


This article was downloaded by: 10.3.97.143

On: 28 Mar 2023

Access details: *subscription number*

Publisher: *Routledge*

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: 5 Howick Place, London SW1P 1WG, UK



## The Routledge Handbook of Metaethics

Tristram McPherson, David Plunkett

### Normativity and Agency

Publication details

<https://www.routledgehandbooks.com/doi/10.4324/9781315213217.ch25>

Hille Paakkunainen

**Published online on: 28 Aug 2017**

**How to cite :-** Hille Paakkunainen. 28 Aug 2017, *Normativity and Agency from: The Routledge Handbook of Metaethics* Routledge

Accessed on: 28 Mar 2023

<https://www.routledgehandbooks.com/doi/10.4324/9781315213217.ch25>

**PLEASE SCROLL DOWN FOR DOCUMENT**

Full terms and conditions of use: <https://www.routledgehandbooks.com/legal-notices/terms>

This Document PDF may be used for research, teaching and private study purposes. Any substantial or systematic reproductions, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The publisher shall not be liable for an loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

## Normativity and Agency

Hille Paakkunainen

Normative facts, such as the fact that you ought to do something or have good reason to do it, are often thought to bear some important general connection to facts about agency. Most non-normative facts, such as facts about the material composition of the Sun, aren't usually thought to do so. Why think that the normative bears some important general connection to agency, where the non-normative, as such, doesn't? And what forms might this connection take?

I start by explicating the phenomena of normativity and of agency at issue ("Normativity and Agency: Some Basics"). I then focus on two putative connections between normativity and agency present in the literature that are, I think, insufficiently understood, but often heavily inform one's further views: certain types of "open question" argument (OQA) ("Agency and Deliberative Open Questions"), and the claim that normative reasons for action are premises in good deliberation ("The Deliberative Constraint: Doubts, Merits, and Conceptual Choices"). I argue that while OQAs don't seem to capture any important general connection between normativity and agency, the Deliberative Constraint does; at least, it captures an important way in which *some* central normative facts depend on agency. The final section, "Further Connections and Topics," briefly connects the Constraint to further putative connections between normativity and agency present in the literature, and sketches some avenues for further exploration.

### NORMATIVITY AND AGENCY: SOME BASICS

We can distinguish various normative phenomena. One broad but helpful contrast is between "formal" *versus* "robust" normativity. What it is for a standard S to be formally normative is for S to be such that one can violate it, or make a mistake by its lights (McPherson 2011: 232). For example, traffic rules and various rules of games are for-

mally normative in this sense. Formal normativity is cheap: we can create new formally normative standards simply by inventing violable rules.

Robust normativity is a seemingly more important phenomenon that many take to be associated with *normative reasons*. Normative reasons for a response  $\Phi$  are facts or true propositions,  $p$ , that genuinely justify, favor, or call for that response, at least *pro tanto*. The fact that a tree fell on my neighbor calls for my response of helping him. (Here, the fact,  $p$ , that is the reason to help is the ordinary non-normative fact that a tree fell on my neighbor. It's the further fact that  $p$  is a reason for me to help that is a normative fact, and  $p$ 's property of *being a reason for action* is the normative property. Cf. Dancy 2006: 137, Parfit 2011 Vol. 2: 330–331.) If I failed to help, I would not only be violating a formally normative standard—the “reasons standard,” as we might call it (McPherson 2011: 232)—but moreover, I'd be violating a standard that has special *authority* or *normative importance* (McPherson 2011: 233). At least, I'd be violating such a standard by failing to help if my reasons to help were *decisive*, so that I *ought* to help, in a robustly normative sense. Whether we act as we have (decisive) reasons to act seems somehow more important than whether we follow or violate the rules of any old game. ‘Robust normativity’, ‘normative authority’, and ‘normative importance’ are labels for this intuitively familiar but difficult-to-characterize type of normativity that is more important than mere formal normativity.

It's robust normativity in particular whose connection to agency I'll discuss. I'll mostly focus on normative reasons and connected ought-facts, assuming with much of the literature that these are central robustly normative phenomena; the section “Agency and Deliberative Open Questions” also discusses the property of *being good for a person*, where this property is linked to reasons in a particular way. I'll stay neutral regarding whether some reasons and oughts have a distinctively *moral* flavor. And I'll focus on reasons for *action*, and oughts concerning actions, where actions are paradigmatic exercises of agency; leaving aside reasons for belief or emotion.

What, then, is agency in the relevant sense? It's not just being a *self-mover*, in the Aristotelian sense of a being that has its principle of movement within itself. Cockroaches are self-movers, since unlike (say) fallen leaves, they can move under their own steam, without being pushed by external forces such as wind. Our concern is with *minimally rational* agency—agency that's rational in the sense of contrasting with *non-rational*—in its various guises: practical reasoning, acting on the basis of considerations, reflecting about what one ought to do and acting on one's subsequent normative judgments. (This is a mere characterization to get an intuitive grip on the phenomenon; not a purported definition.) I'll call minimally rational agency simply ‘agency’. When agents perform some agential activity such as practical reasoning *well*, they perform it in accordance with whatever its proper standards of excellence are. I'll say that these agents are being “rational” (in some respect), where the relevant sense of rationality contrasts with *irrationality*. “Fully rational” agents are hypothetical agents purged of all irrationalities. Different ways of connecting normativity to agency focus on slightly different agential phenomena. I'll comment on the different foci as we go.

