



Philosophical Review

Agency, Shmagency: Why Normativity Won't Come from What Is Constitutive of Action

Author(s): David Enoch

Source: *The Philosophical Review*, Vol. 115, No. 2 (Apr., 2006), pp. 169-198

Published by: [Duke University Press](#) on behalf of [Philosophical Review](#)

Stable URL: <http://www.jstor.org/stable/20446897>

Accessed: 07/01/2014 13:09

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at

<http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Duke University Press and *Philosophical Review* are collaborating with JSTOR to digitize, preserve and extend access to *The Philosophical Review*.

<http://www.jstor.org>

Agency, Shmagency: Why Normativity Won't Come from What Is Constitutive of Action

David Enoch

Hebrew University, Jerusalem

Why care about objective value or ethical reality? The sanction is that if you do not, your inner states will fail to deserve folk-theoretical names. Not a threat that will strike terror into the hearts of the wicked! But whoever thought that philosophy could replace the hangman?

—David Lewis

1. Introduction

There is a fairly widespread—and very influential—hope among philosophers interested in the status of normativity that the solution to our metaethical and, more generally, metanormative problems will emerge from the philosophy of action. What we need, so the thought goes, is to better understand what action and agency consist in. With this understanding in hand, we will then have all that we need in order to give a philosophical account of reasons for action, or of (practical) normativity altogether.

For helpful conversations and comments on previous drafts, I thank Antony Duff, Yuval Eylon, Ido Geiger, Pete Graham, Alon Harel, Andrei Marmor, Sandra Marshall, Andrew McGonigal, Seiriol Morgan, Josh Schechter, Assaf Sharon, Ruth Weintraub, and anonymous referees for the *Philosophical Review*. Previous versions of this essay were presented in a workshop on the philosophy of action at the University of Cyprus in June 2005 and at the annual conference of the British Society for Ethical Theory in July 2005. I thank the participants for the helpful discussions that followed.

Philosophical Review, Vol. 115, No. 2, 2006

DOI 10.1215/00318108-2005-014

© 2006 by Cornell University

The intuitive idea can be put, I think, rather simply: In order to know what it takes for a car to be a good car, we need to understand what cars are, what their constitutive functions are, and so on. A good car is just a car that is good *as a car*, good, that is, in measuring up to the standards a commitment to which is built into the very classification of an object as a car. Analogously, then, perhaps in order to know which actions are good (or right, or reason supported, or rational, or whatever), all we need is a better understanding of what actions are, or perhaps of what it is to be an agent, someone who performs actions. Perhaps the normative standards relevant for actions will fall out of an understanding of what is constitutive of action just as the normative standards relevant for cars fall out of an understanding of what is constitutive of cars.

The hope, I think, has both a normative and a metanormative part. On the normative level, the hope is that we can find out just which standards are the right ones by deriving them from what is constitutive of action. In its most ambitious form, the hope is to derive a full first-order normative (or perhaps just moral) theory from what is constitutive of action. On the metanormative level, the hope is that an attractive second-order theory of normativity (or perhaps just of morality) can be developed starting with the insight that practical normativity is, in some sense, grounded in what is constitutive of action,¹ a metanormative theory that shows practical normativity (and perhaps morality) to be on as solid a ground as the normativity of recommendations about cars.

In this essay, I will argue that these hopes are groundless. I will focus on the metanormative hope, but—as will become clear—showing that the solution to our metanormative problems will not come from what is constitutive of action will also devastate the hope of gaining significant insight into first-order, normative truths by focusing (only) on what is constitutive of action.² Or so, at least, I shall argue.

In the next section, I briefly survey some recent attempts to ground normativity in what is constitutive of action, emphasizing the philosophical motivations underlying them. In section 3, I present the problem that I take such theories to face. The problem is not unex-

1. In some contexts it may be important to distinguish between attempts to ground normativity in what is constitutive of *action* and attempts to ground normativity in what is constitutive of *agency* (see Setiya 2003, 382). Such a distinction will not be relevant for what follows, as my argument applies equally to both attempts.

2. But see note 38, below.

pected—indeed, it is nicely summarized in the quote from David Lewis I started with. And two of the constitutivist theorists whose views I discuss anticipate the problem and try to address it. In section 4, I reject the ways of addressing the problem that can be extracted from Connie Rosati's (and to an extent also David Velleman's) quick remarks on it. In section 5, I argue that no other way of addressing the problem is likely to do much better. After rejecting another possible attempt at addressing the general problem (in section 6), I conclude (in section 7) that normativity will not come from what is constitutive of action. (Nor, I think, will epistemic normativity come from what is constitutive of beliefs, for reasons similar to the ones I put forward in what follows. But I will constrain myself to a discussion of actions and practical reasons.) This conclusion notwithstanding, I believe there is something both true and important about the insights incorporated in the theories I criticize. In the final section, I present a sketch of an alternative way of accommodating the truth in those insights, a way that I develop at length elsewhere (Enoch 2003 and forthcoming).

2. The Hope and Some Underlying Motivations

2.1 Korsgaard's Reply to the Skeptic

Action, so Christine Korsgaard tells us, is self-constitution.³ That your action (partly) constitutes yourself, or perhaps that in your action *you* constitute yourself, is constitutive of what an action is. From this, Korsgaard (1.3.4) continues, an important result follows: "Action is self-constitution. And accordingly, . . . what makes actions good or bad is how well they constitute you." Korsgaard believes that the whole of morality—and, indeed, the whole of practical rationality, and perhaps even rationality more generally⁴—can be extracted from this insight.

3. In this essay I focus almost exclusively on Korsgaard's view as she presents it in her Locke Lectures (2002). For the most part, I will leave out of the picture Korsgaard's theses in *The Sources of Normativity* (1996). I will do so, first, because the relation between the views expressed in the Locke Lectures and in *Sources* is not entirely clear to me, and second, because for the purposes of this essay, the Locke Lectures fit better as an example of the kind of view I will be arguing against. For a discussion of Korsgaard's views in *Sources* that attempts to read them as a constitutivist theory, see Setiya 2003, 381 and O'Hagan 2004.

All references to Korsgaard will be to the Locke Lectures (2002) unless otherwise noted. I will refer to the lectures by section number.

4. In Korsgaard n.d., 4, Korsgaard is clear about applying the constitutive-aim strategy to epistemic norms as well.

Let us postpone for another occasion doubts about self-constitution—whatever exactly is meant by this phrase—as a constitutive feature of action. Still, we may ask: What is the point of grounding the goodness or badness of actions in what is constitutive of action? Korsgaard (1.4.1) clearly thinks that at least one advantage of this constitutivist strategy is that it promises to cope with a (or perhaps even the) skeptical challenge: “The idea of a constitutive standard is an important one, for constitutive standards meet sceptical challenges to their authority with ease.” If someone asks why they should care that their action is morally impermissible, or indeed that it is irrational, just insisting that they should care because it is immoral or irrational will be, so Korsgaard thinks, pointless. But if someone is building a house and if sheltering from the weather is a constitutive standard of being a house, then she cannot sensibly ask: Why should I care if the house I’m building cannot—because of my shoddy work—shelter its inhabitants from the weather? (1.4.2) Skeptical challenges are rarely raised about the normative status of standards that apply to the proper building of houses, and when they are, they can be met⁵ fairly easily and confidently. If skeptical challenges of this kind to morality or to practical rationality can be shown to be on a par with skeptical challenges to house-building norms, it seems that progress has been made: Skepticism about morality and rationality can be seen to be as misconceived as skepticism about house-building norms.

