

The Value of Practical Usefulness

Rob van Someren Greve • University of Amsterdam

Forthcoming in *Philosophical Studies*

Abstract Some moral theories, such as objective forms of consequentialism, seem to fail to be practically useful: they are of little to no help in trying to decide what to do. Even if we do not think this constitutes a fatal flaw in such theories, we may nonetheless agree that being practically useful does make a moral theory a better theory, or so some have suggested. In this paper, I assess whether the uncontroversial respect in which a moral theory can be claimed to be better if it is practically useful can provide a ground worth taking into account for believing one theory rather than another. I argue that this is not the case. The upshot is that if there is a sound objection to theories such as objective consequentialism that is based on considerations of practical usefulness, the objection requires that it is established that the truth about what we morally ought to do cannot be epistemically inaccessible to us. The value of practical usefulness has no bearing on the issue.

1. Introduction

Many hold that one of the criteria of adequacy for moral theories is that such a theory is ‘action-guiding’, or ‘practically useful’. Moral theorizing does not just have a theoretical but also a practical aim, and thus a moral theory should, in some sense of ‘should’, be of use in making practical decisions, or so it is often suggested. The demand for practical usefulness has been wielded against various types of moral theories,¹ but most commonly against objective forms of consequentialism, which hold (roughly) that a morally right action is an action with consequences that are at least as good as those of any alternative to it. Were we

¹ Virtue ethical approaches have for example been criticized for their alleged failure to be action-guiding; see Hursthouse (1999: chapter 1) for discussion.

to try to adopt such a theory as a guide for practical decision-making, we would often, perhaps even always, fail at finding out which action it requires us to perform, and thus such a theory will be of little to no help in deciding what to do. The main reason that such theories fail to be practically useful is that, with respect to most of the actions that we can perform, we know far too little of their consequences to be able to determine whether their consequences are at least as good as those of the alternatives.

Apart from modifying or altogether abandoning her theory, there are two ways in which a proponent of an objective consequentialist theory (or any other putatively useless moral theory) can respond to the objection that her theory fails to be practically useful: on the one hand, she can attempt to show that, appearances to the contrary notwithstanding, the moral theory she proposes *is* practically useful after all; on the other, she can reply that while it indeed fails to be practically useful, this does not show that the proposed theory is incorrect, for the truth about what we ought to do may turn out to be epistemically inaccessible to us, no matter how scrupulous and conscientious our inquiries into an action's consequences have been.² The second line of reply is – unsurprisingly, perhaps – the more popular one.

Faced with the reply that the failure to be practically useful is unobjectionable since the truth about what we morally ought to do may be inaccessible to us, there are (again) two options for a rejoinder open to those who demand that a moral theory be practically useful: she can either try to show that those who propose moral theories which fail to be practically useful are mistaken in thinking that the truth about what we morally ought to do can be epistemically inaccessible to us,³ or she can instead offer the more conciliatory observation that while it is indeed *possible* that the true theory will fail to be practically useful, all parties to the debate at least should be able to agree that being practically useful is a desirable, or 'good-making' feature that a moral theory might possess. Holly M. Smith for instance considers this conciliatory rejoinder:

² For some recent versions of the first reply, see Hare (2011) and Dorsey (2012). The second reply can be found in Bales (1971) and Feldman (2006), among others.

³ Frances Howard-Snyder's (1997) attempted rejection of objective consequentialism provides a good example.

Another option is to view usability as only one valuable feature (among several) of a moral system. On this view we would grade one moral system as better than another, other things being equal, if the first is more widely usable than the second (1988: 102).

While Smith herself does not explicitly come out in support of this way of understanding the demand for a practically useful moral theory, the rationale underlying the demand that she endorses (discussed below) supports this way of looking at practical usefulness. And in a recent article building on Smith's discussion, Pekka Väyrynen develops this conciliatory rejoinder in more detail. After noting that 'it is not *incoherent* to ignore the practical aim in evaluating moral theories', he suggests that conceding this point does not by itself commit us to holding that practical usefulness is not a desirable feature of a theory:

An explanatory ethical theory may fail to give adequate guidance for acting rightly, and yet give a correct account of what is right. ... We might, however, think that even theories that do not *aim* to provide adequate moral guidance are nonetheless better to the extent that they also *do* provide adequate guidance, instead of thinking that doing so is no merit at all (2006: 292, emphases in original).

