



INESCAPABILITY AND NORMATIVITY

BY MATTHEW SILVERSTEIN

JOURNAL OF ETHICS & SOCIAL PHILOSOPHY

VOL. 6, No. 3 | DECEMBER 2012

URL: WWW.JESP.ORG

COPYRIGHT © MATTHEW SILVERSTEIN 2012

Inescapability and Normativity

Matthew Silverstein

It may seem a far cry from Kant to Professor Stevenson ...

—R. M. Hare, *The Language of Morals*

WHEN WE MAKE ETHICAL JUDGMENTS, we invoke a kind of authority or objectivity. When we insist, for example, that people ought not to harm others just for fun, we are not merely claiming that this is so according to *some standard or other*, but rather that it is so according to *authoritative* or *normative* standards.

A familiar worry about our practice of making ethical judgments is that this invocation of normative authority somehow always fails. The worry can arise in a number of different forms. Its most common form was perhaps best expressed by John Mackie in the first chapter of *Ethics: Inventing Right and Wrong* (1977). There Mackie supposes that when we judge that people ought not to harm others just for fun, we *believe* that there is an authoritative standard according to which this is so. Mackie then argues that this belief is always mistaken: “[A]lthough most people in making moral judgments implicitly claim, among other things, to be pointing to something objectively prescriptive, these claims are all false.”¹ In other words, although our ethical judgments presuppose that authoritative standards or objective values are somehow sewn into the fabric of the world, in fact the world contains nothing so “queer.”² Consequently, ethics is not the secure and successful practice it appears to be. Like astrology and witchcraft, ethics – Mackie concludes – turns out to be grounded in nothing more than an illusion.

Over the years, philosophers have attempted to deflect Mackie’s broadside in various ways. Many of these philosophers opt for some form of *realism*. *Nonnaturalists* acknowledge that normative authority is queer but maintain that the fabric of the world contains it nonetheless. These philosophers face familiar metaphysical and epistemological worries. *Naturalists* contend that objective authority is not queer at all and is instead identical or reducible to some natural property. These philosophers avoid the metaphysical and epistemological problems that plague nonnaturalists but instead invite the objection that they have changed the subject from ethics to something else. They face the formidable challenge of showing that the natural property they have identified really is the property we mean to invoke when we make ethical judgments.

As realists, both naturalists and nonnaturalists accept Mackie’s account of the shape and scope of our ethical practices. They both acknowledge that we can vindicate those practices only by showing that objective values or authoritative standards really are built into the fabric of the world. Yet if the

¹ Mackie (1977: 35).

² *Ibid.*, 38.

course of metaethics over the past hundred years tells us anything, it tells us that this is no easy task. Ethics and the fabric of the world do not seem like a very good fit for one another.

Bernard Williams (1995), in his discussion of Mackie's book, outlines a strategy for securing the objectivity of ethics – and thus for vindicating our ethical practices – that does away with the problematic supposition shared by Mackie and his realist critics:

Consider another picture of what it would be for a demand to be “objectively valid.” It is Kant's own picture. According to this, a demand will be inescapable in the required sense if it is one that a rational agent must accept if he is to be a rational agent. It is, to use one of Kant's favorite metaphors, *self-addressed* by any rational agent. Kant was wrong, in my view, in supposing that the fundamental demands of morality were objective in this sense, but that is not the immediate point, which is that the conception deploys an intelligible and adequate sense of objectivity. It seems to have little to do with those demands being a part of the fabric of the world; or, at any rate, they will be no more or less so than the demands of logic – which was, of course, part of Kant's point.³

On the Kantian view sketched by Williams, we need not posit an ethical swatch of the fabric of reality in order to account for the objectivity of ethics. Ethical judgments correctly invoke normative authority not by corresponding to or picking out anything in the world, but rather by expressing demands or commitments that are, from the practical point of view, *inescapable*.

This Kantian strategy has proved quite popular in recent years.⁴ Much of its allure is due to its promise to soothe Mackie-style worries about our ethical practices while avoiding the metaphysical, epistemological and semantic controversies surrounding realism. In fact, as Stephen Darwall, Allan Gibbard and Peter Railton observe in their survey of *fin de siècle* ethics, many philosophers drawn to the Kantian strategy hope to bypass the metaphysics, epistemology and semantics of normative discourse altogether. They aim “to sidestep traditional metaethics,” or even to “render traditional metaethics obsolete.”⁵ These Kantians tend to believe that standard metanormative theories – theories such as expressivism, as well as the nonnaturalist and naturalist varieties of realism mentioned above – somehow miss what is interesting or

³ Williams (1995: 174-75).

⁴ What I am calling “the Kantian strategy” actually goes by many names, including *constructivism* (Korsgaard (2008); Street (2008)), *constitutivism* (Coleman (unpublished manuscript); Ferrero (2009)) and *ethical rationalism* (Setiya (2007)). Each of these labels is used in a number of ways, however, and not all of these uses are compatible with the strategy outlined by Williams. In order to avoid ambiguity, I shall continue to refer to the view that is our subject simply as the “Kantian strategy.”

⁵ Darwall, Gibbard and Railton (1992: 142, 143).

important about ethics, and they offer their approach as an alternative to such familiar but “boring” views.⁶

In this paper, I argue that the Kantian strategy for securing the objectivity of ethics cannot make good on this promise: it cannot vindicate our ethical practices while sidestepping traditional metaethics. Put another way, the Kantian strategy is not an alternative to familiar metanormative theories; the sense of objectivity which it deploys can be, to borrow Williams’ word, “adequate” only with the support of a metaethic.

In section 1, I describe the Kantian strategy in greater detail, drawing on the work of some of its more well-known proponents. I also identify a potentially crippling lacuna in the strategy and argue that the Kantians cannot fill this lacuna without relying on metanormative assumptions. I do not conclude that we should therefore abandon the Kantian approach, however. Instead I introduce two versions of the Kantian strategy which appeal explicitly to metanormative premises: *Kantian expressivism* (in section 2) and *Kantian reductionism* (in section 3). I argue that the former avoids the potentially crippling lacuna and that the latter can easily bridge it. The Kantian strategy *can* succeed, then, but not without metanormative assistance. Finally (in section 4), I consider and reject the possibility that the Kantian strategy involves a metaethic of its own, distinct from both expressivism and reductionism.

1. The Kantian Strategy

In *The Sources of Normativity*, Christine Korsgaard (1996) claims that the authority or normativity to which an agent’s ethical judgments appeal is to be found “in the agent’s own will, in particular in the fact that the laws of morality are the laws of the agent’s own will and that its claims are ones she is prepared to make on herself.”⁷ The laws of morality bind us, in other words, because they consist of demands to which we are already committed: they are demands we make on ourselves. Moreover, their authority is *objective* because *we cannot help* but make these demands; they are “inescapable.”⁸ As Korsgaard sometimes puts it, it is *constitutive* of agency that we value ourselves as agents: “[B]eing human requires it.”⁹ And since “you must value your own humanity if you are to act at all,” Korsgaard concludes, “human beings are valuable.”¹⁰

Here we can see the Kantian strategy outlined by Williams at work. Korsgaard acknowledges that Mackie-style worries about the validity of our ethical practices might lead us to wonder whether “obligations really exist.”¹¹

⁶ Korsgaard (2008: 325, n. 49).

⁷ Korsgaard (1996: 19).

⁸ Korsgaard (1996: 130).

⁹ Korsgaard (1996: 121).

¹⁰ Korsgaard (1996: 123).

¹¹ Korsgaard (1996: 38).

And she argues that a realist account of those practices – an account according to which normative authority is somehow sewn into the fabric of the world – cannot allay such worries. In lieu of realism, Korsgaard offers what she describes as a Kantian account of the sources of normativity. According to this Kantian theory, we should respond to skeptical doubts about the demands of morality not by digging in our heels and insisting, as the realist does, that obligations really do exist, but rather by showing that those demands are inescapable.

