

Original citation:

Peter, Fabienne. (2017) The good, the bad, and the uncertain : intentional action under normative uncertainty. Ethical Theory and Moral Practice .

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/94676>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work of researchers of the University of Warwick available open access under the following conditions.

This article is made available under the Creative Commons Attribution 4.0 International license (CC BY 4.0) and may be reused according to the conditions of the license. For more details see: <http://creativecommons.org/licenses/by/4.0/>

A note on versions:

The version presented in WRAP is the published version, or, version of record, and may be cited as it appears here.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

The Good, the Bad, and the Uncertain: Intentional Action under Normative Uncertainty

Fabienne Peter¹

Accepted: 29 March 2017

© The Author(s) 2017. This article is an open access publication

Abstract My focus in this paper is on a type of bad actions, namely actions that appear to be done for reasons that are not good reasons. I take such bad actions to be ubiquitous. But their ubiquity gives rise to a puzzle, especially if we assume that intentional actions are performed for what one believes or takes to be good reasons. The puzzle I aim to solve in this paper is: why do we seem to be getting it wrong so much of the time? I will argue that we can explain the ubiquity of bad action in light of the practical uncertainties that we face. My claim is not just that the more uncertainty we face as agents, the more likely we are to make a mistake about what counts as a good reason for action, although that is certainly one possible effect of practical uncertainty. My main claim is, rather, that practical uncertainty can have an impact on what counts as acting for a good reason for the agent.

Keywords Uncertainty · Guise of the good · Normative reasons · Reasons for action

1 Introduction

We often act badly, it seems. We are driven by impulses we do not understand and which we do not control well. We also often misjudge the situation we find ourselves in and as a result we pursue our goals in ways that are self-undermining or we pursue the wrong goals altogether. Consider A. A feels he has been slighted and judges that the right way to respond is by asserting himself. In order to assert himself, he emphasizes his achievements. As it happens, his friends feel he is too dominant and respond by ignoring him. Or consider B. B notes how a difficult decision is starting to divide her co-workers. She judges that the best response is to withdraw until the situation has blown over. As it happens, she could have nipped the conflict in the bud. Without an intervention from B, the divisions deepen.

✉ Fabienne Peter
f.peter@warwick.ac.uk

¹ Department of Philosophy, University of Warwick, Coventry CV4 7AL, UK

Some of our behavior is not fully intentional, of course. We may not notice that we feel slighted and or that we are afraid to take a stand. If this is the case, our responses are not transparent to us and, as a result, not fully intentional. This type of behavior is deficient qua intentional action. But the examples above need not be of this type. A and B may judge the situation to be of a certain kind and respond accordingly. Their behavior is thus intentional. Although not deficient qua actions, their actions are nevertheless instances of what I will call “bad actions”. They are bad actions because they appear to be done for reasons that are not good reasons. My focus in this paper is on bad actions of this type.

I take bad actions to be ubiquitous. But their ubiquity gives rise to a puzzle.

The puzzle arises if the following thesis about intentional actions is true: intentional actions are performed for what one believes or takes to be good reasons. Call this the Guise of the Good thesis.¹ Of course, the Guise of the Good thesis is based on the agent’s perspective on what to do, not on good reasons per se. As such, even if it is granted that the Guise of the Good thesis is true, it is not hard to explain the occasional occurrence of bad actions: bad action will result if the agent is mistaken about what good reasons for action are. What is harder to explain, however, if the Guise of the Good thesis is true, is why we seem to be getting it wrong so much of the time. This is the puzzle I intend to solve in this paper.

If the Guise of the Good thesis were false, it would be easy to explain bad actions. For example, suppose that all that is required of intentional action is that the action is performed for *a* reason, not for what one believes or takes to be a *good* reason. It would then not be surprising if many of our actions turned out to be done for reasons that are not good reasons.

Assuming that the Guise of the Good thesis is true, my proposal is that we can explain the ubiquity of bad action in light of the practical uncertainties that we face. I will argue that what we appropriately believe or take to be good reasons depends on how much uncertainty we face when we act. By that I do not just mean that the more uncertainty we face, the more likely we are to make a mistake about what counts as a good reason for action, although that is certainly one possible effect of practical uncertainty. My claim is, rather, that uncertainty can affect what counts as acting for a good reason.

I will argue for this claim on the basis of a two standpoints model. The two standpoints are those of the agent, on the one hand, and an assessor, on the other. The two standpoints may be inhabited by the same person, perhaps at different moments in time, or they may be inhabited by different persons, or by a person and an institution (e.g. a court). What distinguishes the two standpoints is how much uncertainty each faces. The case I will focus on is when the agent faces more uncertainty than the assessor. Based on this two standpoints model, I will argue that what counts as good reasons may differ across the two standpoints and practical uncertainty explains how the two can come apart. Practical uncertainty has this effect, not just because it leads to mistakes about what to do, but because the agent must settle what to take as a good reason for action in this situation. This “settling”, as I will call it, is a requirement of practical rationality. At the same time, it may lead the agent to act on the basis of reasons that, from a standpoint of less uncertainty, do not appear to be good reasons. A bad action, then, is