2.2 Velleman’s *Quasi-Externalism*

David Velleman believes that all actions are caused by beliefs and desires. What is unique to actions—or at least to full-blooded actions—is that they are caused by special desires. The special desire most relevant here is the desire to know ourselves, or (viewed in the other direction) to be

5. Throughout her work—both in *Sources* (1996) and in the Locke Lectures (2002)—Korsgaard systematically equivocates on locutions such as “addressing the sceptic” (for example, 1996, 16). It seems like a reasonable requirement for a theory of normativity that it should defeat the skeptic, that is, that it should show him to be wrong. It is clearly *not* a reasonable requirement that the skeptic be necessarily converted, and, of course, Korsgaard’s own theory doesn’t pass this test (I return to this point later on). By conflating the two—defeating and convincing the skeptic—Korsgaard can employ, in criticizing realism, the stricter requirement with the perceived plausibility of the looser one. For her employment of the stricter requirement against realism, see Korsgaard 1996, 34. For relevant criticism, see Gibbard 1999, 146 and O’Hagan 2004. And for a similar point made in a wider philosophical context, see Hookway 1999, 174–75.

intelligible to ourselves.⁶ This desire kicks in whenever we form a belief about our own future behavior, thereby giving us reason to act as we believe we will act. This aim of self-knowledge or of “having an integrative knowledge of what we are doing” is constitutive of action, and is—or at least can be—the source of many, perhaps all, of our practical reasons (Velleman 2004a, 234).⁷

Postponing doubts about self-knowledge as an aim of action, why insist that this aim is constitutive of action and that it is the source of all our reasons for action? Velleman too, I think, believes that employing such a strategy has antiskeptical implications.⁸ But Velleman has another reason to ground practical normativity in what is constitutive of action, a reason that will be important in what is to come. Velleman (2000e [1998]) is an internalist of sorts about reasons. He (2004b, 279) thinks that it is necessary for one’s having a reason to act in a certain way that one be motivated—perhaps under suitable conditions—to act in that way; at least, “justifying reasons must be capable of entering into explanatory reasons.” But Velleman also recognizes the unattractiveness of too extreme an internalist position, one that renders a person’s (normative) reasons as unstable as—and dependent upon—his desires. So Velleman embarks on what may be thought of as a quasi-externalist⁹ project: He tries to show that even though there is a necessary connection between normativity and motivation, still the externalist—who denies such a connection—can have almost all she wants. In particular, she can have normative reasons that are not as unstable and subjective as garden-variety desires. This is so because some reasons are grounded in the desire constitutive of agency, the desire for self-knowledge.¹⁰ Seeing that no

6. This is a central theme of Velleman 1989. For a more recent statement (or perhaps a restatement) of this view, see Velleman 2000b.

7. At one time, Velleman was happy with thinking of autonomy as the constitutive aim of action. Later on, though, he came to believe that this is not a useful way of putting things. See Velleman 2000b, 30.

8. This is perhaps clearest when Velleman tries to ground *epistemic* reasons in the constitutive aim of *belief*. See Velleman 2000c, and for the analogy between the belief and the action cases see Velleman 2000d [1996].

9. Simon Blackburn’s quasi-realist project—from which I borrow this way of speaking—is, roughly speaking, an attempt to give the moral realist all that she wants but on fundamentally antirealist metaphysical assumptions. See Blackburn’s introduction to his 1993. Velleman’s quasi-externalist motivation is clearest in his 2000d [1996]. On this point, see also O’Hagan 2004, 49.

10. Velleman changed his mind on whether the constitutive aim of action must be an aim of the agent (and so perhaps the object of one of her desires) or whether it

one can even be an agent without having this desire, reasons grounded in it are guaranteed to be universal and not contingent on the specific desires you happen to find yourself with. The grounding of normativity in what is constitutive of agency thus allows the internalist to avoid some of the counterintuitive results that seem to follow from tying one's reasons too closely to one's desires.¹¹

2.3 Rosati's Naturalism

Unlike Korsgaard and Velleman, Connie Rosati (2003) does not offer an answer to the question of what is constitutive of action.¹² But she offers the most explicit account I know of of why it may be a good idea to focus attention on what is constitutive of action in the metanormative context. Rosati starts from Moore's infamous open question argument, arguing that—its many known flaws notwithstanding—it still expresses an insight that should concern ethical naturalists. Rosati's reading of the insight underlying the open question argument is in line with what is now already a tradition of understanding the argument as insisting that normative judgments cannot be reduced to descriptive ones because any such reduction¹³ will lose the very normativity it was supposed to capture.¹⁴ This insight arguably survives the rooting out of Moore's confusions with regard to conceptual analysis and the like and is supposed to pose a threat to synthetic or a posteriori as well as analytic or a priori versions of naturalism (Rosati, 496). Rosati tries to make this line of thought more precise by considering things not from the point of view of the theorizing philosopher but rather from that of an autonomous agent.

is sufficient that some subagential mechanism have it as its aim. For my purposes here, this question need not be decided. For Velleman's most recent view on this matter, see Velleman 2004b, 293, and for his (and Nishi Shah's) most recent view on the parallel issue regarding beliefs, see Shah and Velleman forthcoming.

11. For the claim that the difference between internalism and externalism evaporates if there are universal motivations of the kind Velleman has in mind, see also Moore 1999, 285.

It is not completely clear whether another philosophical motivation underlying Velleman's account is a commitment to naturalism. For his ambivalence on this point, see Velleman 2004b, 294.

12. All references to Rosati's work will be to Rosati 2003 by page number.

13. I am using "reduction" in a fairly liberal sense, one that does not require synonymy or other important a priori relations between the reduced and the reducing.

14. Nicholas Sturgeon (2003, 530, n. 8) is skeptical as to this being a plausible reading of Moore. I do not want to commit myself one way or another regarding this question of Moore exegesis.

Being an autonomous agent consists, Rosati (511) says plausibly enough, in the possibility of critically reflecting on one's desires, stepping back from them and evaluating them rather than just being causally driven by them as a nonagent animal presumably would (actual agents only approximate the ideal of a perfectly autonomous agent [514]). And it is when considering purported naturalist reductions of normative concepts and properties (Rosati is primarily interested in personal good, but the point generalizes) from this angle that they all seem to fail rather miserably, for reductions in terms of my desires cannot help me in critically reflecting on and evaluating (all of) my desires. Hedonistic reductions fare no better because pleasures, pains, and the like still seem—absent any further story, at least—arbitrary from a normative point of view, the only point of view relevant when my autonomous agency is fully engaged in deliberation. And this point—the normative arbitrariness, the arbitrariness from the point of view of the deliberating agent—seems to apply to any naturalist reduction of normativity.¹⁵

But Rosati does not treat this as a refutation of naturalism. Rather, she treats it as a challenge. If they are to avoid this problem, naturalists must come up with a naturalist account “that could survive our reflective scrutiny, that would fit with our capacity for autonomous evaluation and action,” a naturalist account that bears “the proper relation to our agency” (Rosati, 507):

If a naturalistic account of personal good is to fare better, then it will have to find a naturalistic property, relation or causal process that can represent the agent or otherwise secure the proper relation to autonomous evaluation. (510)

Now, though Rosati does not present a detailed naturalist account that would face up to this challenge, she does go a long way toward giving such an account. The main idea she utilizes here—following Velleman—is that there may be some cognitive capacities and motives without which agency is impossible (512). Rosati (513) does not commit herself to what these exactly are and uses the expression “motives and capacities constitutive of agency” as a placeholder for the proper account of the features without which autonomous agency is impossible. These motives and capacities are not unrelated to our agency in the way that,

15. Donald Regan's description of “the moral predicament” is closely related to these observations by Rosati (see Regan 2003, 652). I emphasize the role of deliberation and the problems reductions face from the point of view of the deliberating agent in Enoch 2003 and Enoch forthcoming.

say, simple desires or pleasures are. Indeed, they are tied to the very essence of autonomous agency. These motives and capacities, then, bear the proper relation to our agency.

Assume now that a naturalist account of the motives and capacities constitutive of agency is possible. And assume further that an account of personal good (and of other normative concepts and properties) in terms of the motives and capacities constitutive of agency is likewise possible. Then naturalism can be saved from (Rosati's version of) the open question argument after all, for personal good can be understood in terms of the motives and capacities constitutive of agency, themselves naturalistically understood. The hope of naturalism hinges, so Rosati (524) concludes, on "the prospects of a broadly naturalistic account of autonomous evaluation and agency."¹⁶

According to Rosati, then, the naturalist has her hands full, but at least she is not without hope, and Moore's rejection of *all* naturalist views turns out to be—even today—premature. Indeed, Rosati (520) sketches a naturalist view that does—or may—succeed:

A dispositional theory might, then, incorporate this ideal [of the perfectly autonomous agent], treating personal good as, roughly, the property of being such as to be approved by us when our autonomy making motives and capacities operate effectively and other appropriate conditions obtain. (520)

The theories proposed by Korsgaard, Velleman, and Rosati differ in many ways: Different normative concepts or properties are considered of primary importance (categorical and hypothetical oughts, reasons, and personal good), different features are offered as constitutive of action and agency (self-constitution, self-knowledge, and yet-to-be-specified motives and capacities), and different accounts are given of how exactly these features play their roles in the story of action (Rosati suggests that they do so indirectly, via an ideal advisor model; Velleman believes—or at least believed¹⁷—that what makes something a constitu-

16. Rosati does not acknowledge that there is another necessary condition here, one to which I alluded earlier in the text—personal good (or whatever normative concept we are interested in) must itself be best understood in terms of the motives and capacities constitutive of agency. Perhaps Rosati implicitly assumes that if the naturalist accomplishes so much—giving a naturalistic account of autonomous evaluation and agency—naturalism will be so attractive that such a naturalist-friendly account of personal good will be virtually irresistible.