The Kantian strategy sketched by Williams also features prominently in the recent work of J. David Velleman. According to Velleman, demands are inescapable (and hence objective) by being related to the *constitutive aim* of action. The notion of a constitutive aim of action is the notion of “an aim with respect to which behavior must be somehow regulated in order to qualify as action.”¹² It is, in other words, the idea of a motive or drive that “sets you to reasoning in a way that makes you autonomous, constituting you as an agent.”¹³ Drawing on his work in the philosophy of action, Velleman contends that the constitutive aim of agency is intelligibility. He then argues that this aim generates a standard of correctness for action – a standard that is objectively valid in virtue of being “constitutively inescapable” for any creature asking practical questions about what to do.¹⁴ Put another way, the constitutive aim of action gives rise to demands to which any agent is subject merely in virtue of being an agent. These demands are sewn into the practical point of view, and “anyone to whom they are addressed must already occupy that point of view and must therefore accept whatever demands are woven into it.”¹⁵ And therein, Velleman urges – appealing explicitly to Williams’ Kantian proposal – lies their authority or normativity: “[P]ractical thoughts can be objectively valid by being inescapable.”¹⁶

For both Velleman and Korsgaard, the Kantian strategy culminates in a straightforwardly *normative* conclusion. Korsgaard arrives at the verdict that human beings are valuable, Velleman at the claim that “[action’s] criterion of correctness is intelligibility.”¹⁷ It is not clear, though, how this final stage of the Kantian strategy – the move from inescapability to normativity – is supposed to work. That is, it is not clear how we are supposed to derive normative conclusions about what is actually valuable or about what makes actions correct merely from premises about what we inescapably value or desire. Korsgaard takes the inference to be obvious: “We find ourselves to be valuable. Therefore, *of course*, we are valuable.”¹⁸ But it is *not* obvious at all; the suggestion that we necessarily value our own humanity at least appears to be

¹² Velleman (2004: 234).

¹³ Velleman (2009: 137).

¹⁴ Ibid.

¹⁵ Velleman (2009: 116).

¹⁶ Ibid.

¹⁷ Velleman (2009: 134).

¹⁸ Korsgaard (1996: 124) (emphasis added).

perfectly consistent with the possibility that humanity is not in fact valuable. The most that follows obviously from the inescapability of some commitment or aim is that normative reasoning must always proceed from shared premises. But the fact that normative reasoning must always begin with shared premises does not entail that those premises are the *correct* ones from which to reason. If we are worried – as Mackie was – that our ethical practices are not all they are cracked up to be, it is not clear why the knowledge that certain values or commitments are inescapable should comfort us. Our hope to vindicate ethics consists of more than the desire to know that we are all on the same page; we also want to be assured that we are on the *right* page. And it is not at all clear how we could ever be assured of this simply by appealing to the inescapability of some commitment or aim. There thus appears to be a lacuna in the Kantian strategy.

To see this lacuna most clearly, consider a version of the Kantian strategy according to which what is inescapable for agents is some normative *belief* – say, the belief that we ought to respect one another’s humanity. Would knowing that this normative belief is *inescapable* provide us with any reason to conclude that it is *correct*? Would such knowledge help to allay concerns about the *objective validity* of our practice of judging that we ought to respect one another’s humanity? I cannot see how it would, any more than I can see how learning about the inescapability of some belief regarding the influence of celestial bodies would assure us of the objective validity of some form of astrology. For the standard of correctness for belief is *truth*, and showing that a belief is inescapable does not in any way establish that it is true. That we all believe that *p* is perfectly and obviously consistent with *p*’s being false, for there is always the possibility that we are victims of some sort of collective delusion. Adding that the belief in question is *inescapable* – that we all have this belief merely because we are agents – if anything only strengthens the suspicion that we have all been duped. As William James (1948) observed, it seems rather improbable that the truth should line up so conveniently with our needs and abilities: “In the great boarding-house of nature, the cakes and the butter and the syrup seldom come out so even and leave the plates so clean. Indeed, we should view them with scientific suspicion if they did.”¹⁹

Return now to the Kantian strategy as it is pursued by Korsgaard and Velleman, where what is inescapable is not a belief but rather some practical attitude – a commitment or aim. It is not clear how shifting from an inescapable *doxastic* attitude to an inescapable *practical* one makes matters any better vis-à-vis worries about the validity of our ethical practices. That we are all necessarily committed to respecting one another’s humanity seems to be perfectly consistent with its not being the case that we ought to respect one another’s humanity. We might, in other words, be victims of a collective “practical” delusion, helplessly devoted to pursuing something which actually ought not to be pursued. As it stands, then, the Kantian strategy seems to

¹⁹ James (1948: 103).

provide us with no reason to rule out the possibility that the *correct* or *valid* practical commitments are different from the commitments which are inescapable – or even to rule out the possibility that there are no correct or valid practical commitments at all. The strategy therefore appears to be incomplete. It is difficult to see how premises about the inescapability of some attitude could, by themselves, entail any normative conclusions.

Perhaps such premises are not meant to work by themselves, though. A natural way to augment them would be with some metaethics. We might, for example, be tempted to interpret Korsgaard and Velleman as relying on an unarticulated account of the semantics of normative judgments according to which the truth conditions of such judgments are linked to facts about inescapability. Such an account certainly looks like the sort of thing that could fill the lacuna in their Kantian approach and take us from premises about inescapability to normative conclusions. Alternatively, we might think of Korsgaard and Velleman as embracing an expressivist theory of normative judgments according to which normative judgments are essentially practical, expressing conative states such as plans or desires rather than cognitive states like belief.

I shall have more to say about such metanormative possibilities below. Here let me simply observe that attributing such familiar metanormative theories to Korsgaard and Velleman is not the most natural way to interpret them, given that both evince little interest in metanormative questions and occasionally even show some disdain for metaethics, at least as it is standardly practiced. Korsgaard, for instance, describes her view as a “genuine alternative” to traditional metanormative theories – one that “stands behind the debate between cognitivists and noncognitivists.”²⁰ Given that the distinction between cognitivism and noncognitivism is generally regarded to be exhaustive with respect to metanormative theories, it is reasonable to infer from this description that Korsgaard intends to eschew metaethics. Velleman, meanwhile, advertises his Kantian approach as one that can ground the objective validity of ethics while avoiding the metanormative entanglements of realism. “We don’t need” to posit an ethical pocket in the fabric of the world in order to capture the normative authority of certain demands: these demands are “objectively valid [simply] by being inescapable, in the sense that any agent must accept them.”²¹ Yet Velleman does not frame his theory as a metanormative alternative to realism. Instead, he claims to be offering an alternative to metaethics in general, acknowledging that his account of the foundations of ethics “may strike some philosophers as no metaethics at all.”²² This should not be surprising, given that he has “doubts about the way metaethics is generally practiced”: “Analyzing the semantics, metaphysics, and

²⁰ Korsgaard (2008: 310).

²¹ Velleman (2009: 116).

²² Velleman (2009: 4).

epistemology of ‘Lying is wrong’ will,” Velleman maintains, “reveal very little about the fundamental nature of morality.”²³

Comments such as these support Darwall, Gibbard and Railton’s suggestion that the Kantian strategy involves an attempt to transcend or bypass traditional metaethics. In light of this ambition, it is tempting to conceive of the strategy as an exercise in purely *normative* ethics – albeit normative ethics of a notably fundamental or foundational sort. The suggestion that Korsgaard and Velleman are engaged only in normative maneuvers is certainly consistent with much of what they say about their respective projects. Korsgaard, for example, insists that the question she is trying to answer is a *normative* question—“*the* normative question,” in fact. Are the demands of morality justified? Is there anything that I really ought to do? Metaethics – traditionally understood to operate from a standpoint *outside* of normative discourse – seems ill-suited to answer questions which are so clearly *internal* to that discourse. If we are asking normative questions, then surely what is called for is a *normative* argument or theory.

