.

In sum, the extent to which intuition is to be trusted—if at all—remains unsettled. Our reliance upon moral intuition remains troubling.

Still, the fact remains as well that despite these questions we are all inclined to attend to our case specific intuitions. We worry when our moral beliefs run afoul of them, and we take comfort in the extent to which our moral beliefs accord with them. It may well be, as I believe,

that our moral intuition deserves considerably less respect than it is normally accorded. But it is difficult to believe that we could ever make do without it altogether. No moral argument—no claim, no theory—will ever seem compelling if it has not been subjected to the testing we provide when we think about cases.

Philosophy, Yale University