A final orienting remark. In assuming that there are robustly normative facts, I assume the falsity of important versions of normative nihilism and error theory, on which there are no such facts even if we talk as if there are. But I *don't* assume that robustly normative facts are mind-independent or agency-independent: precisely not. They might be constructed out of, or constituted by, actual or hypothetical psychological facts, or bear some more

modest connection(s) to agents' minds. This is an important terminological warning, as the term 'robust' is sometimes used for a kind of "realism" about normative facts on which they're radically mind-independent (Enoch 2011, FitzPatrick 2008). It's a substantive, disputed question what degree of mind-dependence or independence is needed for capturing 'robust' normativity in the sense of normative authority or importance. In examining forms of agency-dependence below, we examine, in part, this disputed question. (For error theory, see Jonas Olson's chapter "Error Theory in Metaethics." For realism, see David Enoch's chapter "Non-Naturalistic Realism in Metaethics," Peter Railton's chapter "Naturalistic Realism in Metaethics," and Billy Dunaway's chapter "Realism and Objectivity.")

### AGENCY AND DELIBERATIVE OPEN QUESTIONS

How, then, might robust normativity be connected to agency? *Normative judgment internalism* holds that *judgments* about (putative) normative facts bear a necessary connection to agency—specifically, to motivation (see David Faraci and Tristram McPherson's chapter "Ethical Judgment and Motivation"). On a popular way of fleshing out this idea, necessarily, if one sincerely judges that one ought to  $\Phi$ , then if one is fully rational, one is motivated (to some extent) to  $\Phi$ . However, our concern isn't primarily with normative *judgments* and their connection to agency, but with agency and normative *facts* (and the normative *properties* ingredient in those facts). So I set aside normative judgment internalism here.

Still, normative reflection about what one ought to do or has reason to do is plausibly, itself, one kind of exercise of agency, especially if it's such as to engender motivation or action. And some have suggested that such exercises of agency bear an important general connection to normative facts, of roughly the following form:

OQA If the question whether X is F—where F is some normative predicate—remains open from the perspective of idealized normative reflection, then this impugns the claim that X is F. (The details of the idealization depend on the theorist.)

There's a trivially true version of OQA. If the relevant reflective perspective is idealized to include a grasp of the normative truths about whether X is F (and if either X is F or X isn't F), and if the question whether X is F remains open from that perspective, then it's not the case that X is F. But is there a plausible non-trivial version of OQA? This section examines two prominent proposals to the effect that there is: Christine Korsgaard's (1996) and Connie Rosati's (2003). While I cast some doubt on each, Rosati's proposal points us towards questions about agency, normativity, and normative thought that deserve much further exploration.

Korsgaard's discussion of normative reasons is framed by a familiar skeptical doubt that can afflict agents within normative reflection. If your confidence has been shaken in whether some fact,  $p$ —say, that morality requires you to  $\Phi$ —really is a reason for you to  $\Phi$ , then we cannot shore up your confidence merely by pointing to putative reasons out there, without heed to what you think (1996: 34–40). Shoring up your confidence requires engaging your own deepest commitments, including, ultimately, the commitments constitutive of reflective agency as such; and showing that  $p$  is a reason *as judged from the perspective of those commitments* (1996: 93ff). Otherwise, thorough and

clear-eyed normative reflection will always leave it an open question for you whether  $p$  really is a reason. And Korsgaard thinks that such open questions impugn  $p$ 's status as a reason. Reasons are hostage to reflective agents' verdicts (1996: 16, 92–94 *et passim*):

OQA-K: If the question whether  $p$  is a reason for A to  $\Phi$  remains open, as judged from the perspective of A's deepest commitments, including commitments constitutive of reflective agency, then this impugns the claim that  $p$  is a reason for A to  $\Phi$ .

However, it remains unclear in Korsgaard's treatment why we should accept OQA-K. Korsgaard uses OQA-K to reject "substantive realist" views on which reasons exist, regardless of whether they speak to our commitments, and to argue for a Kantian constructivist view on which reasons are a function of agents' deepest commitments. (On constructivism, see Melissa Barry's chapter "Constructivism.") But why think that deliberative open questions are problematic in the way that Korsgaard thinks? We need some reason to think that a clear-eyed working out of the implications of our deepest commitments, even commitments constitutive of reflective agency, doesn't or can't itself embody or engender normative error, or leave room for failures to grasp normative facts that obtain regardless. (Cf. Enoch 2006.)

To be sure, it would be disturbing if clear-eyed normative reflection, regardless of its starting points, *necessarily* left us out of touch with an important portion of normative reality. We may have *prima facie* reason to defend a view of normativity on which this isn't the case. But we can reject OQA-K while holding that some contingent commitments lead to the right normative views *via* clear-eyed normative reflection. There may be a route to normative truth *via* normative reflection, even if not every agent's deepest commitments furnish such a route (cf. FitzPatrick 2008: 177–178). Of course, the easiest way to forge such a route is by building in a grasp of general normative truths or values into the relevant reflective perspective—a grasp that some agents may be lucky to have acquired, perhaps through good moral education (*ibid.*). Again, this front-loads normative truths into the relevant reflective perspective, without seeming to provide any interesting agency-related constraint on normativity.

In sum, we lack a good argument for OQA-K, and while it may seem partly supported by a felt need to view normative facts as not ineluctably outside of agents' reflective reach, OQA-K isn't needed to satisfy this putative need.