¹ The Guise of the Good thesis is controversial, but I lack the space to defend it in this paper. In addition, I will not defend my particular take on the Guise of the Good thesis. The Guise of the Good thesis originated as a thesis about how desires aim for the good and even today, many relate the thesis to desires rather than intentional actions. I will follow those – e.g. Raz (2010) and Gregory (2013) – who interpret the thesis in terms of a link between intentional action and normative practical reasons. For good discussions of different interpretations of the Guise of the Good thesis and its plausibility, or lack thereof, see Velleman (1992), Setiya (2007), Tenenbaum (2007), Schroeder (2008), Raz (2010), and Gregory (2013).

an intentional action performed for what the agent takes to be a good reason that, as recognized from the standpoint of (less un-) certainty of the assessor, does not, however, count as a good reason. This is my proposed solution to the puzzle posed by the ubiquity of bad actions.

My paper is organized as follows. I will start with an explication of the Guise of the Good thesis (Section 2). In Section 3, I will discuss different versions of the thesis, depending on how normative practical reasons are understood (Section 3). I then move on to discuss practical uncertainty and its implications for the Guise of the Good thesis (Sections 4 and 5). In the final section of the paper, I will show how the presence of practical uncertainty can explain bad action.

2 The Guise of the Good

The Guise of the Good thesis, as I understand it here, is a thesis about intentional actions. It claims that intentional actions are performed for what one believes or takes to be good reasons.² The relevant reasons are normative reasons, so the claim is that, when acting intentionally, we are acting for reasons that we believe or take to be reasons that justify our actions. The Guise of the Good thesis thus establishes a link between motivational reasons and normative reasons, or between intentional actions and the grounds of practical justification.

The reasons we act from need not be truly good reasons, of course. For example, we might be mistaken and a consideration that we believe or take to be a good reason for action may not have this normative status. At the same time, the Guise of the Good thesis should be a substantive thesis and not be merely trivially true (Setiya 2007, p. 62; Tenenbaum 2013). The thesis would be merely trivially true if any possible consideration that might motivate an action would count as a consideration that may be believed or taken to be a good reason. For example, if all action-relevant desires counted as good reasons for action, there would be no distinction between what is desirable – what is supported by good reasons – and what is desired. As a substantive thesis, the Guise of the Good thesis must thus rest on an account of what can count as good reasons for action – of what can be believed or taken to be good reasons.

What are good reasons for action? This question divides meta-ethicists and philosophers of action and much of my paper will turn on developing an answer to it. But before I can get to this question in the next section, it will be helpful to briefly address an objection to the Guise of the Good thesis first – not in order to defend the Guise of the Good thesis, but in order to highlight an important feature of the thesis. The objection is that we can account for intentional action without involving normative practical reasons (Setiya 2007). All that is required as a marker for intentional action is that there is an explanation for the action. The objection accepts that intentional actions are actions performed for reasons. But it claims that it is not necessary that the relevant reasons are believed or taken to be good reasons. For an action to be an intentional action proper, it is sufficient that the relevant reasons explain the action.

One advantage of this explanatory account is that it appears to have no problem with bad actions.³ It distinguishes between reasons that explain how we act and reasons that justify action and denies that there is a necessary link between the two. If this account were correct, it would thus not be surprising if many of our actions are performed for reasons that turn out to

² I will comment below on the distinction between “believing” and “taking.”

³ Just to be clear, this is not Setiya’s argument.

be not good reasons. On the explanatory account, the action is fully intentional insofar as we know what we are doing: it is sufficient that we can cite a reason that explains what we are doing. There is no need for a connection between our intentional actions and a reason that would justify the action.

I do not think the explanatory account is sufficient, however, on grounds similar to those of Joseph Raz' defense of the Guise of the Good thesis. The problem with the explanatory account is that intentional action is not just action for any reason; it is action for a particular reason (or a particular set of reasons). As Joseph Raz (2010, p. 127) has argued in defense of the normative account, intentional action involves selecting a reason for action among the many possibilities and this selection process must involve distinguishing between good and bad reasons for action. We have not fully accounted for the intentionality of an action if we only cite a reason that explains the action. To account for intentional action, it is thus necessary to say something about why the agent selected the reason(s) they did and this will involve the agent's perspective on what are good reasons for action in the circumstances.

3 Good Reasons

The key question that needs to be addressed now is how this selection process is supposed to be understood and the answer will depend on a conception of normative practical reasons. There are two main alternative versions of the normative account of the Guise of the Good thesis: a realist version and a constitutivist version. According to the former, normative practical reasons are or are given by facts and such facts are reasons independently of any mental attitudes an agent happens to have – their actual desires or beliefs, for example – or of any explanation in terms of such attitudes.⁴ The realist conception of reasons supports an interpretation of the Guise of the Good thesis that foregrounds the agent's ability to recognize fact-given reasons and/or the actions that such reasons favor. An action is intentional, on this version of the thesis, if the agent acts on the basis of a belief involving fact-given normative reasons. The idea is that fact-given reasons orient the practical agent via their normative beliefs and ground his or her judgment about what to do. Raz, who has been influential in developing the realist conception of reasons, describes this version of the Guise of the Good thesis as follows:

“What marks intentional actions is that they are done because of what their agents believe the action is (including what it may bring about). That means that what the agents believe about the action leads them to do it, and guides their doing of it, all the way to its conclusion, and that suggests that they approve of the action, given what they believe about it” (Raz 2010, p. 116; my emphasis).