The aim of the present paper is to assess this conciliatory rejoinder. More specifically, the aim is to examine what weight we should give to considerations of practical usefulness in determining what moral theory we overall have most reason to believe, when we construe these considerations in the way suggested.⁴ I start with sketching the account Smith and Väyrynen offer of why a practically useful moral theory is a better theory, and explaining that on this account, the reason we have for believing a practically useful moral theory is a pragmatic, not an evidential reason, and thus a 'reason of the wrong kind'. This is one

⁴ A different way to characterize the purpose of this paper, in terms some might find more congenial given the existing literature on the subject, is that it seeks to assess what weight we should assign, respectively, to the theoretical and the practical aims of moral theorizing, provided we understand the practical aim in the way sketched in the text. In effect, the claim that I argue for is that the practical aim of moral theorizing should be taken as subsidiary to its theoretical aim, not by definitional fiat, but in virtue of certain facts about the nature and content of the reasons there are to care about whether a moral theory is practically useful, facts which have important implications for the weight of those reasons.

ground for thinking that taking a conciliatory approach has the implication that the weight that considerations of usefulness can put onto the scale, with respect to what theory we overall have most reason to believe, is limited at best (section 2).

After that, I explain why, even if pragmatic reasons can sometimes carry the day with respect to what we overall have most reason to believe, considerations of practical usefulness are unlikely to do so, for two reasons. On the one hand, the value of being practically useful is arguably conditional on the truth of a moral theory; if that is so, then pragmatic reasons for believing a moral theory provided by considerations of practical usefulness will have to point in the same direction as the evidential reasons, if the former point anywhere at all (section 3). On the other, even if the value of being practically useful is not conditional on the truth of a moral theory, the most weighty pragmatic reasons for believing a given moral theory will stem from its being instrumental towards doing what is morally right (or bringing about the best results), and so it is unlikely that we will have most reason to believe a false-but-useful theory over the true moral theory (section 4).

2. Value and reasons for belief

In what respect can a practically useful moral theory be claimed to be a *better* theory than a rival moral theory, which is not practically useful? Smith suggests that the ‘importance of [practical] usability ... can be explained by reference to the concept of *autonomy*’ (1988: 105, emphasis in original).⁵ She proposes to understand autonomy in the following way: ‘a person acts autonomously insofar as his decision to act is governed by the kinds of considerations that *he* deems most important’ (*ibid.*). In other words, we act autonomously in the relevant sense whenever we decide to act on the basis of what we ourselves take to be the considerations that determine what we morally ought to do in the situation. Being autonomous in this sense is something that matters to us, Smith maintains, and rightly so, as ‘the morally committed person ... [has] an intrinsic interest in carrying out the precepts of his morality’ (*ibid.*).

The value of autonomy, understood as self-governance, explains why a practically useful moral theory is a better moral theory than a theory that fails to be useful, because

⁵ See Väyrynen (2006: 297) for a comparable view. Smith offers a briefer statement of her view on why practical usefulness is a desirable feature of a moral theory in a recent paper (2010: 74).

the first theory ‘allows for’ the achievement of autonomy by those who believe that theory, while the second theory does not. As Väyrynen (who argues along the same lines) puts it, ‘an ethical theory is better to the extent that it makes reliable strategies for acting well available to an agent who accepts it’, because autonomy is valuable, and so failing to be practically useful means that a moral theory ‘fails to protect autonomy understood as self-governance’ (2006: 298). We can summarize this explanation of why a practically useful moral theory is a better theory as follows: given that autonomy (understood as self-governance) is non-instrumentally valuable, and being self-governing involves acting in accordance with one’s moral beliefs, believing a moral theory that is practically useful is instrumentally valuable. In this respect, a practically useful theory is a better theory.

I will consider the details of this suggestion further in the next two sections; first, we should note a quite general problem with the approach both Smith and Väyrynen take. The explanation of why a moral theory is better if it is practically useful just sketched makes no mention of whether a theory’s being practically useful constitutes evidence for that theory; rather, it turns on the value of being in a state of believing such a theory. The respect in which a practically useful moral theory is a better theory on this view suggests that ‘being practically useful’ is a property which provides a *pragmatic* reason to believe a theory which instantiates it, but not an *evidential* reason. The reason to believe a theory that the theory’s having this property provides is, in this respect, similar to the reason to believe a theory that a theory’s having a property such as, say, ‘being declared to be the state’s moral theory in a dictatorship with an Orwellian thought-police’ provides: being in the state of believing a moral theory that instantiates either (or both) of these properties confers a benefit on the agent, and so there is surely *a* respect in which a theory which instantiates either one of these properties is a better theory (compared to a theory which does not instantiate them), but that respect is clearly *not* that of ‘more likely to be true’.