Rosati (2003) employs a different version of OQA, arguing that extant naturalist accounts of the property *being good for a person* fail because they leave open questions within fully autonomous, empirically well-informed normative reflection. Take the view that *being good for a person* is some natural property N—say, *being pleasurable*. Rosati's OQA is supposed to challenge such views, and also help us diagnose what goes wrong with them. Specifically, it purports to help us identify "some feature of our notion *good for person A* that the proposed account misses and that cannot [...] be abandoned without sacrificing our ability to ask and answer the questions that ordinarily concern us when we wonder what is good for us" (2003: 501). The feature of personal good that we mustn't abandon is, roughly, its "fit" with our capacity for autonomous normative reflection and subsequent action (2003: 507). As I read Rosati, the presence of open questions within autonomous reflection is supposed to indicate failure to secure the requisite "fit." In particular, if the question whether X is good for us, given that X is N, remains open within

fully autonomous and empirically well-informed normative reflection, then the relevant “fit” hasn’t been secured, and the account of personal good as Nness must be rejected (ibid.). Generalized for any robustly normative property F and property N:

OQA-R If the question whether X is F given that X is N remains open within fully autonomous and empirically well-informed normative reflection—so that Fness, on the account of it as Nness, doesn’t “fit” with our capacity for autonomous reflection—then this impugns the claim that X really is F (even though X is N) (thereby impugning the theory that Fness is Nness).

To clarify, Rosati isn’t proposing an account of personal good in terms of autonomous reflection, but a necessary condition on accounts of personal good: an adequate account must construe personal good as “fitting” with our agency so as to close the relevant reflective question (2003: 519; cf. Rosati 2016). OQA-R forces us to refine our accounts if, in their present forms, they fail to secure the requisite “fit” (2003: 524). (Notice that nothing in Rosati’s challenge turns on N’s being a natural property. Cf. Rosati 2003: 517.)

Why think that such a “fit” with agency *is* a necessary condition on personal good (or on some other normative property F)? Recall Rosati’s claim that without the relevant “fit,” we sacrifice “our ability to ask and answer the questions that ordinarily concern us when we wonder what is good for us” (2003: 501, 507). Let’s unpack this idea.

As I read Rosati, our “ordinary concern” with personal good has two related facets that together require the relevant “fit” with agency. First, in wondering what’s good for A, we’re partly concerned with what we have reason to do, or ought to do (2003: 516–520). In ordinary thought and talk about personal good, there are “inferential links” between propositions about personal good and about reasons and obligations: for instance, if X is good for me, I have “at least a [*pro tanto*] reason to obtain X;” and “others have at least a [*pro tanto*] obligation to support” my efforts to obtain X (2003: 517). A second, related facet is that judgments of personal good, and related judgments about reasons and obligations, “function as regular guides to action and attitude” (2003: 520). And they do so not by merely causing bodily movements or attitudes, but *via* engaging those reflective, critical capacities that make us self-governing, autonomous agents (2003: 518, 520). It’s through autonomous normative reflection that we can step back from a moment’s impulse and ask what we really have reason to do, or ought to do, and be moved accordingly. (Judgments about wrongness plausibly enjoy similar inferential links, and a similar action-guiding role; if so, Rosati’s argument should generalize to some extent. Cf. Rosati 2016.)

Rosati infers that the “content” of propositions about personal good must be “such as to engage [our] autonomy-making motives and capacities” (2003: 518–519).<sup>1</sup> Our account of these propositions’ content must make sense of why judgments about personal good engage our autonomous agency by engendering related judgments about reasons and obligations, and by motivating action (ibid.). Rosati further infers that accounts of the *property of being good for a person* shouldn’t “leave it mysterious why” the relevant features of ordinary thought and talk obtain (2003: 517, 522). The truth-makers for propositions about personal good should help explain why judgments about personal good exhibit the highlighted inferential and action-regulating behaviors. This, I take it, is the relevant “fit” with our agency that accounts of personal good should secure.



Suppose Rosati is right about the need for such a “fit.” What does this have to do with reflective open questions? The presence of such open questions is supposed to indicate that the relevant “fit” is missing. To illustrate how, Rosati considers a toy account of personal good as blueness. While something’s being good for me provides me with a reason to pursue it, and I would generally judge as much within autonomous reflection, autonomous and empirically well-informed agents can reasonably doubt whether they have any “reason to pursue [their] good as specified by an account of personal goodness as blueness” (2003: 518). The judgment that X is blue doesn’t generally engender the judgment that one has reason to pursue X, nor, relatedly, does it engender motivation to pursue X, within autonomous reflection. Personal good construed as blueness palpably fails to engage our autonomous agency in the way that personal good is supposed to do. The account of personal good as blueness accordingly seems to do nothing to explain the inferential and action-regulating behaviors of judgments of personal good as made by autonomous agents. I take it that this is how the presence of open questions is supposed to indicate that a proposed account fails to secure the relevant “fit” with autonomous agency.

Moving beyond toy accounts, Rosati presses the same problem against accounts of personal good as pleasure, or as a function of informed desires (2003: 518–519). Rosati suggests that we can close the troublesome reflective questions only by construing personal good as a “relational property” that’s partly defined in terms of “some connection to what makes persons self-governing actors” (2003: 521). Rosati suggests that a dispositional theory on which *being good for a person* is “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,” might do the trick (2003: 520).

Rosati’s argument deserves more discussion than I can provide here. One worry is that it’s unclear how the suggested dispositional theory *does* close reflective questions. Can’t a fully autonomous, empirically well-informed agent find it an open question whether X is good for A even though X is such as to be approved by her when her “autonomy-making motives and capacities operate effectively and other appropriate conditions obtain” (ibid.)? One might, after all, doubt whether one’s autonomy-making motives and capacities make one approve of the right things. It’s unclear why mere empirical information and autonomy would dispel the doubt. (Cf. Enoch 2006: 178, 180–185.) (Of course, much may depend on what “appropriate conditions” are.)