It may seem obvious that a normative theory – even a foundational one – could never allay broad, Mackie-style worries about our ethical practices. Such worries involve doubts about *all* of our normative judgments, and it is unclear how we could ever vindicate our practice of making normative judgments by making more normative judgments. We might therefore be tempted to conclude that the Kantian strategy pursued as an exercise in purely normative ethics can never get off the ground.

Such a conclusion would be premature, though. The line of thought that yields it may not do justice to the sort of normative theory Kantians aspire to produce. The goal of the Kantian strategy is to ground the norms of practical reasoning in *nonnormative* foundations. In Velleman’s words, the point of “trying to identify a constitutive aim of action is to find a nonnormative foundation for our norms of practical reasoning.”²⁴ The ambition to derive normative conclusions from nonnormative premises is also evident in the work of another famous proponent of the Kantian strategy, Alan Gewirth. At the outset of *Reason and Morality*, Gewirth promises to provide a “logical derivation of a substantial normative moral principle from the nature of human action.”²⁵ His goal is thereby to show “how ‘ought’-judgments ... can be logically and non-circularly derived from ‘is’-statements which describe . . . facts about the world.”²⁶ Setting aside for the moment any doubts we might have about whether an “ought” can be derived from an “is,” we can see how this Kantian project – were it to succeed – would dispel general doubts about normative discourse. As Gewirth and Velleman envision this project, normative conclusions are supposed to follow merely from premises in the philos-

²³ Velleman (2009: 157).

²⁴ Velleman (2004: 287).

²⁵ Gewirth (1978: x).

²⁶ Gewirth (1973-74: 46).

ophy of action. These premises – because they are nonnormative – fall outside the scope of any general worries we might have about normative discourse. And any conclusions we derive from such premises will be as secure as the premises themselves. So, without providing or presupposing a particular semantics or metaphysics of normative discourse – without providing a metaethic, in other words – the Kantian strategy could show that some of our normative judgments are objectively valid or correct and thus that our normative practices are sound. Of course, the strategy would not tell us where the skeptical arguments that give rise to such worries in the first place go wrong, but it would assure us that they must go wrong somewhere.

Now, however, those doubts we momentarily set aside must reclaim our attention: How are we supposed to derive normative conclusions from purely nonnormative premises without any metanormative assistance? How, in other words, are we supposed to derive an “ought” from an “is”? Hume famously argued that such derivations are “altogether inconceivable.”²⁷ Yet even some philosophers who think that Hume was on to something take the Kantian strategy to be a special case. These philosophers accept the rule that we cannot derive an “ought” from an “is” but maintain there are exceptions – special conditions that license this otherwise specious sort of inference, conditions such as *unthinkability* (Korsgaard), *inconsistency* (Gewirth) or a certain kind of *teleology* (Velleman).²⁸ We need not consider these arguments in detail here, though, since various critics of the Kantian strategy have already shown that even appeals to such special conditions fail to support the relevant normative conclusions.²⁹ Either these appeals leave the inference from nonnormative premises to normative conclusion ungrounded (in which case the lacuna remains), or they achieve the desired derivation but only by relying on some substantive normative assumption (in which case they deviate from the Kantian agenda). Both outcomes leave the Kantian strategy unable to dislodge Mackie-style worries about our normative practices.

This failure confirms what we already suspected: we cannot derive normative conclusions from nonnormative premises about agency without relying on metanormative assumptions concerning the semantics or metaphysics of normative discourse. In retrospect, the failure of this purely normative understanding of the Kantian strategy should not be surprising. Despite her apparent lack of interest in traditional metanormative questions, Korsgaard has repeatedly and explicitly presented her account of the foundations of

²⁷ Hume (2000/1739: 302).

²⁸ For Korsgaard’s appeals to *unthinkability*, see (1996: 258) and (2008: 61-62). For Gewirth’s appeals to *inconsistency*, see (1978: 150-60). For Velleman’s appeal to teleology, see (2009: 125, 136), as well as (2000: 176-82, 246).

²⁹ For two related and particularly insightful discussions of Korsgaard’s move from inescapability to normativity, see Fitzpatrick (2005) and Enoch (2006). For criticisms of Gewirth’s attempt to derive an “ought” from an “is,” see Stohs (1988), Allen (1982) and Adams (1980). Finally, for criticism of Velleman’s appeal to *aims*, see Silverstein, “Normativity from Agency.”

ethics as an alternative to *realism* – not to the relatively thin or “procedural” thesis that there are correct answers to normative questions, but rather to the “substantive” and straightforwardly metanormative view that the existence of normative facts together with our knowledge of them is what explains our ethical practices.³⁰ Moreover, the Kantian strategy – at least as it is outlined by Williams and implemented by Korsgaard and Velleman – is supposed to do more than provide a compelling normative argument for a desired ethical conclusion. It is supposed to unearth the sources of normativity, to explain how there could be such a thing as normativity – how there could be considerations that are genuinely authoritative or normative judgments that are objectively valid. It should be clear that no purely normative theory could ever accomplish this.

Kantians need not despair, though. The lesson we should draw from these failures is not that the Kantian strategy is doomed, but rather that it is a mistake to interpret that strategy as merely an exercise in normative ethics. The Kantian appeal to inescapability cannot stand on its own. When we combine that appeal with a traditional metanormative theory, however, the result is something much more potent.

2. Kantian Expressivism

Perhaps the most natural metanormative complement to the Kantian strategy is expressivism. Expressivism, as I understand it, consists of two central claims: first, that we should understand the meaning of normative statements in terms of the states of mind they express, and second, that those states of mind are primarily or essentially conative rather than cognitive.³¹ According to expressivism, then, normative thought is fundamentally *practical* in nature. For if the attitudes expressed by normative statements are essentially conative, then we cannot grasp the nature of normative thought by approaching it as though it were a kind of theoretical reasoning about the state of the world. If expressivism is correct, we should – at least initially – conceive of normative questions not as questions about what is the case, but rather as questions about what to do, think or feel.³²

One reason expressivism seems like a natural partner for the Kantian strategy is that proponents of that strategy – and Velleman and Korsgaard, in

³⁰ See Korsgaard (1996: 35-48) for her account of the distinction between “procedural” and “substantive” normative realism and her unequivocal rejection of the latter. Compare Velleman (2009: 115-17, 139). See Hussain and Shah (forthcoming: 2-3) for a similar observation about Korsgaard’s opposition to realism.

³¹ Note that I have defined expressivism so that it is consistent with but separable from the quasi-realism with which it is now so frequently associated. If the states of mind expressed by normative statements are only *primarily* or *essentially* conative, then that leaves open the possibility that they are *also* cognitive (as the quasi-realist maintains) and thus candidates for truth and potential objects of knowledge.

³² Compare Blackburn (1998: 69-70) and Gibbard (1990: 75) and (2003: 10).

particular – often make claims that would be quite at home in any expressivist treatise. Velleman at least flirts with expressivism, for instance, when he observes that his Kantian approach “forgoes the step of transposing demands into indicative judgments” and instead “leaves them in their practical form, as demands – or at least as practical thoughts of some kind, such as aims or commitments or projects.”³³ Korsgaard’s work displays a more prominent expressivist streak. In *The Sources of Normativity*, she repeatedly denies that normative questions are “third-person, theoretical” questions about the world and argues instead that they are “first-person,” practical questions which arise only for an agent deliberating about what to do.³⁴ Accordingly, ethical reasoning is just “reflection about what to do, not reflection about what is to be found in the normative part of the world.”³⁵ Elsewhere Korsgaard reiterates this thought: “Practical philosophy ... is not a matter of finding knowledge to apply in practice. It is rather the use of reason to solve practical problems.”³⁶ It would be difficult to find a more textbook statement of the motivation behind expressivism (hence one of its earlier appellations: *noncognitivism*).