17. See note 10 above.

tive aim of action is, roughly, that the mechanisms governing action have that aim as their function, naturalistically understood; and I must confess that it is not clear to me what exactly Korsgaard thinks about the way the constitutive aim of action plays its role). But—as is clear even from the sketchy, somewhat oversimplified, and incomplete presentations of these theories above—they are all attempts to ground normativity in what is constitutive of action.¹⁸ And though the rationales suggested for this attempt differ—addressing the skeptic, accommodating externalist intuitions consistently with internalism, and coming up with a naturalist theory that is immune to the improved version of the open question argument—these rationales are obviously related. For all three of these theories are attempts at accounting for such (related) features of normative (or moral) discourse as its objective purport, its apparent universality (or claim thereto), its nonoptionality and nonarbitrariness, while nevertheless avoiding a more robustly realist metanormative theory, one that stipulates the existence of purportedly queer irreducibly normative facts and properties.¹⁹

But can grounding normativity in what is constitutive of action or agency achieve all that?

3. The Problem

These and similar theories can be criticized piecemeal. Even if the constitutivist strategy is broadly right, still Korsgaard has to show that self-constitution (in whatever sense she gives this expression) is indeed constitutive of action and furthermore that all the normativity she wants (morality, the hypothetical imperative, and so on) can be extracted from this aim of self-constitution. Similarly, of course, for Velleman and the kind of self-knowledge that he argues is a constitutive aim of action. And even without the further details that Rosati owes us about the motives and capacities constitutive of action, still if there are general reasons to reject the ideal advisor model,²⁰ they may justify rejecting Rosati's suggestion as well. But, being interested here in the constitutivist strategy more

18. For references to some more constitutivist views, see Lavin 2004, 453, n. 54.

19. In the background of these—and other—attempts to ground normativity in what is constitutive of action there are, of course, Kantian themes and inspirations (perhaps with an Aristotelian twist). I do not know whether an interpretation along these lines is the best reading of Kant, and so I have no view on the applicability to Kant of my criticism of the constitutivist strategy.

20. Such as the one I pursue in my “Why Idealize?” (2005).

generally,²¹ let me put such worries to one side. Assume a constitutive-of-agency metanormative theory whose details are immune to criticism of this piecemeal sort.²² Can such a theory do the job for which it was hired?

The problems this strategy was supposed to solve—the skeptical challenge, the externalist challenge, and the antinaturalist challenge—all revolve around the idea that the normative, though somehow essentially tied to us and our desires (or at least our will), can nevertheless not be as arbitrary as our (usual) motivations. But then how is the fact that some motivations are constitutive of action supposed to help?

Consider Rosati first. According to Rosati, we are to think of a (perhaps ideally) autonomous agent stepping back from her desires (of whatever order) because she sees them—when deliberating and evaluating—as normatively arbitrary. And we are to think of her as troubled by the fact that it is hard to see what could possibly give the answers she is looking for because all facts of her psychology are just as arbitrary as her desires. She then finds out that some parts of her psychological makeup are unique in that they are such that without them she would not have qualified as an agent at all. Knowing that, is she supposed to be relieved? Why does it matter, as far as the question of normative arbitrariness is concerned, that some parts of her psychology have this necessary-for-agency status? Why shouldn't our agent treat the motives and capacities constitutive of agency as normatively arbitrary? Why shouldn't she treat the very fact that they are constitutive of agency as normatively arbitrary? She is, remember, stepping back from her desires, attempting a kind of detached scrutiny, evaluation, and choice. But then how is the constitutive-of-agency status at all relevant? What is it to her, so to speak, if some of her motives and capacities enjoy such a status?

Or consider Korsgaard's hope of grounding a reply to the skeptic

21. Velleman too seems to be more interested in the general features of the constitutivist strategy than in the specific aim he suggests as constitutive of action. See Velleman 2004a, 234.

22. There may also be other general problems—different from the one I proceed to pursue—that all constitutivist theorists have to face, whatever the details of their theory. Perhaps the clearest example of such a general problem is that of accounting for wrong, or bad, or irrational actions, actions that violate the relevant norms but that nevertheless qualify as actions. Korsgaard (5.4.1) mentions this problem, but it is not clear to me how she ultimately addresses it. For intricate discussions of improved versions of this problem, see Setiya 2003, 374–75 and Shah 2003, 460–65.

The problem about to be discussed in the text is different from this other general problem and, as far as I can see, independent of it.

in what is constitutive of action.²³ We are to imagine, then, someone who remains indifferent when we tell him that his actions are immoral or irrational. He then reads Korsgaard and is convinced that self-constitution is a constitutive aim of action, so that you cannot even count as an agent and your bodily movements cannot even count as actions unless you aim at self-constitution of the kind Korsgaard has in mind. And assume that our skeptic is even convinced that—miraculously²⁴—morality and indeed the whole of practical rationality can be extracted from the aim of self-constitution. Do we have any reason to believe that now he will care about the immorality or irrationality of his actions? Why isn't he entitled to respond along the following lines: "Classify my bodily movements and indeed me as you like. Perhaps I cannot be classified as an agent without aiming to constitute myself. But why should I be an agent? Perhaps I can't act without aiming at self-constitution, but why should I act? If your reasoning works, this just shows that I don't care about agency and action. I am perfectly happy being a shmagent—a nonagent who is very similar to agents but who lacks the aim (constitutive of agency but not of shmagency) of self-constitution. I am perfectly happy performing shmactions—nonaction events that are very similar to actions but that lack the aim (constitutive of actions but not of shmactions) of self-constitution." Has Korsgaard put us in a better spot vis-à-vis this why-be-an-agent (rather than a shmagent) problem than we were vis-à-vis the why-be-moral or why-be-rational challenges with which we—or at least Korsgaard—started? Consider again the example of the house and the shoddy builder, and suppose we manage to convince him that certain standards—standards he previously did not care about and regularly failed to measure up to—are constitutive of being a house. It seems he is entitled to respond: "Very well then, I guess I am not engaging in the project of building a house but rather in the project of building a shmouse, of which these standards aren't constitutive. So what is it to me how you classify my project?"

At times Korsgaard writes as if she thinks no such retort—either in the house case or in the metaethical or metanormative case—is possible. In Lewis's (1996, 60) terms, at times Korsgaard writes as if she believes that the threat that your inner (and outer) states will fail to

23. Velleman's quasi-externalist motivation is a harder case. I return to it below.

24. For a theory of intentional action that is in some respects very close to Velleman's but that argues that action's constitutive aim is just too thin to generate the norms of practical reason, see Setiya 2003.

deserve folk-theoretical names (such as “action”) is indeed a threat that will strike terror into the hearts of the wicked.²⁵ But no support is offered for this surprising claim. And notice that Korsgaard’s problem here is not merely that the skeptic is unlikely to be convinced by such a maneuver. The problem runs deeper than that because the skeptic *should* not be convinced.²⁶ However strong or weak the reasons that apply to him and require that he be moral, surely they do not become stronger when he realizes that unless he complies with morality his bodily movements will not be adequately described as actions.

Notice that the problem is not that action does not have a constitutive aim, or that there are no motives and capacities constitutive of agency. Indeed, I am here granting these claims for the sake of argument. Nor is the problem that such constitutive aims, motives, and capacities are philosophically uninteresting. For all I am about to say, they may be able to explain much that is philosophically important as well as interesting.²⁷ The problem is just that it is hard to see how the constitutivist strategy can serve to ground normativity or to solve the metanormative problems it was supposed to solve.