If all goes well, the agent's beliefs will track the normative practical reasons that apply in a particular situation. Perhaps the agent even knows what reasons apply and this knowledge is the basis of her intentional actions. But the Guise of the Good thesis does not imply knowledge, of course. The normative beliefs that guide an agent may well be false. In this

⁴ See Raz (1999), Dancy (2000), Parfit (2011), and Scanlon (2014), among others. The details of their accounts differ, of course. To keep terminology simple, I will call such reasons fact-given reasons.

case, the content of the beliefs will not be given by the reasons that, in fact, apply in the particular situation. Instead, there will be a gap between the action that normative practical reasons truly favor and the action that the agent believes is so favored.

To clarify, let me note that the realist version of the Guise of the Good thesis can be understood in two different ways, depending on how beliefs involving normative reasons are understood. I lean towards a *de dicto* interpretation of beliefs involving normative reasons. On this interpretation, an agent acts intentionally if they act on the basis of a consideration that they believe is a normative reason for action that applies in the relevant circumstances. On a *de re* interpretation of the relevant beliefs, by contrast, an agent acts intentionally if they act on the basis of a belief that, if it were true, would turn out to be or give rise to a normative reason. *De re* interpretations are sometimes favored because they do not require that the agent has a concept of a normative reason when acting for a reason. I lean towards a *de dicto* interpretation, however, on grounds similar to those Kurt Sylvan (2015) cites against *de re* interpretations: *de re* interpretations tend to neglect the role of an agent's rationality in their accounts of what it means to act for a reason.

In the alternative, constitutivist, account, the agent's rationality in acting for reasons plays a central role. While there are many versions of constitutivism, I will focus here on Christine Korsgaard's Kantian version (Korsgaard 2008, 2009a, b).⁵

Korsgaard would reject even the *de dicto* interpretation of the realist version of the Guise of the Good thesis as not involving the agent's rationality to a sufficient degree or in the right way. On her account, normative practical reasons are not fact-given and not something that rational agents can discover and use to guide their actions. Instead, normative practical reasons constitutively depend on practical rationality, where the capacity for rationality is understood in an active sense, rather than a passive sense. Here is how Korsgaard describes the "activity of reason" (2009b, p. 30):

"The faculty of reason is not identified merely as the ability to recognize and respond to reasons. The faculty of reason is identified rather as the active dimension of the mind, and rational principles are then identified as those that describe or constitute rational activity. They are constitutive principles of rational activity. When those principles are applied by the person who is trying to work out what to believe or to do, they pick out the substantive considerations that we then regard as reasons."

On this active understanding, the capacity for practical rationality is the capacity to entertain and endorse certain considerations as reasons for action and to act on the basis of reasons thus determined.

While Korsgaard rejects the notion that the normative guidance that is at the heart of intentional action stems from fact-given reasons, she does recognize that rational activity is governed by certain standards or norms. According to her, these standards are not themselves reasons, however, they are, instead, standards that constitute rational agency. Specifically, Korsgaard recognizes two main standards that, together, constitute rational agency. Her understanding of those standards is derived from her interpretation of Kant's ethics (Korsgaard 2009a, p. 45ff). The first is the Categorical Imperative (in its different formulations). The other is the Hypothetical Imperative, a standard of instrumental rationality. Other

⁵ For other versions of constitutivism, see Street (2008) and Katsafanas's helpful overview (Forthcoming).

constitutivist views imply different standards of rationality. But they all construe those standards as internal to rational agency – in contrast to the realist view.

Returning to the main topic of interest here, the constitutivist version of the Guise of the Good thesis says that intentional action is action performed for what the agent takes to be a good reason for action, not in the sense of what the agent believes about fact-given reasons, but in the sense of being a consideration that the agent endorses as a good reason. If all goes well, the intentional action is a fully rational action – it fully satisfies the standards of rationality – and the consideration that the agent endorses as a reason is a normative practical reason proper. Like the realist version, this version of the Guise of the Good thesis does not include a success condition either. An agent's deliberation may not satisfy these demanding standards and, as a result, even though the agent acts on the basis of a consideration that they endorse as a good reason for action, the consideration may not be a normative practical reason proper.⁶

Prima facie the two versions of the Guise of the Good thesis appear to be incompatible. On one version, intentional action is driven by beliefs involving fact-given normative reasons and, on the other, it is driven by the active endorsement of certain considerations as reasons. I believe that each captures an important aspect of intentional agency. My aim is thus to argue that we should seek to combine the insights each offers in a hybrid version of the Guise of the Good thesis.