A second (and more important) reason to partner the Kantian appeal to inescapability with expressivism is that the product is a view with considerable promise. To see this, first recall that one point of the Kantian strategy, at least as it is described by Williams and appropriated by Korsgaard and Velleman, is to vanquish the broad skeptical worry that normative discourse is somehow based on an illusion and that consequently nothing really matters. As we have already seen, to the extent that this worry is grounded, as it is for Mackie, in external, metanormative claims about the semantics and metaphysics of ethics, we cannot allay it simply by appealing to further normative judgments. All such judgments fall within the disputed discourse and are thus no more credible than the judgments they are meant to vindicate. This is why the Kantian strategy conceived as an attempt to overcome external doubts about normative discourse through purely normative ethics seems hopeless.

But now let us add an expressivist theory of normative discourse to the mix. According to expressivism, it is a mistake to understand normative judgments, at least initially, as attempts to represent the state of the world. Rather, we should understand them as essentially practical commitments. Even without filling in the details of any particular expressivist theory, we can see that the addition of expressivism places the Kantian strategy in a much better position to dismiss broad external worries about normative discourse. For if ethical statements express essentially practical (or otherwise noncognitive) attitudes, they do not presuppose the dubious metaphysical commitments that invite Mackie’s skepticism. If some form of expressivism

³³ Velleman (2009: 116).

³⁴ Korsgaard (1996: 16).

³⁵ Korsgaard (1996: 116).

³⁶ Korsgaard (2008: 321).

is correct, then we need not posit the existence of normative facts and properties in order to vindicate normative discourse, and consequently, metaphysical doubts about the existence of such facts and properties are beside the point.³⁷ Normative discourse cannot rest on the sort of collective delusion envisioned by Mackie if our normative judgments are not essentially attempts to represent the state of the world.

Notice, though, that the Kantian strategy itself – the appeal to inescapable commitments or aims – plays no role in this victory over Mackie’s skepticism. Expressivism is doing all of the work. This should not be surprising. Mackie’s worries are essentially metanormative. They arise out of semantic claims about the truth conditions of normative judgments and metaphysical concerns about queer normative properties, and these are the purview of theorists examining our normative practices from the outside. We should therefore expect any solution to Mackie’s challenge to depend on metanormative theorizing.

But if expressivism alone is responsible for the defeat of Mackie-style skepticism, is there any work left for the Kantian strategy to do? Is this really, in other words, a vindication of the Kantian strategy? I believe that it can be. If expressivism is correct, then the only doubts about the authority or objectivity of ethics it makes sense to entertain are doubts *internal to ethical discourse*.³⁸ And the Kantian strategy is well situated to address such doubts, or at least to assure us that they will be addressed satisfactorily.

Let me elaborate. In the face of strong challenges to our fundamental ethical commitments, we might sometimes begin to worry about the status of those commitments. In particular, we might wonder why we cannot say more on their behalf. Of course we can always justify one commitment by appealing to another, but in that direction there looms a potentially vicious regress of justification. This realization can quickly lead to the suspicion that our commitments are not what they should be. This suspicion might be merely local – confined to the justification of certain moral principles, perhaps – or it might be more global. In the extreme, we might throw up our hands and conclude that nothing matters at all. Of course, if we are of a suitably philosophical bent, questions about the status of our ethical commitments will feel pressing even if our confidence in them has not been

³⁷ A quasi-realist version of expressivism might be committed to the existence of normative facts and properties. Quasi-realists, however, deny that their metaphysical commitments, such as they are, are of the sort threatened by Mackie-style arguments. Instead they insist that normative facts and properties are merely the minimalist shadows of true normative propositions. For an argument that quasi-realists cannot have their metaphysical and epistemological cake and eat it too, see Street (2011).

³⁸ *Doubt* is a cognitive notion, and it is natural to wonder whether there is room for normative doubt in a theory according to which normative statements express essentially conative states of mind. Here I simply assume that some form of quasi-realism can provide expressivists with the resources they need in order to make sense of the idea of normative doubt. I am indebted to Nishi Shah for helpful discussion of this point.

shaken at all. I need not doubt that murder is wrong in order to wonder – in a philosophical frame of mind – what makes it wrong.

Whether these questions about our ethical commitments stem from nagging doubts about the status of those commitments or merely from idle philosophical curiosity about the particulars of their justification, the important point is that the questions – and the worries that might inspire them – are internal to our practice of making normative judgments. They are, in other words, *normative* questions. They consist of questions or doubts about whether our commitments are the right ones to have, not doubts about whether the practice of having commitments is confused or unsound. Put metaphorically, they express doubts about whether the moves we are making in the normative game are the correct moves to make, not worries that there is something fundamentally amiss with the game itself. And because these questions are internal to normative discourse, the Kantian strategy can help us answer them.

If, for example, you are wondering whether you ought to put any stock in the norms of morality, then showing you that certain moral conclusions follow from commitments you already have is an excellent way of demonstrating that you should. This is how normative ethics is usually practiced: we try to convince one another that our normative theories are correct by arguing that they follow from more basic normative principles we all share. The Kantian strategy offers us a way to make such arguments with a level of confidence we might not otherwise have, since the strategy assures us that there are at least some commitments all of our potential interlocutors must share. If we can identify what those inescapable commitments are – and if we can show that certain conclusions follow from them – then we can engage in ethical debates certain that our opponents will not be able to escape the force of our arguments by claiming to cleave to different fundamental values.

But what if your doubts extend even to those foundational, inescapable commitments? Suppose you wonder, for instance, whether you should value even your own humanity. If Korsgaard is correct, any deliberation on this question will necessarily lead you to conclude that you should. This is just the sense in which the commitment is inescapable. But is this enough to restore your confidence in that commitment? Does it follow from this inescapability that your humanity is actually valuable? Above, when we were examining purely normative versions of the Kantian strategy, we suggested that it does not follow. In that case our reason for rejecting the inference was that there is no obvious connection between the structure or logic of practical reasoning on the one hand and the truth of our normative judgments on the other. However, this is the right way to think about the Kantian strategy reason *only if we assume that our normative statements express beliefs*. If expressivism is correct, then you cannot step outside the practical point of view and still ask meaningful questions about whether your humanity really is valuable. (That would be, as Korsgaard suggests, akin to trying to see the colors I see by cracking open my skull.) Questions about the value of your humanity have signifi-

cance only inside normative discourse, and once you step back inside that discourse, your only way to answer these questions is to engage in precisely the sort of reasoning that will necessarily, according to Korsgaard, yield the conclusion that your humanity is valuable. Even the most serious doubts must give way.

Let me put the point another way. If expressivism is correct, there is something wrong with the question that purely normative versions of the Kantian strategy cannot answer. There is something misguided about asking whether the fact that we inescapably value humanity entails that humanity is valuable. This question presupposes that the Kantian strategy aims to deliver a deductive argument with inescapability as a premise and some normative judgment as a conclusion. But that is not the right way to conceive of the Kantian strategy, at least when it is paired with expressivism. The question whether humanity is valuable is a normative one, and if expressivism is correct this question is answerable only from the practical or first-person point of view. Yet from that point of view what we see is *the value of the end itself*, not ourselves valuing it. The inescapability of our commitment plays no part in the expressivist version of the Kantian argument: the premise to which we appeal is not *that we are inescapably committed to valuing our own humanity*, but simply *that our own humanity is valuable*. And from this premise it obviously does follow that we ought to value our own humanity. Thus, when we combine the Kantian strategy with expressivism, what justifies the Kantian's normative conclusion is the value to which we are inescapably committed, not the fact that we are inescapably committed to it. The fact of inescapability is something we observe from outside the practical point of view. It merely reassures those of us standing on the sidelines of the normative game what we will find justified once we start playing again. You might still be tempted to object that this does not prove that humanity is in fact valuable. But if expressivism is correct, then the only thing you can be wondering when you raise this objection is whether to value humanity. And that is a practical question.³⁹ Korsgaard encourages you to ask it, because she knows how you must answer it.