In what has become standard terminology, pragmatic reasons for belief are ‘reasons of the wrong kind’, whereas evidential reasons are ‘reasons of the right kind’. There is an ongoing debate over the general criteria for being a reason of the right (or wrong) kind,⁶ but we need not settle on any such criteria here; all we need in order to claim that there is a

⁶ Probably the most well-known proposal, and one that seems initially quite attractive, is the one presented by Parfit (2011: volume 1, chapter 2). For an argument for why even this account fails, see Schroeder (2012).

problem for the account that Smith and Väyrynen offer are these three claims: (i) there is such a distinction, (ii) it matters on what side of this divide a particular reason falls, and (iii) pragmatic reasons for belief are ‘reasons of the wrong kind’. These three claims are all fairly uncontroversial, even though how to draw the distinction between kinds of reasons is not. Some who have commented on the issue hold that reasons of the wrong kind are in fact not reasons at all, but it is sufficient for present purposes if we can maintain the view that (iv) ‘reasons of the right kind’ are the reasons that are central to what we all-things-considered have most reason to believe (do, intend, and so on), and this is much typically accepted even by those who *do* take reasons of the wrong kind to be genuine reasons.⁷ If, in brief, pragmatic reasons for belief are reasons of the wrong kind, and if reasons of the wrong kind are not central to what we all-things-considered have most reason to believe, then the reason to believe a practically useful moral theory instead of a theory that fails to be useful will not put much (if any) weight onto the scale. This in turn means that even if it is correct that a moral theory is better if it is practically useful, then that is unlikely to be a something that can support believing a useful theory over a theory that fails to be useful.

The problem (of providing reasons of the wrong kind) does *not* arise because of the specifics of the way in which Smith and Väyrynen explain the value of practical usefulness; it arises for any attempt to explain the respect in which a practically useful moral theory is a better theory which appeals to benefits befalling those who believe a theory that has the property of being practically useful. Only an explanation which holds that the respect in which a practically useful theory is a better theory in the sense of ‘being better supported by the evidence’ (and thus more likely to be true) is capable of avoiding the problem of providing a reason to believe of the wrong kind. However, an explanation of that sort would not, it seems clear, result in a position that deserves to be called ‘conciliatory’: the proponent of a practically useless moral theory (such as objective consequentialism) may well agree that practical usefulness is a desirable feature of a moral theory, yet surely, she will not agree that the true moral theory is likely to be practically useful – for accepting this claim would amount to admitting that theory she is proposing is unlikely to be true, and

⁷ See Reisner (2009: 258-259) for discussion and further references. Reisner defends the view that pragmatic reasons for belief can carry significant weight; see, in addition to the paper just mentioned, also his (2008). I discuss what weight pragmatic reasons that are provided by considerations of practical usefulness might have in a little more detail below.

that is *not* something that we can reasonably expect proponents of such theories to do. Either we seek common ground, in which case the reason to believe a theory because it is practically useful will be a reason of the wrong kind, or we can maintain that it is a reason of the right kind, in which case we will have to give up on the idea that all can agree that being practically useful indeed ‘counts for something’. But we cannot have it both ways.

3. The value of self-governance

Say that we want to maintain the conciliatory spirit of the suggestion, and so we accept that the respect in which a practically useful theory is a better theory provides only a pragmatic reason for believing the better theory. Furthermore, assume for the sake of argument, that pragmatic reasons, despite being reasons of the wrong kind, *do* provide genuine reasons for belief, and are therefore capable of putting significant weight onto the scale, sometimes carrying the day with respect to what we have most reason to believe. Could the pragmatic reason for believing a practically useful moral theory put *sufficient* weight onto the scale? I believe this is not the case; there are two grounds for doubting that it could. This section discusses the first of these, the following section the second one.