How we respond to practical uncertainty plays a key role in my argument. But before I get to this topic, let me first spend a bit more time explaining what the two aspects of intentional agency are. My starting point is the question of how we should understand the role of practical reasoning in intentional action. As I have characterized it, the realist version of the Guise of the Good thesis highlights what I call the orientation function of practical reasoning. On this account, the aim of practical reasoning is to adequately respond to normative truths – truths about fact-given reasons. The key aspect of intentional action that this account highlights is that in determining how to act, we respond to considerations that appear not to be up to us. The important distinction is between what is up to us – acting on one desire rather than another, or on one consideration rather than another – and what is not up to us. What is not up to us, on the realist account, does not just include non-normative features of this world – e.g. the fact that I cannot fly. It also includes normative features that function in the same way, for example the fact that, in many circumstances, there are no reasons to harm others. The aim of the orientation function of practical reasoning is to ensure that in determining how to act, we adequately respond to what is not up to us.

The constitutivist account, by contrast, highlights what I call the settling function of practical reasoning. Rather than focusing on what is not up to us, it takes the opposite perspective and focuses on what, in practical reasoning, is up to us.⁷ The emphasis is on how, in practical reasoning, we can take a step back from our motives to act in particular ways and to reflect on possible alternative grounds for action. On this account, practical reasoning is taken to be a power that can settle what counts as a reason for action and the aim of practical reasoning is to settle what counts as a good reason.

⁶ Lavin (2004) has an excellent discussion of the ways in which Kantian constitutivism can – and cannot – allow for error. Korsgaard discusses defective action – actions that fail to meet the standards of practical rationality – in chapter 8 of *Self-Constitution*. In her view, all bad actions are defective actions – they are “bad *as* actions” (2009a, p. 161, her emphasis). On the view that I will advocate in this paper, there are bad actions that are good as actions but bad in the sense that they are performed for reasons that do not, from a different standpoint, count as good reasons.

⁷ See also Alonso (2017) on this way of understanding intentional action.

4 Practical Uncertainty

How does practical uncertainty bear on intentional action? This question is the topic of this section. Uncertainty about how to act is common and I am assuming that it would be wrong to think that intentional action, action for what we believe or take to be good reasons, is suspended in circumstances of uncertainty. Acting under uncertainty should thus not be bracketed from the scope of the Guise of the Good thesis as a marginal case. If the Guise of the Good thesis is true, it must also be true for actions under practical uncertainty.

When we are uncertain about how to act, the link between practical reasoning and intentional action becomes, if anything, even more important. Uncertainty makes us pause and reflect on how best to act. There is a range of uncertainties that affect our practical reasoning. We may, for example, be uncertain about non-normative properties of the circumstances we are acting in. For example, a person may be uncertain about whether her watch shows the correct time and thus whether she will catch the train. Or she may be uncertain about what would happen if she missed the train – will the meeting take place without her or will it be rescheduled? This type of uncertainty is common, but primarily theoretical. My main focus is on practical uncertainty: uncertainty about how to act.

Practical uncertainty may concern the normative implications of different actions, including of their consequences. I call this normative outcome uncertainty. For example, if I give to a particular charity, will the charity be able to act on its commitments or will it find its efforts thwarted, such that my giving will have adverse effects? Or: will my apology be received as such or will this create further dissonance between my friend and me? Practical uncertainty may also concern the correct description of the circumstances under which one acts. I call this normative circumstance uncertainty. It is generally recognized that intentional action is “action under a description” (Anscombe 1963). But which description? The question is not just whether one has accurately described all the relevant non-normative facts. The question that normative circumstance uncertainty highlights is, rather, whether one has included all the normatively relevant features of the practical situation one finds oneself in, and whether one has thus correctly described the alternative possible actions. For example, you may be uncertain as to whether you are correct to think that the normatively salient feature of the situation you are in is that your friend needs help or whether the correct description is that you are feeling the need to control your friend.⁸

Practical uncertainty, thus understood, creates a particular problem for the realist version of the Guise of the Good thesis because that version links intentional action to normative beliefs. Recall that, on this version of the thesis, actions are intentional if they are performed for what are believed to be good reasons. Both forms of practical uncertainty imply, however, that we cannot just rely on normative beliefs in intentional action.

With normative outcome uncertainty, the challenge that practical uncertainty poses is how to compare alternative courses of action given uncertainty about their normative implications. For example, if you are uncertain whether a charity will be able to act on its commitments, the question you will be asking yourself is whether that implies that you have good reason to go ahead all the same or whether this changes your reason to give such that you now have good reason to not give at all or to give to a different charity, one where a positive outcome of your giving is more certain, perhaps.

⁸ See also Hampshire (1949) for a good account of what I call normative circumstance uncertainty here.