Recall that the point of the Kantian strategy is to secure the *objective validity* of at least some of our ethical judgments. Mackie casts doubt on their objectivity by arguing that ethics is the sort of business which requires metaphysical backing in order to be solvent. Since, according to Mackie, the only possible metaphysical backers are normative properties too queer to be credible, we should invest ourselves in some practical alternative to ethics. If expressivism is correct, though, then the enterprise of ethics requires no metaphysical underwriting, and doubts about its solvency are therefore misplaced. Yet this victory over Mackie's skepticism may not relieve all of our doubts about the objectivity of our values. We might still wonder, for instance, what makes our commitments any more valid than those of the many people with

³⁹ Compare Gibbard (1999: 154).

whom we disagree about ethical matters. The crucial point is that – if expressivism is correct – in wondering this we can only be engaged in straightforward normative reasoning, reasoning internal to the practice of ethics. And this is where the Kantian strategy comes into play. For if the Kantian argument that certain commitments are inescapable is sound, then straightforward normative reasoning always leads us back to those commitments. The only type of reasoning available to us necessarily yields the conclusion that these commitments are the correct ones. And that is the sense in which they are objectively valid. Of course our ethical commitments will not have the kind of objectivity that would come with the metaphysical backing of ontologically prior normative properties and facts, but that sort of objectivity was never more than a chimera, at least if expressivism is true. Instead they will have the only kind of objective validity commitments can have: the kind we cannot help but see whenever we engage in practical deliberation.⁴⁰

If the foregoing is correct, then the Kantian argument about inescapability and the expressivist account of normative discourse are both doing crucial work in this expressivist version of the Kantian strategy. From ex-

⁴⁰ Kantian expressivists might be tempted to argue for objective validity somewhat differently. As we noted above, even after Mackie's external skepticism about ethics has been defeated, we may continue to worry that our values are merely subjective. We may, for instance, fear that those values are largely the product of our own peculiar upbringing and biases. And in the face of such fears, we might try to employ the Kantian strategy as follows:

My values cannot be *merely* the product of my upbringing, since those values are inescapable. In other words, my values cannot be subjective if they are necessarily shared by all agents. And if they are not subjective, they must be objective. Therefore, my values are objectively valid.

The problem with this line of argument is that it confuses *objective validity* with *universality*. The fact that some commitment is inescapable entails that every agent shares that commitment, and thus that the commitment is universal. It does not entail, though, that the commitment in question has any kind of *authority* or *validity*. This is the same lacuna we encountered above when we examined the purely normative forms of the Kantian strategy: It remains unclear why I should take comfort in the fact that my values are shared by everyone else, or even in the fact that they are necessarily so shared. That my values are inescapable seems to be no guarantee that they are the correct values to have.

This line of argument runs into trouble because it abandons the first-person, practical point of view that is so central to the expressivist version of the Kantian strategy. This leads me to suspect that Kantian expressivism may not actually be a viable route for Velleman, despite my suggestion that we can read him as an expressivist. Velleman contends that our inescapable commitment takes the form of a desire or drive, yet he resists the thought that being motivated toward something in this way necessarily involves valuing it (Velleman (2000) and (2008: 410-11)). But if we can desire or aim at something without valuing it, then the fact that action constitutively aims at self-understanding tells us nothing about how our normative reasoning must play out – even reasoning about whether self-understanding is worth pursuing. It tells us nothing about which values or commitments we will embrace within the practical point of view. Of course, if Velleman is right about the nature of agency, the aim of self-understanding will always govern such reasoning, but there is no guarantee that this aim will guide us to some particular normative conclusion.

pressivism it follows that there is only one way we can meaningfully entertain doubts about the value of humanity – namely, from within the practical point of view. And according to Korsgaard’s Kantian argument, this way of entertaining doubts always and necessarily results in our repudiating them. Whenever we take up the question of whether to value humanity, we inevitably arrive at the same answer: we should. If at this point you wonder why it follows from this Kantian story about inescapability that humanity is in fact valuable, you are just engaging once again in the very process the story describes. And if the Kantian story is correct, you will arrive at the same answer. Essentially, the conclusion that humanity is valuable is the only one available to you, for there is no way to step outside of the practical standpoint and ask whether it is *really* the right answer.

This is indeed a vindication of the Kantian strategy, then, although it will not seem so to philosophers whose commitment to that strategy is rooted in a desire to bypass the quagmire of traditional metaethics. A full articulation and defense of Kantian expressivism would drive us headlong into that quagmire, where we would soon confront the Frege-Geach problem, the problem of creeping minimalism, and other familiar and pressing objections to expressivism. Philosophers who find expressivism unpalatable need not abandon the Kantian strategy, though: there is more than one path through the morass of contemporary metaethics for Kantians to follow.

3. Kantian Reductionism

Perhaps the most obvious way to bridge the divide between nonnormative premises about the commitments constitutive of agency and the normative conclusions at which the Kantian strategy aims is via a reductionist account of the truth conditions of normative judgments – in particular, an account that locates those truth conditions in the structure of practical thought or elsewhere in the conditions of agency. Velleman harbored reductionist ambitions in much of his early work, though he has since renounced those ambitions.⁴¹ Gewirth at least flirts with reductionism when he states that normative truth is a function of what agents logically must accept.⁴² Even Korsgaard occasionally gestures in this direction. At one point in *The Sources of Normativity* she proposes that “the normative word ‘reason’ refers to a kind of reflective success.”⁴³ This cryptic remark could be read as an attempt to provide a semantics for talk about reasons. In particular, it suggests a reductionist theory according to which a consideration is a reason for action just in case it figures in or brings about successful reflection on the question of what to do.

⁴¹ For his reductionism, see Velleman (1989: 197-208); for the renunciation, see Velleman (2004: 294-97).

⁴² See Gewirth (1973-4: 60).

⁴³ Korsgaard (1996: 93).

There are popular views in the metaethics literature that comport with these fairly inchoate descriptions. Consider Sharon Street's constructivism. According to Street, "normative truth consists in what is entailed from within the practical point of view."⁴⁴ More specifically, "the fact that X is a reason to Y for agent A is constituted by the fact that the judgment that X is a reason to Y (for A) withstands scrutiny from the standpoint of A 's other judgments about reasons."⁴⁵ Even without introducing the details of Street's constructivist theory, it should be clear how it can bridge the gap between nonnormative premises about our inescapable commitments and normative conclusions about what we have reason to do. If, for instance, we cannot help but value of our own humanity, then the conclusion that our humanity is valuable is entailed from within the practical point of view. That is, the judgment that our humanity is valuable will necessarily withstand scrutiny from the standpoint of our other normative judgments about reasons, since the judgment that our humanity is valuable will necessarily be among those judgments. That is what it means for the commitment to be inescapable. So, when we supplement the Kantian strategy with a metaethic like Street's, we have a clear explanation of how the Kantian strategy can yield the promised normative conclusions.⁴⁶

Once again, though, we might wonder whether this is really a vindication of the Kantian strategy. It may seem, just as it did when we introduced Kantian expressivism in the previous section, that the metanormative half of this new Kantian hybrid is doing all of the heavy lifting, especially when it comes to neutralizing broad, Mackie-style doubts about normative discourse. For if constructivism is the correct metanormative theory, then normative discourse does not presuppose the sort of stance-independent normative properties Mackie finds so queer. It presupposes only that some normative judgments withstand scrutiny from the standpoint of other normative judgments. And that looks both metaphysically and epistemologically unproblematic – nothing queer in sight. Constructivism alone is therefore sufficient to rule out the irrealist hypothesis that our invocations of normative authority inevitably misfire.