However, the presently most relevant worry concerns the role of open questions in Rosati’s argument. What’s ultimately important for Rosati seems to be just the claim that personal good should “fit” with our agency, helping to explain why judgments of personal good, as made by autonomous agents, exhibit the inferential and action-regulating behaviors they do. This demand for a “fit” may be apt, even if the presence of open questions doesn’t always indicate that the “fit” is missing. Suppose that Fness is Nness, but that seeing why requires grasping a convoluted theoretical argument. And suppose that, having appreciated the argument, one can see that the account of Fness as Nness *does* help to explain the inferential and action-regulating behaviors of judgments about Fness. It’s unclear why these suppositions couldn’t be true and the relevant convoluted argument sound, even while fully autonomous, empirically well-informed agents can doubt whether something that is N is also F; and relatedly, can doubt whether something’s being N gives one reasons to act. After all, full autonomy and empirical information presumably don’t ensure grasp of the envisaged convoluted argument. Not having grasped the argument,

one might fail to realize that something's being N amounts to its being F, and thus gives one reasons to act.

Of course, if full autonomy and empirical information *did* ensure grasp of the relevant argument, or belief in the account, then the question whether, given that X is N, X is F, *would* be closed within autonomous reflection. But this doesn't give us any independent grip on whether the account is true, or the envisaged argument sound. Either way, the presence or absence of open questions within autonomous reflection isn't a good adequacy test for accounts of Fness as Nness. Such open questions may function as a fallible heuristic device for directing attention to perceived shortcomings of an account. But whether the shortcomings are real or merely apparent depends on the availability of arguments that agents needn't, just in virtue of being autonomous and empirically well-informed, appreciate.

In sum, I doubt that OQA-R provides a true, non-trivial constraint on accounts of normative properties. If autonomy together with empirical information somehow ensures a grasp of an account of Fness as Nness, then the account trivially passes the test of OQA-R. But to establish the truth of the account, we still need an independent argument. On the other hand, autonomy and empirical information might not ensure a grasp of even a true account that helps to explain (in some non-obvious way) why judgments of personal good engage our agency in the ways they do. If so, reflective open questions don't tell against the account.

None of this challenges Rosati's argument that accounts of normative properties should "fit" with autonomous agency, doing something to explain why judgments about those properties engage agency in the ways they do. This argument deserves further attention. One question to consider in assessing it is whether the nature of *agency* might explain the inferential and action-regulating behaviors of normative judgments within autonomous reflection; perhaps the nature of normative properties needn't play any explanatory role (cf. Rosati 2003: 525–526). Further, our ordinary normative *concepts*, such as the concept of personal good or of a reason for action, might be inferentially articulated in a way that explains the relevant behaviors of judgments deploying those concepts, even if the properties that these concepts pick out can also be picked out by means of other concepts that lack the same inferential articulation, and even if the properties themselves don't explain the inferential and action-guiding behaviors of normative judgments. (See Matthew Chrisman's chapter "Conceptual Role Accounts of Meaning in Metaethics" for some relevant discussion.) Still, I suggest that these questions are best pursued separately from the concern with open questions.

### THE DELIBERATIVE CONSTRAINT: DOUBTS, MERITS, AND CONCEPTUAL CHOICES<sup>2</sup>

OQA connects normativity to idealized normative reflection—idealized reflection on the explicitly normative question whether X is F. A different and common way to connect normativity to agency focuses on agents' deliberation in view of, and motivation by, the non-normative facts, *p*, whose status as reasons explicit normative questions can concern. For instance, Mark Schroeder proposes the "Deliberative Constraint" that when an agent



“is reasoning well, the kinds of thing about which he should be thinking are his reasons” (2007: 26; cf. 33). Kieran Setiya sees it as a “harmlessly illuminating” starting point for further theorizing to construe reasons for  $A$  to  $\Phi$  as premises for “sound reasoning to a desire or motivation to  $\Phi$  whose further premises are available to  $A$ ” (2014: 221). Jonathan Way says that it’s “near platitudinous” that “a reason for you to  $\Phi$  must be an appropriate premise for reasoning towards  $\Phi$ -ing” (2015: 1). There are different ways to flesh out these in this area. But the basic idea I’m interested in is that reasons are considerations that non-normatively well-informed good deliberation takes into account; and if the reasons are decisive, it’s part of good deliberation to be moved to act on them, in the way that they support. In slightly more detail, and borrowing Schroeder’s label:

**Deliberative Constraint:**

If  $p$  is a reason for  $A$  to  $\Phi$ , then there’s a possible course of non-normatively well-informed good deliberation such that, were  $A$  to undergo it,  $A$  would take  $p$  into account; and if  $p$  is a decisive reason to  $\Phi$ , then it’s part of the relevant course of good deliberation to be moved to  $\Phi$  on the basis of  $p$ .

Those who explicitly articulate versions of the Constraint often see it as rather obvious. But even more often, the Constraint operates quietly in the background, unremarked-upon. Much of the literature on “existence internalism” about reasons is an example of this. (On internalism, see Errol Lord and David Plunkett’s chapter “Reasons Internalism,” and Paakkunainen forthcoming.) On Bernard Williams’ (1981) famous internalist view,  $p$  is a reason for  $A$  to  $\Phi$  only if there’s a broadly instrumental “sound deliberative route” from  $A$ ’s “subjective motivational set”  $S$ —from  $A$ ’s current contingent set of desires, dispositions, and projects—to  $A$ ’s  $\Phi$ -ing or being motivated to  $\Phi$  on the basis of the consideration,  $p$ , that is the reason. Early influential critics of Williams take issue with the ideas that sound deliberation must start from merely contingent motivational elements, and that the standards of soundness in deliberation are merely instrumental, suggesting that soundness in deliberation may involve satisfying standards of Kantian universalization or ethical virtue that demand more than mere instrumental cogency (Korsgaard 1986, 1996, 2009; McDowell 1995). Neither Williams nor these critics, however, question the idea that reasons are premises in sound deliberation. Indeed, this is why they take the question of what sound deliberation involves to impact the fate of views about reasons.