I take it that in your practical reasoning in this case, you balance the reasons you believe you have for particular courses of action against the degree of outcome uncertainty that attaches to these alternative courses of action to determine what you have good reason to do. For example, if it is likely that the charity can uphold its commitments, your original reason to give prevails. If it is very unlikely that they can, however, you might then judge that the reason to give to them fades in importance relative to other reasons for action you might consider. In responding to normative outcome uncertainty, you take a step back from your beliefs and weigh reasons that favor alternative courses of action with the likelihood of their implications. This type of uncertainty will thus affect which action you will take to be supported by good reasons. It also makes your practical reasoning more complicated than in cases of certainty, because your judgment of what counts as a good reason for action is now not just dependent on your beliefs about normative reasons for action, but also on how you take your reasons for action to be affected by the uncertainty. Because normative outcome uncertainty makes your practical reasoning more complicated, it may well increase the likelihood of mistakes about the reasons that apply in those circumstances and thus widen possible gaps between what you believe or take to be good reasons for action and the normative practical reasons that truly apply in those circumstances. But, more importantly, normative outcome will also impact on the link between intentions and beliefs involving fact-given normative practical reasons. In your practical reasoning, you may still be guided by your beliefs in relation to fact-given reasons. But, as we saw, you cannot just rely on the guidance that your beliefs about what these reasons are provide. You will also have to make a judgment about how uncertainty interacts with those reasons. What I earlier called the orientation function of practical reasons need thus not be disabled by normative outcome uncertainty, but you cannot solely rely on it when deliberating about how to act.

With normative circumstance uncertainty, the problem that uncertainty poses runs even deeper. Before I explain what I have in mind here, it will be helpful to consider an excellent example for what I mean by normative circumstance uncertainty: Iris Murdoch's description of M's moral challenge and her response (Murdoch 1971, p. 16f):

A mother, whom I shall call M, feels hostility to her daughter-in-law, whom I shall call D. M finds D quite a good-hearted girl, but while not exactly common yet certainly unpolished and lacking in dignity and refinement. D is inclined to be pert and familiar, insufficiently ceremonious, brusque, sometimes positively rude, always tiresomely juvenile.... Thus much for M's first thoughts about D. Time passes, and it could be that M settles down with a hardened sense of grievance and a fixed picture of D, imprisoned (if I may use a question-begging word) by the cliché: my poor son has married a silly vulgar girl. However, the M of the example is an intelligent and well-intentioned person, capable of self-criticism, capable of giving careful and just *attention* to an object which confronts her. M tells herself: 'I am old-fashioned and conventional. I may be prejudiced and narrow-minded. I may be snobbish. I am certainly jealous. Let me look again.' Here, I assume that M observes D or at least reflects deliberately about D, until gradually her vision of D alters. If we take D to be now absent or dead this can make it clear that the change is not in D's behavior but in M's mind. D is discovered not to be vulgar but refreshingly simple, not undignified but spontaneous, not noisy but gay, not tiresomely juvenile but delightfully youthful, and so on.

Murdoch's story suggests that M gradually discovers the correct description of her daughter-in-law. M starts off with a description of the relevant circumstances that is biased under the influence of her own feelings towards her son, but then distances herself from that description and discovers the correct description of the practical situation she finds herself in. I am less interested in that happy outcome, however, than in the moment where M takes a step back and asks herself whether the way in which she first described the situation to herself is indeed the correct one. This moment – when she questions her initial description but before she discovers the correct description – is a moment of normative circumstance uncertainty.

Murdoch's own portrayal of M's moral challenge focuses on M's perception of the situation, not on M's practical reasoning. But it is easy to see how uncertainty about the correct description of the circumstances translates to uncertainty about the normative practical reasons that apply. If M's first description of the situation were correct, actions that support D's further personal development and further education might appear as supported by good reasons. If the second description were correct, actions that aim at sharing D's joyfulness might appear as supported by good reasons.

The main problem that normative circumstance uncertainty creates is not just that we are uncertain as to which of a certain number of known possible alternatives is the correct description of the normative circumstances. If we know what the possible alternatives are and are merely unsure about which will obtain, we have a case of outcome uncertainty. The distinct problem that normative circumstance uncertainty creates is that we are uncertain about which possible alternatives to include in our deliberation about how to act. In my interpretation of M's case, the problem is her uncertainty about what alternative ways of seeing her daughter in law are correct. M might ask herself the following. Is her original perception correct and should she thus try to educate D? Or is D more accurately described as joyful, etc., and should M thus try to learn from her? Or, now going beyond Murdoch's own portrayal of M's challenge, is there yet another way of seeing D, one that M has not yet conceptualized, which would be truer to D and which would have yet other implications for how best to respond?

However M answers these questions, it seems that she cannot settle the question of how to act on the basis of her normative beliefs. Whatever guidance her normative beliefs offer, that guidance is only relative to a description of the circumstances. Their guidance is thus not sufficient to help M settle on an intention when she is uncertain about how to describe the normative circumstances she finds herself in. Generally speaking, if we are uncertain about what the correct description is of the normative circumstances we find ourselves in, then what I earlier called the orientation function of practical reasoning will be inhibited or even disabled altogether.