4. Attempts at Addressing the Problem

Quite surprisingly, Korsgaard nowhere—to the best of my knowledge—addresses this worry explicitly. And though Velleman—in a very recent text—does mention it, it is not clear to me what exactly he says in reply.²⁸ Rosati (521) anticipates the worry:

The complaint about both brute and hedonistic naturalism was that they treat as normative certain natural forces or tendencies that lack sufficient normative credentials. But why think that the motives and capacities that render persons agents do not have the same problem?

25. The “tyrannical person does not really choose *actions*, in the technical sense I have defined” (Korsgaard, 5.5.2).

26. Emer O’Hagan (2004, 56) notes that the skeptic is not likely to be convinced by Korsgaard’s reasoning. He does not note that the skeptic *should* not be convinced.

27. For Velleman’s list of the phenomena that his theory about the constitutive aim of action helps to explain, see Velleman 2004b, 288–89.

28. See Velleman 2004b, 290–92. To the extent that I understand these remarks, Velleman here suggests the responses that I discuss below, in section 4.3 and in section 6.

Perhaps Velleman doesn’t address this worry more carefully because—in some moods, at least—he doesn’t consider it a serious worry. In the case of the constitutive aim of belief—in which Velleman hopes to ground all reasons for belief—Velleman (2000d [1996], 186) writes: “The question whether to be inclined toward the truth

Of course, not any desire or aim is constitutive of agency, and so those that are are unique in some way, but the question—which Rosati puts in terms of the antinaturalist challenge she is interested in—is whether this status is *normatively* relevant, whether it is relevant in the way needed in order to make the motives and capacities constitutive of agency nonarbitrary in the intended sense.

Here is how Rosati (522) attempts to address this worry:

Unlike our other motives and capacities, our autonomy-making motives and capacities are not arbitrary but, rather, make self-governance possible: they are motives and capacities without the effective operation of which we would not be agents and evaluators at all. Insofar as we are agents, the effective exercise of these motives and capacities matters to us, and our caring about them involves no identifiable mistake. Their operation, we might say, is self-vindicating, and efforts to challenge them cannot even get going without relying on them.

It seems to me that there are in this passage three distinct (though perhaps related) lines of thought:²⁹ According to the first, the status of being constitutive of agency renders the relevant motives and capacities nonarbitrary, thus normatively vindicating them. According to the second, what normatively vindicates these motives and capacities is that they matter to us, and “our caring about them involves no identifiable mistake.” According to the third, there is something self-defeating about challenges to the capacities and motives constitutive of agency, and so they have a self-vindicating status. Let me discuss these three suggestions in turn.

4.1 Nonarbitrariness

The first can be dismissed fairly quickly. Of course, if being constitutive of agency renders the relevant motives and capacities nonarbitrary (in the intended, normatively relevant, sense), then the problem is solved. But this, remember, is exactly where the problem was located. The question was *why* it is that being constitutive of agency renders these motives and capacities nonarbitrary. Just saying that they are does not, of course, constitute a satisfactory answer. And notice that Rosati—like anyone else

on some topic—and hence whether to be subject to reasons for belief about it—is left entirely open.”

29. Rosati’s paragraph continues, but only—it seems to me—to express the same points in a different way.

who acknowledges the force of the antinaturalist challenge—cannot just dig in her heels, relying on the intuitive plausibility of the normative nonarbitrariness of the motives and capacities constitutive of agency. For intuitively speaking, pleasures and pains—reduction to which Rosati (rightly, I think) rejects—seem much less arbitrary and much stronger as candidates for normative significance than any motives and capacities plausibly considered constitutive of agency.

4.2 *We Do Care and without Mistake*

The second line of thought found in the passage quoted seems to be that, first, we care about the effective exercise of these motives and capacities, and second, this caring does not seem to involve any mistake on our part. The conjunction of these two claims is presumably thought of as defusing the worry that these motives and capacities are just as arbitrary from the point of view of the deliberating agent as all others are. But it is hard to see how this can be made to work: First, it is not at all clear that we do care about whatever is constitutive of agency. As Allan Gibbard (1999) notes, we may value many things without valuing our capacity to value them and to act accordingly. Second, and more importantly, the mere lack of an identifiable mistake is surely not sufficient to vindicate *other*, non-constitutive-of-agency motives and capacities from the charge of arbitrariness, or else the critique of, say, hedonistic naturalism collapses (for what identifiable mistake is involved in our caring about pleasures and pains?). If such lack of an identifiable mistake is to vindicate the motives and capacities constitutive of agency, then this must be because there is some *other* thing normatively special about them compared to other motives and capacities that do not involve an identifiable mistake. But, again, it is exactly this normative uniqueness of the motives and capacities constitutive of agency about which we are now raising doubts. Rosati is not entitled to assume these doubts away.

4.3 *Self-Vindication*

It seems, then, that if an answer to our worry is to be found in the quoted passage, it must be in the idea that the motives and capacities constitutive of agency are self-vindicating (or, as Rosati says later in the paragraph, self-supporting). Let me again quote the relevant sentence from Rosati (522):

Their operation, we might say, is self-vindicating, and efforts to challenge them cannot even get going without relying on them.

The idea—also hinted at by Velleman (2004b, 290–91)—seems to be this. We are in the process of trying to evaluate the motives and capacities—whatever exactly they are—that are constitutive of agency. But the very challenge presupposes, it seems, that it can be met, because the activity of evaluating these motives and capacities—presumably an instance of one of the activities agency consists in—already presupposes them and their significance. Take away the motives and capacities constitutive of agency and you have taken away agency itself and with it the challenge to the normative significance of these motives and capacities. The challenge to them is thus rendered unstable, and so their status is self-vindicating.

The idea is familiar from the discussion of skeptical challenges in other contexts and is perhaps most clearly pressed and developed by Thomas Nagel in *The Last Word* (1997).³⁰ Nagel rightly notes, for instance, that in order to launch a skeptical attack on logic, the skeptic is going to have to use some logic. Nagel then argues that the skeptic's dialectical position is thus rendered unstable. Logic and—perhaps more generally—Reason herself are simply unavoidable; even the skeptic, while launching her skeptical attack, still reasons. So skepticism about logic and Reason is unstable. Our use of logic is immune to global skeptical challenge and can thus be thought of as self-vindicating.³¹

But this line of thought—however influential—nevertheless fails.³² Skeptical challenges—some, at least, including the ones relevant here—are best seen, I think, as highlighting tensions within our own commitments, as paradoxes arguing for an unacceptable conclusion from premises we endorse, employing rules of inference to which we are committed. In responding to such challenges, we must not yield to the temptation of the “adversarial stance” (Wright 1991, 89): The philosophical challenge is not to defeat a real person who advocates the skeptical view or

30. See also Thomas Nagel's emphasis on the inescapability of the demands of both prudence and ethics in *The Possibility of Altruism* (1970, for example, 18–19).

31. The case of logic is not merely an analogy here: Rosati (522) mentions the capacity to reason logically as a cognitive capacity possibly constitutive of agency.

32. For discussions of what follows (though in a different context), I am indebted to Josh Schechter and Paul Boghossian.

Let me also note here that, first, despite the criticism that is to follow, there is much in the line of thought that Nagel (1997) develops in *The Last Word* that I find both plausible and important (more on this in section 8 below), and second, in conversation it has become clear to me that Nagel's view (today) is more complex than—and not as vulnerable to the criticism in the text as—the view *The Last Word* is naturally understood as expressing.

occupies the skeptical position (*what* view or position?) but, rather, to solve the paradox, to show how we can avoid the unacceptable conclusion at an acceptable price. If we must think of the situation in dialectical terms, we should think of skeptical challenges as *ad hominem* arguments, with all of us as the relevant homini. “The skeptic” is entitled to use, say, logic because *we* are committed to the legitimacy of doing so. And he is entitled to engage his motives and capacities that are constitutive of agency even while putting forward a critique of them because *we* are (purportedly) committed to the legitimacy of him so doing. In other words, the skeptic is entitled to use our own weapons against us. If, using these weapons, he can support a conclusion we are not willing to swallow—one stating, for instance, that the very weapons he is using are not ones we are entitled to use—then it is *we* who are in trouble, not him, because we have been shown to have inconsistent commitments. Think of the situation as an analogue of a *reductio ad absurdum*: a good *reductio* argument assumes a certain claim only to prove it false. And just like a sound (*reductio*) argument establishing that a certain claim is false may assume that it is true, a sound argument concluding that a certain method of reasoning is incorrect or unjustified may employ that very method.³³ And the analogous point applies, it seems to me, to the reliance on the motives and capacities constitutive of agency. If the motives and capacities (or indeed aims) constitutive of agency can be turned against themselves in this way, surely this is a problem for their friends, for those, like Rosati, who advocate the normative significance of these motives and capacities, not for their enemies. That the skeptic—if there actually is such an interlocutor—cannot avoid engaging the motives and capacities constitutive of his agency while mounting his attack only shows that he is in the same boat as the rest of us (surely, not a surprise). It does not show that these motives and capacities are self-vindicating or nonarbitrary from a normative point of view.³⁴

33. So if, for instance, we can construct a slippery slope argument the conclusion of which is that we should not use slippery slope arguments, it is the friends, not foes, of slippery slope arguments who are in trouble. See Enoch 2001.