Smith and Väyrynen endorse what we can call a ‘content-neutral’ view of autonomy: an agent acts autonomously in the relevant sense whenever an agent’s decision to act in a certain way is based on an application of the moral view that the agent takes to be correct, irrespective of whether the moral view that this agent takes to be correct *is* correct. Yet we can wonder: is there really value in agents acting in accordance with their moral beliefs, irrespective of the contents (and the truth-value of those contents)? To get a grip on this question, imagine two worlds, W_1 and W_2 , both of which are populated by a variety of individuals; let us focus on just one of these individuals, whom we can call ‘Adele’. Adele performs the same actions in both W_1 and W_2 , but in W_1 she acts on the basis of what she takes to be the correct account of morality, while in W_2 she does not (instead, let us assume, she tosses coins to make her practical decisions, because she believes that this is as likely to get her to do what is right as trying to think matters through). If content-neutral autonomy is indeed of non-instrumental value, then we should expect to have the intuition that W_1 is a better world than W_2 irrespective of what moral theory we imagine Adele to believe in that world. But this is not our reaction – or at least, it is not mine.

If I imagine that Adele is a Nazi in W_1 , I am not at all inclined to judge that there is any added value to her acting on the basis of her moral beliefs in that world, as compared to W_2 , in which she tosses a coin in order to make her decisions. If anything, I am inclined to say that Adele's basing her decisions on her Nazi beliefs makes W_1 a *worse* world than W_2 . I find this particularly plausible if we imagine that Adele's actions are not monumental moral wrongs (e.g., committing genocide, starting a race riot, and so on), but instead are of a petty nature.⁸ Imagine for instance that in W_1 , guided by her Nazi beliefs, Adele quite often buys bratwurst but never buys hummus, she often reads papers that are published in the *Deutsche Zeitschrift für Philosophie* but never reads papers published in *Philosophia*, and often listens to *schlager* music but never to *klezmer* music; she does the very same things based on coin tosses in W_2 (where the coin toss consistently favors buying bratwurst instead of hummus, and so on). The view about the value of autonomy as self-governance Smith and Väyrynen appeal to implies that W_1 is better than W_2 , since the view imposes no restrictions at all on the content of the moral view that agents rely on when making practical decisions. But W_1 is not a better world than W_2 (and may well be a worse world). Therefore, this view should be rejected.

There is an easy fix to the problem posed by cases such as that of Adele: we can drop the assumption of content-neutrality, and hold instead that there is value in deciding what to do on the basis of the considerations that the moral theory one accepts singles out as settling the matter, *provided* the moral theory one accepts is correct. We might put this by saying that although there is value in doing 'the right thing for the right reasons', there is no value (and perhaps even disvalue) in doing 'the wrong thing for the wrong reasons', even if one mistakenly takes oneself to be doing the former instead of the latter. While this revision allows us to retain the claim that there is value in acting on the basis of the considerations that one takes to be morally decisive, by the same token, it undercuts whatever dialectical force the appeal to the idea that a practically useful moral theory is a better theory could have even if pragmatic reasons are taken to be capable of putting significant weight onto the scale. Here is why.

⁸ Some of those to whom I presented this case in conversation reported that they think it is worse to commit horrible moral wrongs on the basis of a coin toss than on the basis of deeply misguided moral beliefs, hence the elaboration of the example offered in the text.

On the revised account, the value of autonomy as self-governance is *conditional on the truth* of the theory by which autonomous agents guide their actions; if its value is conditional in this way, then the fact that a moral theory is better if it is practically useful provides no independent reason for believing one theory rather than another, over and above the evidential reasons that support these respective theories. Unless you have good reason to believe that a given theory is true, you have no reason to believe that anything of value will be realized if you were to believe this theory, and once you have the first set of reasons, the second set is redundant. In other words, the pragmatic reasons for believing a given moral theory, if these reasons stem from the value of autonomy (understood as self-governance) can only point in the same direction as the evidential reasons which support that theory, insofar they are to point anywhere at all. If the evidence available supports the verdict that a moral theory which fails to be practically useful is correct, then there will be neither evidential nor pragmatic reasons to believe a competing theory which is practically useful. Therefore, the fact that a useful theory is a better theory is, in effect, irrelevant to what we have most reason to believe. In brief, if the value of guiding oneself by one's moral beliefs is conditional on the truth of these beliefs, we can simply focus on the evidential reasons supporting competing moral theories, and ignore the issue of practical usefulness.