To a greater or lesser extent, both forms of practical uncertainty thus imply that we cannot simply rely on our beliefs involving fact-given reasons for orientation in forming an intention.⁹ This raises a question about how we should revise the realist version of the Guise of the Good thesis to accommodate intentional action under practical uncertainty. As I have argued above,

⁹ Both forms of practical uncertainty may also be described as cases of epistemic underdetermination. When facing normative outcome uncertainty, we do not know enough about the normative implications of our actions. When facing normative circumstance uncertainty, we lack, in addition, knowledge about the correct description of the circumstances we find ourselves in. Epistemic underdetermination is distinct from normative underdetermination. It is possible that normative practical reasons “run out”, to use a phrase by Ruth Chang (2013). As Chang has argued, rightly in my view, practical reasoning cannot be limited to considering fact-given reasons as such reasons might underdetermine the decision problem. But the problem I am addressing here is different: it is not that fact-given reasons have run out; we just do not know what action they favor.

if the Guise of the Good thesis is true, it must also be true for intentional action under practical uncertainty. So there must be another way of understanding the link between action and normative practical reasons than the one that the realist version of the Guise of the Good thesis offers.

5 Intentional Action under Practical Uncertainty

In this section, I will argue that the constitutivist version of the Guise of the Good thesis offers us a clue for how to think about intentional action under practical uncertainty – that is, under circumstances where reliance on the orientation function of practical reasoning is either limited or blocked altogether. Recall that the constitutivist version claims that intentional actions are performed for what are taken to be good reasons for action. I interpreted this in terms of considerations that are endorsed as normative reasons for action and called the power to confer normative status to such considerations the settling function of practical reasoning.

Of course, the constitutivist version of the Guise of the Good thesis is normally presented as an alternative to the realist version, not as a complement. As I have explained earlier, the underlying claim is that normative practical reasons are not facts or fact-given, and hence not the object of normative beliefs; they are the result of appropriately conducted practical reasoning. While I cannot argue for this here, I agree with the defenders of the realist version of the Guise of the Good thesis that normative beliefs – and what I have called the orientation function of practical reasoning – matter. The problem with the claim about normative practical reasons that is underlying the constitutivist version of the thesis is that it invites too much arbitrariness: without anchoring in normative facts, even well conducted practical reasoning may fail to generate correct conclusions about how we should act (Shafer-Landau 2003).

I thus reject the ambitious version of the constitutivist project, which aims to reconstruct practical normativity entirely from within practical reasoning. But I think that a more modest interpretation of the core idea of this project can help us overcome the deficiencies of the realist version in dealing with intentional action under practical uncertainty. Constitutivists have rightly highlighted the question of how we can rationally form an intention when we cannot turn to normative facts for guidance. With their work on what I have called the settling function of practical reasoning, they have captured an important function of practical reasoning and an important aspect of practical normativity. Their insights can be preserved in a more restricted interpretation of the role of constitutivism, one that reconciles it with the contribution of the realist project. The important contribution of the constitutivist project, as I see it, is to help us understand how to think about rational guidance in cases where the orientation function of practical reasoning is inhibited or even disabled.

While the restrictive interpretation rejects many of the claims that constitutivists make about the nature of practical normativity in general, it upholds the claim that the settling function of practical reasoning can offer rational guidance, at least in some circumstances. Of course, not all exercises of the settling function will pick out normative practical reasons proper. Just like relying on the orientation function may lead us astray when our beliefs are false, so does relying on the settling function may lead us astray when our practical reasoning is faulty. Even the restrictive interpretation maintains, however, that taking a consideration as a good reason for action, will, in the right conditions, truly pick out a normative practical reason, albeit a reason of a different kind than fact-given reasons. It rejects the more ambitious claim of

constitutivists, however, that appropriately taking a consideration as a normative practical reason is the only source of practical normativity.

How are we supposed to imagine the interplay between the orientation function and the settling function, and hence a hybrid version of the Guise of the Good thesis that accommodates elements from both the realist version and the constitutivist version? Before I can answer that question, let me comment on an important difference between practical reasoning and theoretical reasoning. If we face uncertainty in the context of theoretical reasoning, it is possible to leave many of our beliefs unsettled. For example, if skeptical arguments have succeeded in making me uncertain about whether or not I have hands, I can, in a purely theoretical context, suspend belief on the matter. Practical reasoning cannot end in the suspension of intention. It is a distinctive feature of the practical context that we must always act. So we must act even when we face uncertainty about what to believe about the situation we find ourselves in.

This feature of the practical context explains the significance of the settling function. When normative practical uncertainty is severe, that is, when the orientation function of practical reasoning is too inhibited or even disabled to offer rational guidance, another function of practical reasoning must be able to fill the gap or the Guise of the Good thesis cannot be upheld. When, in deliberating about how to act, we realize that relying on the orientation function does not deliver an intention, we need to settle the matter in a different way. If we do not, we will fail to act intentionally and, instead, act on the basis of a consideration that we recognize is not a good reason for action. That would be a bad way of settling how to act, one that undermines the Guise of the Good thesis.

A better way of settling how to act in circumstances where the orientation function of practical reasoning leaves matters unresolved is to ask oneself, what consideration one might be able to endorse as a good reason for action in those circumstances. Let me illustrate the thought with the miner example that is sometimes thought to create a problem for the Guise of the Good thesis (quoted from Raz 2010, p. 112f):

“The management proposes to close the colliery. The miners vote on whether to accept the proposal and the redundancy pay which goes with it or to oppose it. You talk to one of the miners: ‘You are voting to stay put’. – ‘Sure’, he says. ‘So you must have some hope [of keeping the mine open]’. – ‘No hope. Just principles’.”