Beyond debates over internalism, the background influence of the Constraint should be evident to anyone who goes looking for it in the vast literature on reasons for action. But we lack a good sense of why, if at all, the Constraint is true. Why should the normative support relations between considerations and the actions they support correspond to good deliberation, or to patterns of rational response to information? It’s a fairly common view that the ultimate moral justifications of actions, as stated by general moral theories, aren’t what morally good deliberation focuses on. Morally good deliberators are often moved by considerations of loyalty or love, such as “she’s my wife,” not by considerations such as “this act would maximize utility,” or whatever the ultimate right-makers are supposed to be (Railton 1984). If normative reasons can’t likewise be considerations that good, rationally excellent deliberation would ignore, it’s unclear why.

Indeed, some have proposed examples that look to defeat the Constraint.

*Surprise Party* (paraphrased from Schroeder 2007: 33)

The fact,  $p$ , that there's a surprise party waiting for Nate at home is a reason for Nate to go home, but only if Nate doesn't believe  $p$ . For while Nate loves successful surprise parties thrown in his honor, he hates unsuccessful ones: his believing  $p$  would destroy  $p$ 's status as a reason to go. Since taking  $p$  into account in deliberation requires at least believing  $p$ , whatever else it involves,  $p$  is a reason for Nate only if Nate doesn't take it into account in deliberation; and so only if Nate doesn't take it into account in good deliberation.

*Deluded Belief* (versions of this are in Markovits 2014: 41, Smith 2009: 523, Johnson 2003: 575)

The fact,  $p$ , that I have the deluded belief that I'm Jesus is a reason for me to seek psychiatric help. But  $p$  is a fact, and so a reason, only if I don't believe it, and so only if I don't take it into account in deliberation. For were I to believe that I have the deluded belief that I'm Jesus—that is, were I to think of my belief that I'm Jesus *as* deluded—this would do away with my delusion.

*Emergency Landing* (paraphrased from Markovits 2014: 48)

Captain Sullenberger performed a successful emergency landing of a commercial airliner on Hudson River. The fact,  $p$ , that so many lives were at stake was a decisive reason for Sullenberger to attempt emergency landing. But he didn't think about this reason on the way to performing the landing—and a good thing too, for he shouldn't have. Were he to have thought about it, he would likely have become very unnerved and endangered the success of the landing.

In *Surprise Party* and *Deluded Belief*, the putative reasons,  $p$ , are facts that A *can't* take into account in good deliberation without destroying their status as reasons, or as facts (and so as reasons).<sup>3</sup> In *Emergency Landing*, the putative reasons are ones that A *shouldn't* take into account; we might infer that it can't be part of good deliberation for A to take them into account and be moved by them. If the Constraint survives these examples, we must spell out why.

The rest of this section explains why these examples don't defeat the Constraint, and sketches a positive argument for the Constraint that seeks to explain its appeal.

Start with *Emergency Landing*. Plausibly the fact,  $p$ , about the lives at stake is a reason, indeed a *decisive* reason, for Sullenberger to attempt emergency landing. For given  $p$ , Sullenberger definitely *ought* to attempt it, in the robustly normative sense. Sullenberger would be failing an important authoritative demand on him if he didn't attempt landing but, say, started doing crossword puzzles instead. However, even if Sullenberger would likely get dangerously unnerved by considering the reason-giving fact,  $p$ , it doesn't follow that it wouldn't be good deliberation if he took  $p$  into account and were thereby moved to act in the way that  $p$  supports. It merely shows that Sullenberger is unlikely to deliberate well, given his nerves. (Compare: it's no part of Williams' internalism to deny that agents can fail, or be likely to fail, to undergo the sound deliberative routes to which their reasons correspond.)

Of course, if one is very likely to get unnerved by considering the reason-giving facts, perhaps in some sense one ought to ignore them. Still, this is a case where one ought, in some sense of ‘ought’, to not deliberate well; not a case of reasons for action that good deliberation would ignore. It’s intuitively a case of good deliberation, a recognizably rational response to the fact about lives at stake, to attempt emergency landing on the basis of one’s appreciation of that fact. The possibility of irrational or non-rational responses to the reason-giving facts shouldn’t obscure this point.

Further reflection on the case supports rather than challenges the Constraint. The lives at stake are a reason for the *whole* action of attempting landing. Once this action is in progress, more relevant are facts about how specific maneuvers will affect the plane’s behavior, about what one needs to communicate to whom, etc. These further facts are reasons for sub-actions, the actions by means of which one lands. It’s to be expected that good deliberation while the landing is in progress focuses on these further facts, and doesn’t dwell on the initial reason to attempt landing. And surely it *is* intuitively a case of good deliberation to take account of these further facts as one is landing, and to be moved by them to perform the maneuvers that they support. Certainly, this is intuitively better deliberation than choosing one’s maneuvers based on, say, facts about today’s crossword puzzle, or about the state of the stock market.

One might object that these claims are hard to assess without a specific account of “good deliberation” in view. Further, if good deliberation is just a matter of responding to the reason-giving facts by doing what they support, then the Constraint looks trivial and uninformative.