4. Weighing pragmatic reasons for believing a moral theory

The argument of the previous section rested on an intuition about the value of autonomy as self-governance which some may not share. If that argument does not succeed, we need not accept that considerations of practical usefulness can put weight onto the scale, as there is a second ground for concluding such considerations can be ignored when we are trying to make up our mind about which moral theory we should believe.

The value of autonomy (understood as self-governance) lies, presumably, not just in having the *capacity* for self-governance, but also – perhaps even primarily – in *exercising* that capacity, viz., in acting in the ways that are supported by one's moral convictions. Imagine that you believe a certain moral theory that is practically useful, rely on it in your practical deliberations, and thus act autonomously in the sense at issue. Imagine as well that the moral theory you believe is incorrect, and that this theory is incorrect not just in the sense that it provides incorrect explanations of why actions have the deontic status that

they do, but also in implying that various particular actions are right while in fact these actions are wrong, and *vice versa*. In brief, the theory you believe is not just *intensionally* but also *extensionally* incorrect. As a result of your accepting of and relying on this false-but-practically-useful moral theory, you will realize the value of acting autonomously. This, let us assume, provides you with a pragmatic reason to believe the theory that you believe.

However, since you will act in ways that the theory prescribes, you will act in ways that are not morally right. This makes it plausible that the pragmatic reasons for believing this theory will be outweighed, for the following reason. If some form of objective consequentialism is true, doing what is morally right is instrumentally valuable, because morally right actions are actions that are instrumentally best (or that are among the actions which are instrumentally best, in those situations where there is a tie). Since morally right actions are instrumentally best, you have most pragmatic reason to guide your actions by – and thus believe – whatever moral theory will lead you to acting rightly on every occasion; the less often you will act rightly if you believe a given theory, the weaker your pragmatic reasons to believe that theory are. Your belief in this false-but-practically-useful theory leads you to act in ways that are not instrumentally best, yet by definition, you can do better than you would do, if you were to believe the false theory.⁹ If there is some moral theory such that guiding your actions by it would lead you to act rightly more often than you would if you believe the false-but-practically-useful theory that you currently believe, you will have a stronger pragmatic reason to believe this competing theory. And if there is some theory such that guiding your actions by it would lead you to always act rightly, then *that* would be the theory you have most pragmatic reason to believe.

Perhaps there is no theory such that guiding your actions by it would lead you to always act rightly – I do not have an argument to show that there is. However, we can note that it is only when believing some false-but-practically-useful theory makes it *no less likely* that you will act rightly than believing the true theory that considerations of practical usefulness will come into play, for the value of doing what is morally right provides you with (much) stronger reasons to believe whatever theory leads you to do what is right, but

⁹ Consequentialism incorporates the ‘ought’ implies ‘can’-maxim into its fundamental deontic principle: you ought to do *the best you can*, not the best that you cannot do. So by definition, you always can do whatever the theory implies that you ought to do. This claim has been contested by Frances Howard-Snyder (1997), but I believe her argument does not succeed, for roughly the reasons Erik Carlson (1999) gives.

that theory, whatever it is, will not be an extensionally incorrect theory, or so it seems reasonable to assume. Provided that accepting it does not lead to a lower-than-possible number of occasions on which we act rightly, it would seem that the theory we have *most* pragmatic reason to believe is the *true* moral theory.¹⁰ And therefore, pragmatic reasons to believe some theory that are provided by the fact that this theory is practically useful are, it seems, likely to be outweighed by the pragmatic reasons to believe the *true* moral theory. If that is correct, then once again, we can conclude that the reasons for belief which stem from considerations of the value of practical usefulness are irrelevant to the question which moral theory we should, overall, believe.