On my interpretation, the example illustrates the limitations of the orientation function of practical reasoning and the need to rely on the settling function. The miner seems to be expressing an awareness that the situation is too uncertain to rely on a belief of the sort that, because staying put will keep the mine open, there is a good reason to stay put. So his decision to stay put is not determined by the orientation function. Instead, it is based on what he takes to be a good reason to act given the uncertainty of the situation. What he takes to be reasoning in the circumstances, in my reading of the case, is a principle such as integrity or solidarity – i.e. a consideration that is consistent with what he has endorsed in related circumstances.

To illustrate the thought further, going back to my earlier example, M, when she is unsure about what the correct way of looking at D is, and hence about the normative reasons for action that apply in the circumstances, might ask herself the following. “Should I give her the benefit of the doubt and opt for being tolerant towards her, given that I am not sure whether I am seeing her in the correct light?” “Instead of focusing on her faults, should I try to find

common ground and focus on interests we both share?” However she answers those questions, if she relies on the settling function of practical reasoning, her reason for acting in a particular way towards D will not simply reflect her normative beliefs, but will reflect which consideration she takes to be endorsable as a reason for action in those circumstances. And note also that even if she were to end her deliberation with “what the hell, trying to help her become a better person can’t be wrong”, this conclusion, although in accord with the normative belief she originally held, is still the result of her taking this to be a good reason for action in the circumstances, not merely a reflection of her normative beliefs. Even then she thus relies on the settling function, not on the orientation function of practical reasoning.

In short, I am proposing a hybrid version of the Guise of the Good thesis, according to which both the orientation function, as identified by the realists, and the settling function, as highlighted by the constitutivists, have a role to play in explaining the process of selecting reasons for action that is at the heart of intentional action. Switching from the orientation function to the settling function is, on my proposal, an adequate response to the epistemic constraints that are faced in practical deliberation. If uncertainty is limited, selecting a reason for action can be guided by beliefs about normative truths. If uncertainty is severe, however, the selection process depends on the guidance of practical reasoning itself and on which considerations an agent can endorse as reasons for action.

6 Explaining Bad Action

My aim in the last two sections of this paper has been to show how practical uncertainty affects intentional action. And my conclusion was that while practical uncertainty need not undermine the normative account of intentional action, it affects how we should interpret the Guise of the Good thesis, specifically, with regard to the conception of normative reasons for action. In this final section of the paper, my goal is to apply the hybrid version of the Guise of the Good thesis that I have developed in this paper to explain the possibility of bad action in a way that is consistent with the thesis, but that does not trivialize bad actions as mere mistakes.

To cut straight to the chase, the hybrid version of the Guise of the Good thesis accommodates bad actions in the following way: bad actions are actions that an agent performs for what he or she takes to be good reasons but that, from a standpoint of (less un-) certainty, are revealed to be not good reasons. This account hinges on a distinction between two standpoints: the standpoint of the agent and the standpoint of the assessor. The two standpoints may be occupied by one and the same person, perhaps at different points in time, or by different persons. Importantly, the two standpoints differ with regard to practical uncertainty. Two scenarios are possible. In the first, the assessor is more uncertain than the agent. The agent is, for example, better informed about the relevant circumstances than the assessor and can rely on her normative beliefs for guidance in a way that would elude a less informed assessor. In the second scenario, evidence is available to the assessor about the reasons-facts that apply that is not available to the agent. The evidence available to the assessor removes or at least diminishes normative practical uncertainty and reveals that the reasons for which the agent performs the action are in conflict with the fact-given normative reasons that apply in the agent’s situation.

The second scenario is the one I want to focus on. If the evidence available to the assessor is not available to the agent in the moment of the action, bad actions may result, even though the agent is acting under the Guise of the Good. In other words, on the account I offer here, bad actions are the result of the settling function commendably identifying some consideration as a

good reason for action that, from a standpoint of less uncertainty, turns out to conflict with normative truths.

Such bad actions are fully compatible with the Guise of the Good thesis because they are actions that are performed for what the agent takes to be good reasons. It is only from a standpoint not available to the agent in the moment of action that it is recognized that the action is not based on a good reason, at least not in the relevant context. Bad actions are possible because of the gap between the standpoint of the agent and the standpoint of the assessor that opens as a result of normative practical uncertainty. As most of our actions arguably take place in contexts of more or less practical uncertainty, we should expect bad actions to be common. We should expect many of our actions to be based on considerations that the settling function of practical reasoning identifies as good reasons for action, even though these considerations may well be in conflict with the normative facts.¹⁰

Bad actions as I have been characterizing them are not mere mistakes. Mistakes occur if practical reasoning fails to meet the norms that govern it. For example, I may misinterpret the evidence I have and mistakenly identify a consideration as a fact-given reason for action. Or, in relying on the settling function, I might endorse a consideration as a reason for action, mistakenly thinking that it is consistent with a judgment I made in a similar situation.