Is there any work left for the Kantian appeal to inescapability to do? As before, I believe there is. Street's constructivism initially seems to entail *subjectivism* about reasons for action. Since what A judges or takes to be a reason is presumably a contingent feature of A 's psychology, the normative judgments which withstand scrutiny from the standpoint of A 's other judgments are

⁴⁴ Street (2010: 367).

⁴⁵ Street (2008: 223).

⁴⁶ There are views in addition to Street's that should be able to do the same work. See, for instance, Coleman (unpublished manuscript); Schroeder (2007); and perhaps even Williams (1981). Which metanormative view is most appropriate may depend on what form the inescapable commitment takes, that is, whether it is a desire, an intention, a purpose or – as we have been supposing – a judgment. For a metanormative view that pairs well with a Kantian appeal to constitutive *aims*, see Silverstein, "Ethics and Practical Reasoning."

likely to be different from those which withstand scrutiny from the standpoint of *B*'s other judgments. That is, the normative judgments entailed from within *A*'s practical point of view are likely to be different from those entailed from within *B*'s practical point of view. It would therefore be surprising if all of the considerations that count as reasons for *A* also count as reasons for *B*. Moreover, even claims that are universally valid – even considerations that are normative for all actual agents – will be only contingently so; we will always be able to imagine agents with different normative commitments, agents who would thereby escape the force of what had appeared to be objectively valid normative claims. The picture we are left with is one in which there are no objectively authoritative or valid normative claims: normative claims that are necessarily valid with respect to any agent. Now, however, let us reintroduce into the mix Kantian premises about inescapable commitments. In particular, let us suppose that there are certain considerations any agent must accept as reasons for action. The normativity of those considerations will possess an objectivity that was missing from the constructivist picture, which included only contingent commitments. These considerations will be necessarily authoritative for any possible agent. And so even though we began with a metanormative view that appeared to allow only for subjective reasons for action, combining that view with the Kantian strategy makes room for ethical objectivity.

Once again, we have vindicated the Kantian strategy – except in the eyes of those who hope to secure the objectivity of ethics without recourse to traditional metaethics. But just as Kantian expressivism is vulnerable to various familiar criticisms of expressivism, so too is this second Kantian hybrid open to well-known metanormative objections. To see this, notice that Street's constructivism involves a kind of *reductionism* about normativity. Normative facts turn out to be facts about what is entailed from within the practical point of view, and the sense of entailment at work here does not presuppose any normative notions. As Street puts it, "to explain this sense of entailment, we need not make any substantive normative assumptions – for example, about what anyone *should or ought* to do or infer, or about what counts as a *normative* reason for what."⁴⁷

The fact that Street's view is a form of reductionism means that any take on the Kantian strategy which incorporates her constructivism must confront some variant of G. E. Moore's famous open question argument. Roughly, the Moorean challenge is this: Why should we think that normative facts are as the constructivist describes them? Why should we think that the property of withstanding scrutiny from the standpoint of our other normative judgments is the property we mean to invoke when we make normative claims about what we ought to do? As Street acknowledges, facts about what is entailed from within the practical point of view are a combination of psychological and logical facts. Why should we think that in settling these mat-

⁴⁷ Street (2010: 367).

ters we are also settling the question of what we ought to do? Put another way, why should we think that what we have reason to do is a function of what we take ourselves to have reason to do? Settling the facts about what follows from *what we take our reasons to be* does not seem to close (or even necessarily bear on) the question of *what our reasons actually are*. I can acknowledge that my current normative commitments entail a particular normative conclusion but still wonder whether that conclusion is true. Even after the full set of my normative commitments is laid bare, the question of what I ought to do appears to remain open, and this suggests that normative questions are about something other than what follows from within the practical point of view.

Any reductionist version of the Kantian strategy must answer such questions, and that will inevitably require wading into murky metanormative waters. Of course, just as the Kantian expressivist can hope to solve the problems that plague expressivism, so can the Kantian reductionist offer responses to the objections facing reductionism. And nothing I have argued here suggests that those responses will be unsuccessful. I actually have high hopes for some form of Kantian reductionism. My point here is simply that although Kantians can dodge the many objections to expressivism by embracing reductionism, they cannot thereby circumvent metanormative obstacles altogether.

4. A Third Way?

In section 1 I suggested – following Darwall, Gibbard and Railton – that the most natural way to interpret various remarks made by proponents of the Kantian strategy is as expressing the ambition to bypass traditional metaethics. I then argued that the Kantian strategy so understood cannot succeed: if the Kantian strategy is to be salvaged, the effort to leave metaethics behind must be abandoned. Perhaps, however, I have misunderstood the Kantians' aspiration with respect to metaethics. Perhaps we should interpret their goal to avoid ordinary metaethics not as the desire to eschew metaethics altogether, but rather as the desire to do *extraordinary* metaethics. If that is their ambition, then Kantians will be unsatisfied with the conclusions of sections 2 and 3, for there is nothing extraordinary about expressivism and reductionism in metaethics. We must therefore consider the possibility that the Kantian appeal to inescapability is meant to be paired with a metanormative view distinct from both expressivism and reductionism as well as from realism – a *sui generis* Kantian metaethic.

It is tempting to read Korsgaard as offering such a metaethic, especially when she promises to offer a “theory of normative concepts” that is a “genuine alternative” to the one that “stands behind the debate between cognitiv-

ists and noncognitivists.”⁴⁸ I believe we should resist this temptation. Although Korsgaard *intends* to offer an extraordinary metaethic – a third way – she fails to distinguish her theory from the ordinary metanormative views canvassed above. At times she leans heavily in the direction of expressivism, and at other times her remarks suggest some form of reductionism. At no time, however, does she point toward a view distinct from these two.

As I noted at the outset of section 2, many of Korsgaard’s attempts to articulate her account of the sources of normativity suggest a distinctly expressivist understanding of normative discourse. According to Korsgaard, normative questions are simply questions about “what to think, what to like, what to say, what to do, and what to be.”⁴⁹ Her criticisms of realism corroborate this expressivist reading. She insists that the problem with realism is that it presupposes that our worries about ethics are worries about the loss of ethical *knowledge*. Realists never pause to ask “whether knowledge of ethical objects, or indeed *any sort of knowledge at all*, is really what we want here in the first place.”⁵⁰

Korsgaard sounds most like an expressivist when she is attempting to describe the mistake of asking normative questions from outside the first-person point of view.

Value, like freedom, is only directly accessible from within the standpoint of reflective consciousness. And I am now talking about it externally, for I am describing the nature of the consciousness that gives rise to the perception of value. From this external, third-person perspective, all we can say is that when we are in the first-person perspective, we find ourselves to be valuable, rather than simply that we are valuable. There is nothing surprising in this. Trying to see the value of humanity from the third-person [external] perspective is like trying to see the colors someone sees by cracking open his skull. From the outside, all we can say is why he sees them.⁵¹

Compare this to Simon Blackburn’s expressivist account of where we can find ethical truth.

The answer is that it is not anywhere that can be visible from this sideways, theoretical perspective. It is not that this perspective is illegitimate, but that it is not the one adapted for finding ethical truth. It would be if such truth were natural truth, or consisted of the existence of states of

⁴⁸ Korsgaard (2008: 310). Commentators who appear to read Korsgaard as offering a *sui generis* metaethic include Schneewind, who labels her view “non-reductive naturalized Kantian foundationalism” (1997: 791) and FitzPatrick (2005: 657-60).

⁴⁹ Korsgaard (1996: 9).