Ruth Weintraub (1997) agrees that skeptical challenges are best understood as *ad hominem* arguments but is nevertheless more sympathetic to the claim that the skeptic self-defeats. But she too agrees, I think, that even if in the dialectical situation described in the text the skeptic doesn’t win, still we lose.

34. Perhaps Nietzsche (1967, sec. 515) had something like the argument in the text in mind when writing: “Not being able to contradict is proof of an incapacity, not of ‘truth.’”

Even if it is true, then, that any effort to challenge the motives and capacities constitutive of agency “cannot even get going without relying on them,” this is still not enough to save them from normative arbitrariness: when an agent deliberates in the kind of reflective and detached frame of mind Rosati describes such that even, say, facts about her own desires or her pleasures and pains cannot but seem to her normatively arbitrary, the mere fact that some of her motives and capacities are such as to be necessarily shared by anyone challenging their normative significance does not suffice to present them as nonarbitrary in the intended sense.

5. Why Normativity Cannot Be Grounded in What Is Constitutive of Action

Rosati’s suggestive remarks do not, I conclude, show promise as ways of addressing the problem—they do not show that being constitutive of agency renders the relevant desires and capacities (or aims) nonarbitrary in the intended sense. But can some other line of thought, one not mentioned by Rosati, be more successful? I think this is highly unlikely.

To see why, think of an analogy Velleman (2000d [1996], e.g., 178) often uses when discussing the aim and motives he takes to be constitutive of agency and action. Games, it seems plausible to say, often have constitutive aims. Perhaps, for instance, if you are not aiming to checkmate your opponent, you are not really playing chess. The aim of scoring a goal (or perhaps outscoring the other team or something of this sort) is plausibly considered constitutive of the game of soccer, and so on. Velleman emphasizes that normativity naturally emerges from these constitutive aims—so long as you are playing chess, you have a reason, it seems, to attempt to checkmate your opponent, to want to checkmate her, to sacrifice a pawn when doing so increases the likelihood of your eventually checkmating her, and so on (Velleman 2000d [1996], 187–88). Velleman’s hope—and Korsgaard’s and Rosati’s too—is to show how the most general norms of action likewise emerge from constitutive features, this time from the constitutive features of action and agency.

But the game analogy, I think, is more complex. In order to have a reason to checkmate your opponent, it seems to me it is not sufficient that you do in fact play chess. Rather, it is also necessary that you *have a reason to play chess*³⁵ (or to play the relevant specific game of chess).

35. For a similar point made in a different context, see Marmor 2001, 30. Shah

Suppose you somehow find yourself playing chess (or, if we are going to be tendentiously picky about what qualifies as actually playing chess, *seemingly* playing chess), but you do not care about the game and about who wins, nor do you have any reason so to care. It seems rather clear to me that you have no reason whatsoever to attempt to checkmate your opponent. And if a metanormative (or metachess) theorist then comes along, explaining to you that attempting to checkmate your opponent is constitutive of the game of chess, so that unless you engage in such attempts your activity will not be classifiable as chess playing, it seems to me you are perfectly justified in treating this information as normatively irrelevant. After all, what is it to you how your activity is best classified? If you have no reason to be playing chess, then that some aim is constitutive of playing chess gives you no reason at all, it seems to me, to pursue it, and this whether or not you are in fact playing chess. (Analogously: If you have no reason to be building a house, then that some standard is constitutive of being a house gives you no reason at all to measure up to it, and this whether or not you are in fact building a house.)

If a constitutive-aim or constitutive-motives theory is going to work for agency, then, it is not sufficient to show that some aims or motives or capacities are constitutive of agency. Rather, it is also necessary to show that the “game” of agency is one we have reason to play,³⁶ that we have reasons to be agents rather than shmagents (analogously: that we have a reason to build a house rather than a shmouse). And this, of course, is a paradigmatically normative judgment. This means, first, that there is no hope of rendering a naturalistic reduction of the normative respectable by focusing on what is constitutive of action because the constitutivist line cannot succeed unless we have a prior reason to be agents, and this prior reason cannot be naturalistically reduced by the constitutivist theory itself. And second, even discarding naturalistic aspirations (that are of central importance to Rosati, but not, it seems, to Korsgaard and perhaps not to Velleman either), if we need a normative judgment—that we have a reason to be agents rather than shmagents—in order for the constitutive-of-agency strategy to kick in, then the constitutivist strategy cannot give us the whole story of normativity. It is no longer possible that

(2003, 458) quotes from an unpublished manuscript by Gideon Rosen entitled “Meaning, Normativity, and All That.” In the passages there quoted, Rosen seems to be making a similar point (though he seems primarily interested in motivation not normativity).

36. Velleman (2004b, 290–91) concedes that the question whether we have such a reason remains “in a sense” open.

all practical reasons are grounded solely in what is constitutive of agency. Normativity, in other words, cannot be grounded in what is constitutive of agency.³⁷

Things would have been different, of course, had it been possible to invoke here a normative truth, robustly realistically understood, to the effect that we all have a reason to be agents. With this normative claim in hand, the rest of the constitutive-of-agency line can be pursued rather safely. And nothing I say here goes to deny that even with such a normative claim in hand questions about the constitutive aspects of action may be interesting and indeed normatively relevant. But introducing such a normative truth—itsself not accounted for by the constitutivist strategy—defeats the motivations underlying this strategy: it renders Korsgaard unable to cope with the skeptic who is left indifferent to the purported reason she has to be an agent; it defeats Velleman's quasi-externalist project because this normative reason, the reason to be an agent, must either be understood itself as an external reason or else must be held hostage to the motivations of specific agents, so that it does not apply to those who lack the motivation to be agents; and it defeats Rosati's naturalism because the normative statement—that we have a reason to be agents—is not itself naturalistically reduced by the constitutivist line.³⁸

6. Why the Unavoidability of Agency Won't Help

Perhaps, though, I have been misusing the game analogy. For it may be thought that whatever intuitive plausibility there is to the claim that your reasons to attempt to checkmate your opponent depend not just on your playing chess but also on your having a reason to play chess comes from

37. For a related discussion, see Railton 1997, 75–79. Rosati (512, n. 59) refers to this essay, but she does not, I think, address Railton's worries satisfactorily.

38. Invoking an independent reason to be agents thus defeats the metanormative rationales of the constitutivist strategy. But this may not be true with regard to the *normative* rationale mentioned in the introduction. Indeed, if we have an independent reason to be agents, then it seems like asking questions about what agency consists in is exactly the way in which to proceed in order to gain insight into our reasons.

This may be so, but: First, even if this is so, the point in the text—about the *meta*-normative philosophical motivations for constitutivism—stands. And second, note how implausible the claim is that (all?) our normative reasons are based on the fundamental reason we have to be agents. If I had to choose candidates for ultimate normative or moral truths, I would consider such things as the badness of pain, the significance of autonomy, the moral status of love, the duty or duties of respect for persons, the obligation to stand by one's word. I certainly wouldn't mention anything about a reason to be an agent.

the fact that you can avoid playing chess. You can choose not to play chess in the first place, and even if you find yourself playing chess, you can opt out of the game. This is why it makes sense for you to ask for a reason to play chess and not to acknowledge the chess-related reasons you have (to sacrifice a pawn and so on) unless you have a reason to play chess.