Bad actions as I have been characterizing them may fully respect the norms that govern the two functions of practical reasoning. While I lack the space here to develop substantive proposals for what the norms that govern practical reasoning are, let me briefly comment on this issue. Each function of practical reasoning is governed by different standards or norms. The orientation function, which connects practical reasoning to normative facts, is governed by the appropriate norm of belief. I have to leave open here what this norm is. The settling function, in the restricted interpretation that I have advocated, is constrained in its application by the norm of belief, but it is not itself governed by that norm. Its function is to settle our intentions in the absence of guidance from fact-given reasons. It is governed by its own norms, norms that determine well-functioning practical reasoning. As we saw above, Korsgaard defends very ambitious norms for what I call the settling function, including the Categorical Imperative. On the restricted view that I favor, the norms will be less demanding requirements of consistency in reasoning, for example.

Whatever exactly the right account of the norms that govern the two functions is, bad actions as I have characterized them are actions that result from practical reasoning that potentially respects both the norms of the orientation function and of the settling function, but that, from a different standpoint, are recognized as based on reasons for action that clash with the fact-given reasons that apply in the circumstances. So bad actions are not mere mistakes.

By the same token, bad actions as I have explained them need not be defective qua actions. An action is defective if it is not fully intentional in the sense that it is not performed for what are believed or taken to be good reasons. One type of action that is defective qua action is an action that is performed for a reason that is believed to be not a good reason, for example if the settling process is aborted too soon, as explained earlier. Generally speaking, defective action results if the switch from the orientation function to the settling function is not made even though the agent is aware that practical uncertainty would warrant it or if the settling function is not exercised appropriately.

¹⁰ If I am right about the ubiquity of practical uncertainty, we should also expect diverging judgments about what counts as a good reason for action to be common.

My account of bad actions does not presuppose such defects, although such defects may occur, of course. I have argued that the settling function of practical reasoning appropriately takes over in circumstances of severe practical uncertainty. If practical reasoning functions well in those circumstances, belief about the applicable reasons-facts is suspended. There is thus no assumption that the action would be performed for what one believes to be not a good reason. The role of the settling function is to identify what can be endorsed as a good reason for action in a situation of unsettled normative beliefs. It functions well if the process of identifying a consideration as a good reason for action fully satisfies the norms that govern it – whatever exactly those norms are. The result is an action that is not defective as an action. What my account shows, however, is that even if the settling process works appropriately in circumstances of practical uncertainty, it might still identify considerations as good reasons for actions that are not, in fact, good reasons.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Alonso F (2017) Intending, settling, and relying. In: Shoemaker D (ed) *Oxford studies in agency and responsibility*, volume 4. Oxford University Press, Oxford
- Anscombe E (1963) *Intention*. Harvard University Press, Cambridge
- Chang R (2013) Grounding practical normativity: going hybrid. *Philos Stud* 164:163–187
- Dancy J (2000) *Practical reality*. Oxford University Press, Oxford
- Gregory A (2013) The guise of reasons. *Am Philos Q* 50(1):63–72
- Hampshire S (1949) Fallacies in moral philosophy. *Mind* 58(232):466–482
- Katsafanas P (Forthcoming) Constitutivism about practical reasons. In: Star D (ed) *The Oxford handbook of reasons and normativity*. Oxford University Press, Oxford
- Korsgaard CM (2008) Acting for a reason. In: Korsgaard CM (ed) *The constitution of agency*. Oxford University Press, Oxford, pp 207–229
- Korsgaard CM (2009a) *Self-constitution*. Oxford University Press, Oxford
- Korsgaard CM (2009b) The activity of reason. *Proceedings and Addresses of the APA* 83(2):23–43
- Lavin K (2004) Practical reason and the possibility of error. *Ethics* 114(3):424–457
- Murdoch I (1971) *The sovereignty of good*. Routledge, London
- Parfit D (2011) *On what matters*. Volumes I and II. Oxford University Press, Oxford
- Raz J (1999) *Engaging reason*. Oxford University Press, Oxford
- Raz J (2010) On the guise of the good. In: Tenenbaum S (ed) *Desire, practical reason, and the good*. Oxford University Press, Oxford, pp 111–137
- Scanlon T (2014) *Being realistic about reasons*. Oxford University Press, Oxford
- Schroeder M (2008) How does the good appear to us? *Soc Theory Pract* 34(1):119–130
- Setiya K (2007) *Reasons without rationalism*. Princeton University Press, Princeton
- Shafer-Landau R (2003) *Moral realism*. Oxford University Press, Oxford
- Street S (2008) Constructivism about reasons. In: Shafer-Landau R (ed) *Oxford Studies in Metaethics*, Vol 3. Oxford University Press, Oxford, pp 208–245
- Sylvan K (2015) What apparent reasons appear to be. *Philos Stud* 172(3):587–606
- Tenenbaum S (2007) *Appearances of the good: an essay on the nature of practical reason*. Cambridge University Press, Cambridge
- Tenenbaum S (2013) Guise of the good. In: LaFollette H (ed) *The international encyclopedia of ethics*. Wiley-Blackwell, Hoboken, pp 2262–2271
- Velleman JD (1992) The guise of the good. *Noûs* 26(1):3–26