In response, it is indeed part of the operative notion of good deliberation that it involves taking account of the reason-giving facts, and responding to them by doing what they support. The Constraint isn’t supposed to link reasons to some notion of good deliberation that’s completely independent of the idea of appropriate responsiveness to one’s reasons. Instead, the Constraint articulates a dependence between two putatively interconnected notions: normative reasons for action, on the one hand, and rational responsiveness to the considerations that are the reasons, on the other. In doing so, it articulates (part of) a *conception of reasons as linked to such rational responsiveness*. As we’ve seen, one might doubt this conception: one might doubt that reasons, as properly conceived, need correspond to anything that looks like an intuitively rational response to the reason-giving facts. Examples such as *Surprise Party* and *Deluded Belief* make this most evident. Further, the idea of rational responsiveness to the reason-giving facts isn’t empty. We have some grasp of when a response to considering some facts is a rational one, and when it’s an irrational or non-rational one. One can deploy this grasp in assessing my claims about the *Emergency Landing* case, above.

Of course, there’s a sense in which, *if* we accept the conception of reasons that the Constraint articulates, then the Constraint will look trivial as applied to reasons thus conceived. But the question of interest is whether and why we should accept the conception; or whether there’s some good reason, based perhaps on some independent, contrary intuitions concerning normative reasons, to reject the conception. Thus far I’ve argued, in effect, that *Emergency Landing*-style cases provide no good basis for rejecting the conception of reasons as linked to good deliberation.

None of this rules out that “good deliberation” can be further precisified. Does it involve normative ascent—thinking of the reason-giving facts under a normative guise, *as*

reasons? Does it proceed in accord with an instrumental rule of deliberation, such that one is moved by those considerations that reveal courses of action as promoting the objects of one's desires? There may be good arguments for some such precisifications—perhaps partly based on what deliberation as such can involve. For example, if deliberation as such can only take an instrumental form, the same applies to good deliberation. If the Constraint holds, such further constraints on good deliberation are also constraints on reasons. Still, we can make some progress in assessing the basic idea of the Constraint without examining such further precisifications.

What, then, of cases such as *Surprise Party* and *Deluded Belief*? In these cases, there's supposed to be no such thing as responding rationally to the reason-giving facts, while the relevant reasons are in force. If so, some reasons violate the Constraint. If we find this verdict intuitive, we find in these cases some intuitive basis for rejecting the Constraint's conception of reasons.

However, we should consider different interpretations of our intuitive reactions to these cases. In *Surprise Party*, there's clearly something to be said for Nate's going home. But it's unclear why we should interpret this intuition as indicating the presence of the specific alleged normative reason for Nate to go, instead of indicating (merely) some other normative or evaluative phenomena. The relevant fact,  $p$ , might be an explanatory reason, one that explains why Nate's going home would be a good outcome from the perspective of Nate's preference-satisfaction. Or it might be a normative reason for Nate's friends to urge him to go; or a normative reason for Nate to be glad, should he end up home: a reason for an affective response, should a pleasant outcome occur. The fact that Nate would be glad if he went might itself be a reason for Nate to go, a reason that Nate *can* take into account and act on. (*Mutatis mutandis* for *Deluded Belief*.) These hypotheses are all compatible with the Constraint. It's not pre-theoretically clear that we should adopt an anti-Constraint interpretation of the cases instead. Our tendencies to apply the word 'reason' may, if unexamined, elide important distinctions between different normative or evaluative notions that the Constraint actually helps us mark. And even if we think, say, that agents' reasons are ultimately a function of their preference-satisfaction—even a function that allows for violations of the Constraint—this is a theoretical claim that's at least as controversial as the Constraint.

If so, the examples on their own don't defeat the Constraint. Still, is there a good positive argument for the Constraint, and so for interpreting the examples in Constraint-friendly ways?

The rest of this section sketches an argument that there are certain key jobs in theorizing and in everyday normative practice that the notion of reasons is often pressed to perform, and that these jobs are best performed by reasons as conceived by the Constraint. This leaves open that there are other important theoretical or everyday jobs that only some incompatible conception of reasons can perform. But if so, we should be explicit about what these jobs are.

Recall that part of what's important about reasons is their connection to normative authority. If  $A$  has decisive reasons to  $\Phi$ , then  $A$  ought to  $\Phi$ , in the robustly normative sense; decisive reasons pose authoritative demands. A qualification: some reasons, even if not outweighed by contrary reasons, might merely "entice" or recommend (Dancy 2004: 21). Still, it's an important job for the notion of reasons that some reasons can impose authoritative demands. Consider disputes about morality's authority. These disputes usually proceed on the assumption that moral requirements have authority on  $A$  if, and only

if, A has some normative reason to do what morality requires. Part of the thought is that, without such reasons, A is free to take or leave morality as a guide to action. However, if even strong reasons to abide by moral requirements left it entirely normatively optional whether to do so—as they would, if they merely recommended without demanding anything—this wouldn't be much of a vindication of morality's authority. We would *still* be free to take or leave morality as a guide to action. This suggests that the operative conception of reasons in disputes about morality's authority is, at least implicitly, one on which reasons can impose not just recommendations but demands. And plausibly, this concern with “reasons to be moral” has roots in common-sense concerns about morality's authority. There are important theoretical and everyday jobs for a conception of reasons as capable of imposing authoritative demands.

The Constraint is plausibly a condition of making sense of reasons as imposing authoritative demands. This is because of a link between authoritative demands and *reasonable expectations*. If the facts in a situation demand of A that she  $\Phi$ , then there must be some possible condition, X, that A might be in in the situation, such that if A encounters the situation while in that condition, then we can reasonably expect of A that she will  $\Phi$ . The relevant notion of reasonable expectations is partly predictive, partly normative. The predictive part is that, if A is in condition X in a situation, it's predictable that A will end up doing precisely what the facts in the situation demand of her: condition X *well-equips* A to meet authoritative demands. Without some such possible condition X, the idea that A is under an authoritative demand to  $\Phi$  seems to lapse. A might of course happen, by sheer accident, to do precisely  $\Phi$ . And we might apply some positive evaluative predicate to such happy accidents. But authoritative demands to  $\Phi$  require that there be some possible condition that A might be in (even if she's not in that condition, in fact) that well-equips A to meet those demands. So it seems to me.