⁵⁰ Korsgaard (1996: 47). I am hardly the first reader of Korsgaard’s work to interpret her as an expressivist. In his review of *The Sources of Normativity*, Gibbard lays out Korsgaard’s account of normative concepts and then adds: “This is the doctrine I call *expressivism*” (1999: 141). See also Schroeder (2010: 164) and Hussain and Shah (forthcoming).

⁵¹ Korsgaard (1996: 124).

affairs in the real world. That is the world seen from the viewpoint that sees different and conflicting moral systems – but inevitably sees no truth in just one of them. To “see” the truth that wanton cruelty is wrong demands moralizing, stepping back into the boat or putting back the lens of a sensibility.⁵²

For Korsgaard, just as for Blackburn, normativity is not something we can find (or even talk about) from the third-person, theoretical point of view. Normative talk – talk of values and reasons – comes alive only from the perspective of someone engaged in first-person reasoning about what to do, what to think or what to feel. Only here – “within the standpoint of reflective consciousness” or, as Blackburn puts it, behind the “lens of a sensibility” – are normative judgments appropriate or meaningful.

Korsgaard, of course, repeatedly denies that she is an expressivist. The closest she ever comes to praising expressivism is when she declares that it, like realism, is true, but “only in a way that makes it boring.”⁵³ Given her assessment of realism, this is hardly a heartfelt endorsement. Her reasons for thinking that her view differs from expressivism are not at all compelling, though.

The distinction [between cognitivism and noncognitivism] suggests that a moral judgment either articulates a description of some fact or is a disguised version of some alternative use of language – either expressive or prescriptive. But where does this leave theories like Aristotle’s or Kant’s, according to which moral judgments are the conclusions of practical reasoning? A conclusion of practical reasoning is not obviously a description of a fact about the world, but it hardly seems like some sort of emotional expletive either. Where do these theories fit?⁵⁴

It is not clear why Korsgaard supposes that rejecting cognitivism and thereby denying that ethical judgments are statements of fact requires one to conclude that they are merely emotional expletives. This may have been the view of one famous early noncognitivist – A. J. Ayer – but it would be a mistake to describe the varieties of expressivism fashionable today in such terms. In fact, the genealogy of modern-day expressivism – beginning with Charles Stevenson and proceeding through R. M. Hare to Simon Blackburn and Allan Gibbard – is very much the story of philosophers working to develop an account of ethical judgments according to which those judgments – even though they are not straightforward statements of fact – can nonetheless be employed in deliberation that deserves to be called *reasoning*. Despite Korsgaard’s claim to the contrary, then, it appears that the distinction between cognitivism and noncognitivism leaves plenty of room for theories like Aristotle’s and Kant’s (as she understands them).

⁵² Blackburn (1993: 372).

⁵³ Korsgaard (2008: 325, n. 49).

⁵⁴ Korsgaard (2008: 309).

Another reason Korsgaard might deny that her view is tantamount to expressivism is that she wants to leave room for talk of *truth* and *facts* in the normative realm. She believes that “there are right and wrong ways to answer” normative questions, and that this is sufficient for there to be “normative truth.”⁵⁵ But expressivism also leaves room for such talk, at least when it is paired – as it standardly is – with quasi-realism. There is not space here to lay out the details of any of the various quasi-realist strategies for securing talk of normative truth and objectivity within an expressivist metaethic. The basic idea is familiar enough, though: expressivists can countenance talk of normative truth and objectivity by interpreting appeals to those notions simply as expressions of complex or higher-order but still essentially conative attitudes. One might wonder at this point whether, if quasi-realism is sufficient to license talk of objectivity in the normative realm, there is any work left for the Kantian appeal to inescapability to do. There is. As I argued in section 2, what the Kantian strategy adds to expressivism is not the right to speak of normative truth or objectivity in general, but rather the resources to argue for particular normative truths, such as those concerning the value of humanity.

What ultimately makes the expressivist reading of Korsgaard so attractive is her repeated insistence that – despite the legitimacy of talk about normative truths and facts – the primary function of normative concepts is not to represent the world. Even if we can incorporate normative concepts into what look like straightforward assertions that describe reality, the fundamental role of these concepts is to help us engage in genuinely practical reasoning or, as Korsgaard often puts it, to help us solve practical problems.

Many readers of Korsgaard are nevertheless reluctant to interpret her as an expressivist. Even if such reluctance is warranted, however, we should not conclude that she is offering a *sui generis* Kantian metaethic. We have already found fault with Korsgaard’s reasons for denying the exhaustiveness of the distinction between cognitivism and noncognitivism. And so, if she is not an expressivist, she must be a cognitivist of some variety. What is more, there are good reasons to think that she must be a reductionist as well. How else can we understand her assertion that normative concepts “are the names of” solutions to practical problems and that “the normative word ‘reason’ refers to a kind of reflective success”?⁵⁶ For that matter, how else can we make sense of her confidence that facts about the inescapability of certain normative judgments entail the truth or correctness of those normative judgments? As I argued in section 1, if Korsgaard is a cognitivist, then absent some reductive hypothesis linking normative statements to statements about the logic of practical thought – or linking normative facts to facts about the logic of practical thought – there is no way to move from the claim that we inescapa-

⁵⁵ Korsgaard (1996: 35).

⁵⁶ Korsgaard (2008: 322) and (1996: 93). Commentators who read Korsgaard roughly along such cognitivist lines include Smith (1999: 386) and Street (2010: 366).

bly judge humanity to be valuable to the conclusion that humanity is in fact valuable, at least without the help of some auxiliary normative premise. And for reasons I have already explained, the Kantian strategy cannot accomplish what its proponents hope to accomplish if it incorporates such a premise. Reductionism, then, seems to be the only avenue remaining.

Korsgaard's insistence that normative concepts do not describe "external reality" might lead some to conclude that she cannot be a cognitivist.⁵⁷ But the suggestion that our normative concepts do not describe *external* reality is perfectly consistent with cognitivism, since those concepts might instead describe *internal* reality – that is, the reality of what does and does not survive practical reflection, or of what does and does not solve practical problems. Thoughts employing such concepts would be straightforwardly cognitive, but they would also be able to play the role in practical reasoning Korsgaard reserves for normative judgments. They would, for example, help us think and talk about the solutions to our practical problems.

Another worry one might have about this reductionist reading of Korsgaard's view is that it misses what is distinctive about her metanormative outlook, namely, her claim that normativity is essentially first-personal – that normative force is something which can be felt only from the standpoint of a deliberating agent. Korsgaard might be a cognitivist, the worry goes, but the essence of her view is her distinctly practical account of normativity. This is what constitutes her *sui generis* Kantian metaethic; any reductionist account of normative concepts or properties is at most an afterthought – an offering to metaethicists interested in "boring" questions about the semantics of normative judgments.

The reductionism cannot be an afterthought, though. Korsgaard's distinctive metanormative contribution may very well be her account of the essentially practical nature of normative force. But if Korsgaard is a cognitivist, that account can do the work she wants it to do only if it is paired with a reductionist semantics that links normative judgments to the force we experience as deliberating agents. After all, it is our everyday normative questions – questions about what we ought to do – that Korsgaard seeks to address. And without a reductionist semantics, we have no assurance that the peculiarly practical force Korsgaard has identified is the kind of force we invoke when we ask such questions. Put another way, without the reductionist semantics, we have no way of ruling out the possibility that in focusing on the force we feel from the practical point of view, Korsgaard has just changed the subject. On any cognitivist understanding of Korsgaard's view, then, a reductionist semantics is indispensable.