The game of agency, however, is crucially different. Quite plausibly, it is a game you find yourself playing, and one you cannot opt out of. You can, of course, choose to end your action-performing life by simply ending your life. But far from opting out of the game of agency altogether, this very decision will be a major move in that game (Velleman 2004b, 291). Perhaps, then, for games you cannot avoid playing, there is no need for a reason to play the game in order to have the reasons internal to the game. And indeed, it is a central theme throughout Korsgaard's work that agency, action, reflection, and deliberation are not things we can discard or avoid: "Human beings are *condemned* to choice and action" (Korsgaard, 1.1.1). Korsgaard never says so explicitly (as far as I know), but it seems to me the emphasis on the unavoidability of agency is meant to preempt something like the objection in the previous section.³⁹

But it is hard to see how such unavoidability can help.⁴⁰ "Perhaps," Korsgaard's skeptic may say, "I cannot opt out of the game of agency, but I can certainly play it half-heartedly, indeed under protest, without accepting the aims purportedly constitutive of it as mine." The kind of necessity the game of agency has to enjoy in order to solve the problem we are now in is *normative* necessity. Invoking other necessities here will just not do.⁴¹ Perhaps noticing this problem, Korsgaard is never clear about the kind of necessity she invokes.⁴²

39. In correspondence and in conversation (several years ago) Velleman also suggested that a solution to my problem is to be found in the unavoidability of action. For an emphasis on unavoidability in the context of a constitutivist account of epistemic norms, see Rysiew 2002, 450.

40. Note that the unavoidability point is different from the self-vindicating point discussed in section 4.3. There the claim was that the fact that any challenge to the normative significance of the motives and capacities constitutive of agency seems to engage them shows that they are self-vindicating or otherwise immune from criticism. Here the claim is that the fact that *we* cannot discard whatever is constitutive of agency plays a role in saving the constitutive-of-action strategy from trouble. These two claims may be related, but they are certainly distinct.

41. Sharon (n.d.) makes a similar point.

42. Korsgaard (1.1.1) says that action is necessary but neither logically nor causally nor rationally. She doesn't say what the notion of necessity involved is, except to

Think again about finding yourself playing a game of chess, and assume for now that for some reason you cannot quit—not that you *should* not quit but that you *cannot* quit. And assume that sacrificing a pawn is the thing you have most chess-related reason to do (it best promotes your chances of checkmating your opponent or some such). Well, do you have a reason to sacrifice a pawn? Not, it seems to me, if you don't have a normative reason to play or win the game, and this even if you can't quit.⁴³ For you can continue playing or “going through the motions,”⁴⁴ grudgingly, refusing to internalize the aims of the game. And absent some normative reason to play the game, there need be nothing irrational about such an attitude. Nonoptionality of this constitutive (rather than nor-

say that “It is our plight.” And later on Korsgaard (1.3.4) writes: “The principles of practical rationality bind us because we must constitute ourselves as unified agents . . . The principles of practical reason are normative for us, then, simply because we must act.” Similarly, Korsgaard (n.d., 7) talks of the relevant necessity as a necessity “that we confront.” Nowhere, as far as I know, does Korsgaard clarify these ways of talking or comment on the relations between all these different ways of talking about possibly different notions of necessity.

Alan Gewirth (1978, 158) too seems to conflate different modalities when defending a somewhat similar move: “the dialectically *necessary* aspect of my method gives the PGC [Principle of Generic Consistency]’s ‘ought’ a more than contingent or hypothetical status.”

43. Some people think that the possibility of failing to stand up to a normative standard is a necessary condition for the standard's status as normative. If this is so, we can make a stronger claim: If you can't quit playing the game, then not only doesn't it follow that you have a reason to play (or to make a certain move), but also it *does* follow that you do *not* have such reasons (or at least not the reason to play the game). But I am not sure it is true that the possibility of noncompliance is indeed necessary for normativity. For intricate discussion of these matters, including a discussion of the (in)consistency of constitutivism with this purported necessary condition for normativity, see Lavin 2004.

44. There is a complication here. If you don't count as playing chess unless you're trying to win, then the description in the text is incoherent: For “going through the motions” without trying to win does not count as playing chess, and then if you can do this, it is not after all true that you cannot opt out of the chess game—you just did! This shows that there is something tendentious about refusing to describe someone going through the motions as genuinely playing chess—certainly, in everyday parlance we would be happy saying that he's playing chess even though he doesn't care about winning. (Here is Korsgaard [4.6.2] employing a similarly tendentious reading of “to deliberate”: “an agent who is deliberating about what to do cannot completely ignore the Kantian imperatives and still recognizably be deliberating about what to do, because what it *means* to deliberate is to be guided by those imperatives.”) If we nevertheless insist that he is not genuinely playing chess because you can't be playing chess without trying or wanting to win, then we must conclude that if it is (nonnormatively) necessary that I play chess, it is also (nonnormatively) necessary that I want to win. Returning

mative) kind is thus simply neither here nor there.⁴⁵ The antinaturalist challenge raises its head again: nothing short of an explicitly normative claim seems fit to settle normative questions.⁴⁶ (If you're unconvinced, think of all the features plausibly considered nonoptional in this non-normative way: Do you think normativity emerges from all of them?)⁴⁷

The point can be restated in terms closer to the general problem Rosati reads into the open question argument. As an autonomous agent, you sometimes reflectively scrutinize the desires or the dispositions to feel pleasure or pain you find yourself with, stepping back from them to achieve normative evaluation. But this is true also of the games you find yourself playing. Just as you, as an autonomous agent, cannot settle, at least when your agency is fully engaged in deliberation, for the desires you happen to find yourself with, so too you cannot settle for the games you just find yourself playing. What you need to see is that the relevant desire is worth having, or that its object is worth pursuing, or that the relevant game is worth playing. But none of these are, as they stand, nat-

now to the game of agency, this would mean that the constitutivist theorists insist that we (nonnormatively) necessarily care about constituting ourselves, knowing ourselves, or whatever, and so presumably that we all already care about being rational, moral, and so on. Perhaps—I am not sure—this is a result Velleman will be happy with. But it is hard to think of this as a result that Korsgaard and Rosati—concerned as they are with the moral and rational skeptic and the failures of naturalism—can accept as the ground of normativity, because it shows that the skeptic is not wrong but rather impossible, and it after all relies on a naturalist reduction in terms that seem normatively arbitrary. And independently, this claim—that a skeptic of this sort is impossible—seems like a very implausible result indeed. For a discussion of a closely related point, see Hussain 2004, 273.

45. Indeed, nonoptionality of this nonnormative kind seems relevantly similar to psychological nonoptionality. So compare Hookway's (1999, 178) discussion of the psychoanalyst who tells you that some of your beliefs are such that you cannot discard them. Have you now gained justification for holding them?

46. For claims that the constitutivist strategy—at least as employed in the past by Velleman—loses normativity, see Shah 2003, 461; Mele 2004, 259; and Hussain 2004, 267. See Velleman 2004b, 291–97 for his attempt at a reply. And for the claim—in a different context—that the unavoidability of participating in a practice does not exempt one from the need to support such participation by reasons, see Marmor 2001, 37–39.

47. What if, as an anonymous referee suggested, it's metaphysically impossible to just grudgingly go through the motions of the game of agency without thereby also *doing* something else, like performing some other actions whose aims I do internalize? If so, then there may be a sense in which the inescapability of action (or better, it seems to me, of deliberation) can ground a commitment to normative truths and facts. But this way will be very different from the one pursued by constitutivists. Indeed, it will be the line hinted at in section 8 below.

uralistically respectable. That some of the desires you find yourself having or the games you just find yourself playing are such that if you cease to have or play them, you will no longer merit being called “an agent” is beside the point, as is the fact that some of these desires are ones you can’t discard and some of these games are ones you can’t opt out of.⁴⁸ Absent some further, normative argument, this is just another normatively arbitrary fact about you.⁴⁹

Perhaps—I am not sure—constitutivists can strategically retreat, conceding much of what I’ve said without giving up on the importance of the constitutivist strategy altogether. Perhaps, for instance, Korsgaard can argue that given that we cannot opt out of the game of agency, the

48. Earlier on I presented the skeptic as not caring about how his bodily movements are classified, as not caring whether he merits being called an agent. Perhaps someone could have objected that I interpret a metaphysical point (about what is constitutive of agency) as a metalinguistic one (about the meaning of ‘agency’). The discussion in the text shows, of course, that metaphysical necessity here doesn’t help.