Here is the normative aspect of the notion of reasonable expectations: if A is under an authoritative demand to  $\Phi$ , then it must be reasonable to demand it *of* her that she  $\Phi$ s, and to criticize her if she fails to  $\Phi$ —at least absent excusing conditions. (For instance, non-culpable ignorance of relevant non-normative facts plausibly excuses.) The condition X that well-equips A to meet authoritative demands must be of a type that makes sense of the appropriateness of such criticism. And roughly, it seems that A is criticizable only for what she does or fails to do under her own steam, *via* the exercise of her capacities for agency. We no more criticize A for involuntary non-rational twitches than we criticize rocks for their behavior—even though there are conditions of rocks that well-equip them for certain behaviors in certain circumstances. It seems, then, that X must be some deliberative condition of A's: a condition in which A acts on considerations, takes in information, and makes choices in its light, in a way aptly describable as an exercise of agency.

In sum, X must be (a) a possible deliberative condition of A's that (b) well-equips A to do precisely what the facts in the circumstance demand of her. Without some such X, the facts in the circumstance can't demand anything of A. What deliberative condition X best satisfies (b)? The obvious answer is the condition of deliberating well, where this involves taking account of the facts that impose the relevant demands, and being moved on their basis to do what they demand. It's certainly harder to see why deliberating *badly* (e.g., drawing bizarre conclusions from relevant reason-giving facts), or taking into account only facts that do nothing to support  $\Phi$ -ing, would generally lead agents to  $\Phi$  whenever the reasons so demand.



While the argument is abstract, its conclusion isn't surprising. It's *prima facie* odd to think that the facts in a situation might authoritatively demand  $A$  to  $\Phi$ , even though no possible rational response to those facts would lead  $A$  to  $\Phi$  on their basis. We usually think there's a point to trying to figure out facts relevant to what we should do before making important decisions, and to trying to make our decisions in light of those facts. We expect deliberating well in light of the reason-giving facts to help us to do what we should do (cf. Schroeder 2007: 132). The argument above gives some explanation of why decisive reasons and connected oughts are linked to good deliberation in these intuitive ways.

The argument is merely a sketch. A proper treatment would respond to objections, consider different types of reasons, and so on. (For instance, if there are "ultimate" and "derivative" reasons, perhaps good deliberation need only take account of one or the other; cf. Star 2015.) Recall, too, that the argument leaves room for alternative conceptions of reasons that deny the Constraint. For all I've said, there may be further important theoretical and everyday jobs that reasons can perform only on some such alternative conception. If so, we should spell out these jobs. Either way, I doubt we can resolve disputes about reasons without such explicit attention to issues of conceptual choice, and solely on the basis of attending to when we're willing to deploy the term 'normative reason' in imagined, or even real life, cases.<sup>4</sup>

I've argued that we have intuitive need for the Constraint's conception of reasons as linked to the agential activity of good deliberation. The next section briefly connects the Constraint to further putative connections between agency and normativity, and suggests some topics for further exploration.

### FURTHER CONNECTIONS AND TOPICS

The Constraint states a necessary connection between reasons and agency, but the connection is fairly modest. It assumes nothing about the contours of good deliberation beyond its involving rational responsiveness to, and motivation in light of, the reason-giving facts. Nor does it assume that agents have a current motivational propensity to undergo the good deliberative routes to which their reasons correspond, only that there's a possible world in which they do so, while the reasons remain in force. But we might wonder whether the argument for the Constraint supports a stronger thesis. Specifically, if  $A$  isn't aptly criticizable for failures to  $\Phi$  for the reason that  $p$ , where  $\Phi$ -ing for this reason is outside of  $A$ 's current motivational reach, this may suggest a stronger, "internalist" constraint on reasons. (Cf. Lord 2015 for a somewhat similar argument for an ability condition and a further epistemic condition on reasons.)

Whether the argument might be so extended depends in part on which conditions merely excuse from otherwise reasonable criticism, and which ones make the criticism, and the associated demand to  $\Phi$ , lapse. Where the absence of appropriate motivational propensities merely excuses, the demand, and so the reasons that impose it, may stay in force. This would allow that thoroughly vicious people have decisive reasons to act in non-vicious ways, even if they lack the current motivational wherewithal to do so. Further, if it's reasonable to criticize mature agents for failures to acquire non-vicious motivational tendencies, lacking those tendencies might not even excuse. Whatever the



case may be, I suggest that the precise contours of the link between reasons and good deliberation depend partly on such questions about excuses and reasonable criticism.

One notable consequence of the Constraint is that the Advice Model of reasons is inadequate as a model of reasons capable of imposing authoritative demands. On the Advice Model, there's (decisive) reason for A to  $\Phi$  (if and) only if A+, A's fully rational and non-normatively fully informed counterpart, would, after thinking about what A is to do, *advise* A to  $\Phi$ . But A+ wouldn't be moved to  $\Phi$ , herself (Smith 1995). The Advice Model allows that there may be *no* version of A who is moved to  $\Phi$  by rationally considering A's reasons to  $\Phi$ . Since reasons in the sense associated with authoritative demands require that there be such a version of A, the Advice Model fails as a conception of reasons in this sense.