There is, then, no *sui generis* Kantian metaethic on offer in Korsgaard's work. Nor have we been given any reason to think that one is even available. Does this spell doom for the Kantian strategy? Hardly. If I am correct, any viable version of the Kantian strategy is going to face familiar metanorma-

⁵⁷ Korsgaard (2008: 324).

tive challenges, yet nothing I have written here suggests that those challenges cannot be overcome. My target in this paper is not the claim that the Kantian strategy can help us secure the objective validity of our normative judgments, but rather the claim that it can do so and also stay clear of metanormative debates about the semantics and metaphysics of normative discourse.

5. Conclusion

At the beginning of *The Sources of Normativity*, Korsgaard describes two different standards that any satisfying theory of normative concepts must meet: a standard of *normative* or *justificatory adequacy* and a standard of *explanatory adequacy*. Korsgaard's complaint about traditional metanormative theories is that they are concerned only with explanatory adequacy. They purport to explain our ethical practices from the outside, and they therefore leave the urgent matter of normative adequacy unaddressed. Consequently, they can provide neither comfort nor guidance to people struggling with the question of what they ought to do.

This is precisely what we should expect from traditional metanormative theories. Any traditional metaethicist should tell us that if we are agonizing about what we ought to do or about what we have most reason to do, we should seek guidance not from metaethics but from normative ethics. We should, in other words, step inside our ethical practices and take part in the conversations that constitute them. After all, the question of what we have reason to do is, as Korsgaard repeatedly reminds us, a *normative* question, not a metanormative one. Korsgaard is correct, then, that if we are concerned with the normative or justificatory adequacy of our ethical practices, we need more than traditional metaethics. We need something like her Kantian strategy.

What Korsgaard and other Kantians fail to recognize, though – and what I have been trying to demonstrate here – is that their strategy for showing that our practices meet the standard of normative adequacy can succeed only if Kantian arguments about inescapability are coupled with the right sort of metanormative theory. For it is a metanormative theory concerned with explanatory adequacy that tells us *what we are doing when we are asking about the normative adequacy of our ethical practices* or are otherwise engaged in ethical discourse. Only with the support of a metaethic can we be sure that those Kantian arguments about inescapable commitments are relevant to questions of normative adequacy. If, according to our best metanormative theory, ethical judgments are essentially beliefs about something other than the structure of practical thought or what follows from within the practical standpoint, then Kantian premises about what we inescapably value or constitutively aim at will be entirely beside the point. Matters of normative adequacy will remain pressing, and Kantian arguments will not be able to help us address them.

Kantians should therefore be a bit less cavalier in their attitudes toward traditional metaethics. The success of their strategy for securing the objectivity of ethics hinges on a number of metanormative outcomes. “Boring” or not, the quagmire of traditional metaethics awaits. There is no Kantian route around it.⁵⁸

Matthew Silverstein
New York University Abu Dhabi
matthew.silverstein@nyu.edu

⁵⁸ For comments on earlier drafts of this paper I am indebted to David Enoch, Matt Evans, Mark Schroeder, Russ Shafer-Landau, Peter Railton, Joshua Silverstein and David Velleman, as well as the participants in the Mid-Atlantic Reading Group in Ethics, various anonymous referees and the editors of the *Journal of Ethics & Social Philosophy*. Special thanks to Stephen Darwall and Nishi Shah for many helpful and encouraging conversations and suggestions.

References

- Adams, E. M. (1980) "Gewirth on Reason and Morality," *Review of Metaphysics* 33(3): 579-92.
- Allen, P. (1982) "A Critique of Gewirth's 'Is-Ought' Derivation," *Ethics* 92(2): 211-26.
- Blackburn, S. (1998) *Ruling Passions*, Oxford: Clarendon Press.
- ____ (1993) "How to Be an Ethical Antirealist," in *Essays in Quasi-Realism*, New York: Oxford University Press.
- Coleman, M. C. "Exploring Metanormative Constitutivism," unpublished manuscript.
- Darwall, S., Gibbard, A. and Railton, P. (1992) "Toward Fin de Siècle Ethics: Some Trends," *Philosophical Review* 101(1): 115-89.
- Enoch, D. (2006) "Agency, Shmagency: Why Normativity Won't Come from What Is Constitutive of Action," *Philosophical Review* 115(2): 169-98.
- Ferrero, L. (2009) "Constitutivism and the Inescapability of Agency," *Oxford Studies in Metaethics* 4: 303-33.
- FitzPatrick, W. J. (2005) "The Practical Turn in Ethical Theory: Korsgaard's Constructivism, Realism, and the Nature of Normativity," *Ethics* 115(4): 651-91.
- Gewirth, A. (1978) *Reason and Morality*. Chicago: University of Chicago Press.
- ____ (1973-74) "The 'Is-Ought' Problem Resolved," *Proceedings and Addresses of the American Philosophical Association* 47: 34-61.
- Gibbard, A. (2003) *Thinking How to Live*, Cambridge, MA: Harvard University Press.
- ____ (1999) "Morality as Consistency in Living: Korsgaard's Kantian Lectures," *Ethics* 110(1): 140-64.
- ____ (1998) *Wise Choices, Apt Feelings: A Theory of Normative Judgment*, Cambridge, MA: Harvard University Press.
- Hume, D. (2000/1739) *A Treatise of Human Nature*, D. F. Norton and M. J. Norton, eds., Oxford: Oxford University Press.
- Hussain, N. J. Z. and Shah, N. (forthcoming) "Metaethics and Its Discontents: A Case Study of Korsgaard," in C. Bagnoli, ed., *Constructivism: For and Against*, Cambridge: Cambridge University Press.
- James, W. (1948) "The Will to Believe," in *Essays in Pragmatism*, New York: Hafner Press.
- Korsgaard, C. (2008) *The Constitution of Agency: Essays on Practical Reason and Moral Psychology*, Oxford: Oxford University Press.
- ____ (1996) *The Sources of Normativity*, Cambridge: Cambridge University Press.
- Mackie, J. L. (1977) *Ethics: Inventing Right and Wrong*, London: Penguin Books.
- Schneewind, J. B. (1997) "Review of *The Sources of Normativity*," *Mind* 106(424): 791-94.
- Schroeder, M. (2010) *Noncognitivism in Ethics*, London: Routledge.
- ____ (2007) *Slaves of the Passions*, Oxford: Oxford University Press.
- Setiya, K. (2007) *Reasons without Rationalism*, Princeton: Princeton University Press.
- Silverstein, M. "Ethics and Practical Reasoning," unpublished manuscript.
- ____ "Normativity from Agency," unpublished manuscript.
- Smith, M. (1999) "Search for the Source," *Philosophical Quarterly* 49(196): 384-94.
- Stohs, M. D. (1988) "Gewirth's Dialectically Necessary Method," *Journal of Value Inquiry* 22: 53-65.
- Street, S. (2011) "Mind-Independence Without the Mystery: Why Quasi-Realists Can't Have It Both Ways," *Oxford Studies in Metaethics* 6: 1-32.
- ____ (2010) "What Is Constructivism in Ethics and Metaethics?" *Philosophy Compass* 5(5): 363-84.
- ____ (2008) "Constructivism about Reasons," *Oxford Studies in Metaethics* 3: 207-45.
- Velleman, J. D. (2009) *How We Get Along*, Cambridge: Cambridge University Press.
- ____ (2008) "A Theory of Value," *Ethics* 118(3): 410-36.
- ____ (2004) "Replies to Discussion of *The Possibility of Practical Reason*," *Philosophical Studies* 121(3): 277-98.
- ____ (2000) *The Possibility of Practical Reason*, Oxford: Clarendon Press.
- ____ (1989) *Practical Reflection*, Princeton: Princeton University Press.

- Williams, B. (1995) "Ethics and the Fabric of the World," in *Making Sense of Humanity*, Cambridge: Cambridge University Press.
- (1981) "Internal and External Reasons," in *Moral Luck: Philosophical Papers 1973–1980*, Cambridge: Cambridge University Press.