49. Velleman (2004b, 293) is now aware of the problem. In response to it, he now endorses—following Nishi Shah (for instance, Shah 2003, 475 and 481, n. 41)—a revised view, in which *action* (and *belief*) are normative concepts, the possession conditions of which require endorsing the relevant norms, with “endorsing norms” understood along norm-expressivist lines. For the most recent statement of this view (with regard to belief), see Shah and Velleman forthcoming. But this will not save the constitutivist strategy from the criticism in the text, for the following reasons. First, on this account, the ultimate metanormative story is the norm-expressivist one, not the constitutivist one. So even if this account works, it does not show that normativity can be grounded in what is constitutive of action. Second, such a revision is subject, of course, to whatever objections norm-expressivism in general is subject to. And third, and most importantly, this revision doesn’t seem to solve the problem. For even if I cannot master the concepts *action* or *belief* without endorsing the relevant norms, and even if (therefore) I cannot deliberate about what to do or believe without endorsing these norms, this does not give me a normative reason to endorse them, it doesn’t give me a normative reason to deliberate about actions and beliefs rather than about shmac-tions and shmeliefs or a normative reason to deliberate rather than to shmeliberate about actions and beliefs, or something of this sort. This revision, in other words, is just another attempt at the unavoidability response to the arbitrariness problem, and it fails for similar reasons. This last point applies also, I think, to the points made in Shah 2003, 465–74.

Shah and Velleman’s interesting account also raises other, philosophy-of-language-related questions because it employs the idea of possession conditions of concepts as the basis of the relevant norms (Shah and Velleman forthcoming). For some relevant doubts here, see Schechter and Enoch forthcoming. Indeed, it seems to me that a precisely analogous problem to the one highlighted in this essay can be raised with regard to an account in terms of possession condition of concepts. For it appears that Shah and Velleman—as they will probably agree—cannot give reason to employ the concept *belief* (rather than some competing concept). But I digress.

skeptical challenge—though not defeated—is still shown somehow not to be relevant, not to be a real option for us.⁵⁰ There is a sense in which this claim seems to me clearly true—for it does seem that we cannot genuinely refrain from committing ourselves to normative judgments. But it is not clear what follows from this. In section 8, I will argue that a proper understanding of this point does have important metanormative implications, but ones that are not at all constitutivist.

I do not want to overstate the point. Perhaps a strategic retreat (to a “skeptical solution,” maybe) can save something from the insights underlying some constitutivist suggestions. But then, first, we are owed the details of such a retreat (including an argument showing it is not a full surrender), and second, whatever the details, such a retreat will indeed be a retreat, a concession that the constitutivist strategy cannot supply all the goods it is thought by its supporters to supply.⁵¹

7. Conclusion

We are now in a position to conclude, I think, that normativity cannot be grounded in what is constitutive of agency. Perhaps a case can be made that some desires, capacities, or aims are constitutive of agency. And perhaps this even makes them reason supported—if, that is, there is an independent reason to be agents. Perhaps, in other words, thinking about what is constitutive of agency can do some work in our metanormative theory. But it cannot do the fundamental work it was supposed to do.

Returning now to the philosophical motivations of the three constitutivist theorists I have been using as my examples: It is clear that the constitutive-of-agency strategy has absolutely no effect on Korsgaard’s skeptic, nor should it, for the skeptic can, should, and probably will care just as little about that which is constitutive of agency as he does about what is morally or rationally required. Philosophy, then, cannot replace

50. I thank an anonymous referee for pressing me on this point.

51. Many of the argumentative moves in this essay have close parallels in discussions of transcendental arguments. And this is no surprise, as both the constitutivist strategy I criticize and the alternative I am about to suggest (in section 8) may be thought of as variants of transcendental arguments. Of course, I cannot do justice here to the vast literature discussing transcendental arguments. I hope to pursue the relations between some of the points made in this essay and this literature elsewhere. Let me just emphasize that even if general objections to transcendental arguments fail, and some transcendental arguments do support their conclusions, still the objection to constitutivism developed in this essay stands.

the hangman. It is also clear that if Rosati is right—as I think she is—that most suggested naturalist reductions lack sufficient normative credentials, then naturalist reductions in terms of the capacities and desires constitutive of agency fare no better.

How about Velleman's quasi-externalist motivation for focusing on the motives that are constitutive of action and agency? Its success depends, I think, on how ambitious it is. True, if a plausible case can be made that all agents share some motives, then universality—in some sense of this word—can be had even on internalist assumptions. And if a plausible case can be made that all *possible* agents share these motives, then the internalist can go even further in accommodating externalist intuitions. But this is not all the externalist wants. Typically, externalists insist that one's normative reasons do not depend on one's desires, that such dependence is already objectionable unless the desires have something normative going for them. And if—as I argued above—being constitutive of agency does not qualify as a normative story of the kind needed to vindicate such dependence, then Velleman cannot after all give the externalist all that she wants.⁵²

In a sense, I have been arguing for a limited result. All I have shown is that grounding normativity in what is constitutive of action cannot help the constitutivist theorist cope with the skeptical, the externalist, and the antinaturalist challenges. And one may argue, I guess, that none of these “challenges” is a genuine challenge or that there are other promising ways of addressing them.⁵³ If so, perhaps something of the constitutivist strategy can be saved. But then it is hard to see how it can be philosophically motivated. If, for instance, a desire need not even be constitutive of agency in order to be normatively nonarbitrary, then normativity can be reduced to usual desires, and there is neither a need nor a reason to restrict such naturalist reduction to just those desires that are constitutive of agency. Either way, then, the solution to our metanormative problems will not emerge from an understanding of what is constitutive of action.

52. See Moore 1999, 285 and Shafer-Landau 2003, 170, n. 7 for related points.

53. Myself, I take the externalist and antinaturalist challenges *very* seriously—indeed, I am an externalist and a nonnaturalist. But I think that the skeptical challenge as Korsgaard seems to understand it—roughly, the challenge that comes down to “What reason do I have to do what I have reason to do?”—is confused and not a genuine challenge at all (for a similar point see, for instance, Shafer-Landau 2003, 179–80). As I proceed to say in the text, this doesn't save the constitutivist.

Philosophy of action, of course, does not need to have metanormative payoffs in order to be worth engaging in.⁵⁴ And for all I have said, action and agency may have constitutive aims, desires, or capacities. It's just that normativity will not come from what is constitutive of agency.

8. An Alternative

For these reasons, then, I don't think that normativity can be grounded in what is constitutive of agency. Nevertheless, I think that the attempts along these lines incorporate some important insights, ones that can show us a better metanormative way in which to proceed.

Consider first Rosati's insistence on the point of view of the reflective, deliberating agent as the one most relevant here. No one would find even an initial appeal, for instance, in an attempt to ground a conception of the good in, say, what is constitutive of being a state of affairs (such that for a state of affairs to be good is just for it to be good *as* a state of affairs). Rosati's attempt to ground normativity in what is constitutive of agency is—even if it fails, as I argued above—not *as* misguided as such an attempt. And this means that there is something importantly right about Rosati's insistence on the point of view of the deliberating agent.

Furthermore, I think Korsgaard is right in emphasizing the unavoidability (for us) of deliberation or reflection, and that she is also right in taking reasons or other normative facts to be needed for deliberation. But I think she is mistaken in thinking that any of this counts in favor of grounding normativity in what is constitutive of agency or of action.

The metanormative significance of the point of view of the deliberating agent and of the unavoidability of deliberation is, I suggest, not metaphysical but rather epistemological.⁵⁵ Normative truths do not depend on our beliefs or desires, be they constitutive of agency or other-

54. As Alfred Mele (2004, 259) notes in a similar context.

55. Thus, I agree with at least one way of understanding Korsgaard when she writes: "In order to answer that question [whether our incentives give us reasons] we need principles, which determine what is to count as reasons" (4.2.4). But then she (*ibid.*) proceeds to conclude that this tells us something about the metaphysics of normativity: "And so it is in the space of reflective distance, in the internal world created by self-consciousness, that reason is born." As I argue in what follows in the text, I believe the need for normative facts is relevant epistemologically, not metaphysically. Certainly, the "so" in the last quote from Korsgaard is at the very least premature: Nothing has been said in support of the claim that the metaphysical claim about reasons follows from our need for reasons.

wise. But *our reason for believing* in normative truths does indeed depend on the deliberative project and on it being a nonoptional project for us.