The mistake is inherited by theories of reasons that build on the Advice Model without incorporating the Constraint, such as Smith's constitutivism (2015). Smith's constitutivism seeks to ground reasons in the nature of ideal agency, where the advisor, A+, is an ideal agent. Whatever else this ideal involves, if it *doesn't* involve responsiveness to reasons by doing what those reasons demand, we should reject the account as an account of reasons as capable of imposing authoritative demands. Other constitutivists may fare better in this regard: Velleman (2009) and Korsgaard (1996, 2009) both seem to accept versions of the Constraint, at least implicitly. Whether their brands of constitutivism are otherwise adequate is of course a further question. (On constitutivism, see Michael Smith's chapter "Constitutivism.")

There is surely some need for a conception of reasons as figuring in well-informed advisors' advice. One important open question in this area is the extent to which reasons, as figuring in advice, *must* come apart from reasons as imposing authoritative demands: can we capture reasons' intuitive role in advice compatibly with the Constraint? Or will only some incompatible conception of reasons do? However we end up answering these questions, the Constraint, while modest, is an important choice point in much theorizing about reasons; and captures a central way in which some important normative phenomena depend on agency.

#### NOTES

1. Rosati assumes a form of cognitivism about normative judgments. On cognitivism, see Matthew S. Bedke's chapter "Cognitivism and Non-Cognitivism;" cf. Matthew Chrisman's chapter "Conceptual Role Accounts of Meaning in Metaethics."
2. This section draws on a more detailed treatment in Paakkunainen (ms).
3. We might doubt that reasons are facts or true propositions, but I'll continue to assume this with much of the literature I'm engaging.
4. Thanks to David Plunkett and Daniel Star for helpful discussion here.

#### ACKNOWLEDGMENTS

I'm grateful to David Sobel, Nate Sharadin, Ben Bradley, Daniel Star, Kim Frost, Tristram McPherson, and David Plunkett for extremely helpful discussion and feedback, and to the Murphy Institute's Center for Ethics and Public Affairs at Tulane for supporting part of the work on this chapter.

## RELATED TOPICS

Chapter 1, “Non-Naturalistic Realism in Metaethics;” Chapter 2, “Naturalistic Realism in Metaethics;” Chapter 20, “Reasons Internalism;” Chapter 23, “Constitutivism;” Chapter 24, “Constructivism.”

## REFERENCES

- Dancy, J. 2004. *Ethics Without Principles*. Oxford: Oxford University Press.
- Dancy, J. 2006. Nonnaturalism. In Copp, D. (ed.) *Oxford Handbook of Ethical Theory*, 122–145. Oxford: Oxford University Press.
- Enoch, D. 2006. Agency, Shmagency: Why Normativity Won't Come from What Is Constitutive of Action. *The Philosophical Review* 115: 169–198.
- Enoch, D. 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- FitzPatrick, W. 2008. Robust Ethical Realism, Non-Naturalism and Normativity. In Shafer-Landau, R. (ed.) *Oxford Studies in Metaethics*, Vol. 3: 159–205. Oxford: Oxford University Press.
- Johnson, R. 2003. Internal Reasons: Reply to Brady, Van Roojen and Gert. *The Philosophical Quarterly* 53: 573–580.
- Korsgaard, C. M. 1986. Skepticism about Practical Reason. *The Journal of Philosophy* 83: 5–25.
- Korsgaard, C. M. 1996. *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Korsgaard, C. M. 2009. *Self-Constitution: Agency, Identity, and Integrity*. Oxford: Oxford University Press.
- Lord, E. 2015. Acting for the Right Reasons, Abilities, and Obligation. In Shafer-Landau, R. (ed.) *Oxford Studies in Metaethics*, Vol. 10. Oxford: Oxford University Press.
- McDowell, J. 1995. Might There Be External Reasons? In Altham, J.E.J. and Harrison, R. (eds.) *World, Mind, and Ethics*. Cambridge: Cambridge University Press.
- McPherson, T. 2011. Against Quietist Normative Realism. *Philosophical Studies* 154: 223–240.
- Markovits, J. 2014. *Moral Reason*. Oxford: Oxford University Press.
- Paakkunainen, H. Forthcoming. Internalism and Externalism about Reasons. In Star, D. (ed.) *Oxford Handbook of Reasons and Normativity*. Oxford: Oxford University Press.
- Paakkunainen, H. (ms) Can There Be Government House Reasons for Action?
- Parfit, D. 2011. *On What Matters*. Oxford: Oxford University Press.
- Railton, P. 1984. Alienation, Consequentialism, and the Demands of Morality. *Philosophy and Public Affairs* 13: 134–171.
- Rosati, C. 2003. Agency and the Open Question Argument. *Ethics* 113: 490–527.
- Rosati, C. 2016. Agents and “Shmagents”: An Essay on Agency and Normativity. In Shafer-Landau, R. (ed.) *Oxford Studies in Metaethics*, Vol. 11. Oxford: Oxford University Press.
- Schroeder, M. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- Setiya, K. 2014. What Is a Reason to Act? *Philosophical Studies* 167: 221–235.
- Smith, M. 1995. Internal Reasons. *Philosophy and Phenomenological Research* 55: 109–131.
- Smith, M. 2009. Reasons With Rationalism After All. *Analysis* 69: 521–530.
- Smith, M. 2015. The Magic of Constitutivism. *American Philosophical Quarterly* 52: 187–200.
- Star, D. 2015. *Knowing Better: Virtue, Deliberation, and Normative Ethics*. Oxford: Oxford University Press.
- Velleman, J. D. 2009. *How We Get Along*. Cambridge: Cambridge University Press.
- Way, J. 2015. Reasons as Premises of Good Reasoning. *Pacific Philosophical Quarterly* 97(2): 1–20.
- Williams, B. 1981. Internal and External Reasons. In Williams, *Moral Luck*. Cambridge: Cambridge University Press.