The intuitive idea is fairly simple: Just as arguments from explanatory indispensability—instances of inference to the best explanation—can justify a commitment to, say, electrons, so too arguments from deliberative indispensability can justify a commitment to, say, normative reasons, or normative truths or facts. The fact that our best attempts at making sense of the world—our best scientific theories, for instance—require the existence of electrons is commonly (and rightly, I think) taken as reason to believe in the existence of electrons. Similarly, the fact that our best attempts at deciding what to do and how to live our lives require normative facts (indeed, irreducibly normative ones) gives us just as good a reason to believe in normative facts. Rosati's insights about the problems facing naturalist reductions of the normative once they are viewed from the point of view of the deliberating, autonomous agent (together with the realization that reductions in terms of what is constitutive of agency will do no better) can serve to show that our best attempts at deciding what to do do require irreducibly normative facts. And Korsgaard's insight about the unavoidability of deliberation can help to show that deliberative indispensability is in no way less respectable than explanatory indispensability, so that arguments from deliberative indispensability work if inferences to the best explanation do.

Now, if the details of such an argument can be filled in, it will yield a fairly robust version of metanormative realism, one that is not internalist enough for Korsgaard, Velleman, and Rosati, and not naturalist enough for Rosati (and probably Velleman as well). But I do not think this is a shortcoming.

This is all very sketchy, of course.⁵⁶ For the argument to ground the intended conclusion—Robust Metanormative Realism—it must be shown that the deliberative project is indeed nonoptional in the relevant sense; that irreducibly normative truths are indeed deliberatively indispensable; that indispensability arguments—whether explanatory or deliberative—can justify existential conclusions;⁵⁷ and that no other—less than robustly realist—metanormative view can adequately accommodate deliberation. And for the defense of Robust Realism to be complete, a host of familiar objections to it must be addressed (objections from the

56. I develop this line of thought in Enoch forthcoming and in more detail in Enoch 2003. And see Regan 2003 for a closely related account of normativity.

57. For more on this, see Enoch and Schechter n.d.

purported queerness of normative facts; from their supervenience; from the epistemology of the normative; from the relations between normativity and motivation; and from normative disagreement, to mention the most influential ones). I believe these are all challenges that can be satisfactorily met, but I do not want to pretend that they are not genuine challenges.

Even without the fuller account, though, it should be clear that much that is of interest in the writing of the theorists who attempt to ground normativity in what is constitutive of agency can be accommodated by a very different—robustly realist—picture of the normative and indeed can be utilized in an argument supporting this alternative picture. That this is so serves to weaken whatever intuitive plausibility the constitutivist line may have had. And that normativity cannot be grounded in what is constitutive of agency can serve to strengthen the suspicion that nothing less than a robustly realist view of the normative can withstand the scrutiny of the deliberating agent or—consequently—the criticism of the metanormative theorist.

References

- Blackburn, Simon. 1993. *Essays in Quasi-Realism*. Oxford: Oxford University Press.
- Enoch, David. 2001. "Once You Start Using Slippery Slope Arguments, You're on a Very Slippery Slope." *Oxford Journal of Legal Studies* 21: 629–47.
- . 2003. "An Argument for Robust Metanormative Realism." PhD diss., New York University, law.mscc.huji.ac.il/law1/newsite/segel/enoch/index.html.
- . 2005. "Why Idealize?" *Ethics* 115: 759–87.
- . Forthcoming. "An Outline of an Argument for Robust Metanormative Realism." In *Oxford Studies in Metaethics*, vol. 2. Oxford: Oxford University Press.
- Enoch, David, and Joshua Schechter. n.d. "How Are Basic Belief-Forming Methods Justified?" Unpublished manuscript.
- Gewirth, Alan. 1978. *Reason and Morality*. Chicago: University of Chicago Press.
- Gibbard, Allan. 1999. "Morality as Consistency in Living: Korsgaard's Kantian Lectures." *Ethics* 110: 140–64.
- Hookway, Christopher. 1999. "Modest Transcendental Arguments and Sceptical Doubts: A Reply to Hooker." In Stern 1999, 173–87.
- Hussain, Nadeem. 2004. "Discussion—'The Guise of a Reason'." *Philosophical Studies* 121: 263–75.

- Korsgaard, Christine. 1996. *The Sources of Normativity*. Cambridge: Cambridge University Press.
- . 2002. "Self-Constitution: Agency, Identity, and Integrity." The Locke Lectures given at Oxford University, May and June 2002, www.people.fas.harvard.edu/~korsgaard.
- . n.d. "Normativity, Necessity, and the Synthetic a priori: A Response to Derek Parfit." Unpublished manuscript, www.people.fas.harvard.edu/~korsgaard/Korsgaard.on.Parfit.pdf.
- Lavin, Douglas. 2004. "Practical Reason and the Possibility of Error." *Ethics* 114: 424–57.
- Lewis, David. 1996. "Desire as Belief II." *Mind* 105: 303–13. Reprinted in *Papers in Ethics and Social Philosophy*, 55–67. Cambridge: Cambridge University Press, 2000.
- Marmor, Andrei. 2001. *Positive Law and Objective Values*. Oxford: Clarendon.
- Mele, Alfred. 2004. "Discussion—Velleman in Action and Agency." *Philosophical Studies* 121: 249–61.
- Moore, Adrian. 1999. "Conative Transcendental Arguments and the Question Whether There Can Be External Reasons." In Stern 1999, 271–92.
- Nagel, Thomas. 1970. *The Possibility of Altruism*. Princeton, NJ: Princeton University Press.
- . 1997. *The Last Word*. Oxford: Oxford University Press.
- Nietzsche, Friedrich. 1967. *The Will to Power*, trans. Walter Arnold Kaufmann and R. J. Hollingdale, ed. Walter Arnold Kaufmann. New York: Vintage.
- O'Hagan, Emer. 2004. "Practical Identity and the Constitution of Agency." *Journal of Value Inquiry* 38: 49–59.
- Railton, Peter. 1997. "On the Hypothetical and Non-Hypothetical in Reasoning about Belief and Action." In *Ethics and Practical Reason*, ed. Garrett Cullity and Berys Gaut, 53–79. Oxford: Clarendon.
- Regan, Donald. 2003. "How to Be a Moorean." *Ethics* 113: 651–77.
- Rosati, Connie. 2003. "Agency and the Open Question Argument." *Ethics* 113: 490–527.
- Rysiew, Patrick. 2002. "Reid and Epistemic Naturalism." *Philosophical Quarterly* 52: 437–56.
- Schechter, Joshua, and David Enoch. Forthcoming. "Meaning and Justification: The Case of Modus Ponens." *Noûs*.
- Setiya, Kieren. 2003. "Explaining Action." *Philosophical Review* 112: 339–93.
- Shafer-Landau, Russ. 2003. *Moral Realism: A Defence*. Oxford: Oxford University Press.
- Shah, Nishi. 2003. "How Truth Governs Belief." *Philosophical Review* 112: 447–82.
- Shah, Nishi, and J. David Velleman. Forthcoming. "Doxastic Deliberation." *Philosophical Review* 114, no. 4.

- Sharon, Assaf. n.d. "Unreflected Glory: A Challenge to Metanormative Constructivism." Unpublished manuscript.
- Stern, Robert, ed. 1999. *Transcendental Arguments: Problems and Prospects*. Oxford: Oxford University Press.
- Sturgeon, Nicholas. 2003. "Moore on Ethical Naturalism." *Ethics* 113: 528–56.
- Velleman, J. David. 1989. *Practical Reflection*. Princeton, NJ: Princeton University Press.
- . 2000a. *The Possibility of Practical Reason*. Oxford: Oxford University Press.
- . 2000b. "Introduction." In Velleman 2000a, 1–31.
- . 2000c. "On the Aim of Belief." In Velleman 2000a, 244–81.
- . 2000d [1996]. "The Possibility of Practical Reason." In Velleman 2000a, 170–99 (originally appeared in *Ethics* 106: 694–726).
- . 2000e [1998]. "Is Motivation Internal to Value?" In Velleman 2000a, 85–98.
- . 2004a. "Precis of *The Possibility of Practical Reason*." *Philosophical Studies* 121: 225–38.
- . 2004b. "Replies to Discussion on *The Possibility of Practical Reason*." *Philosophical Studies* 121: 277–98.
- Weintraub, Ruth. 1997. *The Sceptical Challenge*. London: Routledge.
- Wright, Crispin. 1991. "Scepticism and Dreaming: Imploding the Demon." *Mind* 100: 87–116.