It could be objected that the foregoing argument unfairly assumes that some form of consequentialism is true. In response, note that we can also make trouble for the idea that the pragmatic reasons for believing a practically useful moral theory can put significant weight onto the scale in a somewhat different way. Compare, again, two worlds: in the first of these, W_3 , the agents inhabiting it all act rightly, but they all fail to be self-governing; in the second world, W_4 , the agents inhabiting it are all self-governing in the relevant sense, but these agents never act rightly. Which of these worlds is better, W_3 or W_4 ? On the face of it, W_3 appears to be a far more desirable world than W_4 , as in W_3 the actions performed will have whatever features make actions morally right: utility is maximized, promises are kept, rights are respected, maxims can be universalized, or some such thing (or collection of things). There is, we can grant, something that makes W_4 good which is lacking in W_3 , but at the same time, there are many *more* things, and many *more important* things, that make W_3 good which are lacking in W_4 . The value that results from (or is inherent in) doing what is morally right is greater – and arguably *much* greater, even – than the value that results from autonomous actions that are not morally right, even if we do not understand right action in consequentialist terms. (While the non-consequentialist will of course deny that the value of an action's consequences exhaust its right-making features, presumably the value of consequences does matter to the deontic status of actions, and for at least a range of right actions, more good will result if they are performed than if they are not performed.)

¹⁰ It could be the case that, for some reason or other, believing the true moral theory has such an effect; it is this sort of concern which leads Henry Sidgwick to propose that utilitarianism should be an 'esoteric morality'. See Sidgwick (1980 [1907]: book IV, chapter V).

If this assessment is correct, the pragmatic reasons to believe a false-but-practically-useful moral theory are likely to be outweighed by the pragmatic reasons to believe a theory that leads to a greater number of morally right actions, even if morally right action is not understood along consequentialist lines.

Finally, we can ask: what if we assume instead that the false moral theory is merely intensionally incorrect, but not extensionally? On that assumption, accepting the false theory would not lead to the performance of wrong actions (since the theory singles out actions as right that are in fact right, etc.) and, moreover, the false theory would also be able to accommodate the value of autonomy, understood as self-governance. It appears that the pragmatic reasons to accept this type of false theory would be stronger than those to accept the true moral theory, for that theory is not backed by the reasons stemming from the value of autonomy. If this is correct, then the value of practical usefulness *is* capable of playing a significant role in shaping what theory we overall have most reason to believe, a critic might suggest. The reply to this worry is simple: if there is such an alternative theory, then the true moral theory would not be practically useless, for then the true theory would be *indirectly* useful, instead of directly. Take classical utilitarianism. Perhaps, due to time constraints and such, we cannot always identify which action (of those open to us) will in fact maximize utility. This does not all by itself imply that classical utilitarianism fails to be practically useful: as J. S. Mill already noted, we may well be able to formulate 'secondary rules', the following of which is a means to maximizing utility, but which are not couched in terms of utility, or related concepts. Now, these secondary rules are justified insofar as they are a means to this end, and what they do is put us in a position to single out, on many occasions, the action that we ought to perform, according to the utilitarian standard. But if we can do that, then utilitarianism counts as practically useful after all.

Generalizing, the point is that either the true moral theory is not practically useful, in which case the alternative theory is not extensionally correct, or the alternative theory is extensionally correct, in which case the true theory is practically useful (albeit perhaps only indirectly, by way of being capable of being supplemented by secondary rules). In the first case, the pragmatic reasons for believing the useful theory provided by considerations of practical usefulness are likely to be outweighed by the pragmatic reasons for believing the true theory. In the second case, the question of whether we should accept some

competing, incorrect moral theory on grounds of considerations of practical usefulness is moot, for the alternative theory is not a genuine alternative, but instead a statement of the secondary rules that help us guide our actions by the true moral theory.

It is perhaps worth emphasizing the general character of the argument developed in this paragraph: *any* putative account of why we should care about whether a moral theory is practically useful, irrespective of its details, runs into the problem that the value inherent in or generated by acting in accordance with the true moral theory is presumably much greater than the value of having a theory that can be used to make practical decisions (and acting on the basis of that theory), and thus the pragmatic reasons to believe a moral theory that are provided by the fact that this theory is practically useful are quite likely to be outweighed irrespective of the details of the account of the value of practical usefulness. This difficulty can be avoided only by identifying acting rightly with guiding one's decision-making by some set of moral beliefs or other, for only then will there be no conflict between the value of acting in accordance with the true moral theory and the value of being able to guide one's decisions by the moral theory one accepts. However, as this would amount to a radical form of moral subjectivism that, presumably, few are willing to accept, I believe we can safely set this option aside.

5. Conclusion

I have argued for three claims in this paper. First, the respect in which a practically useful moral theory is a better theory provides only a pragmatic reason – a 'reason of the wrong kind' – to believe such a theory. Second, if we follow Smith and Väyrynen and try to explain why practical usefulness is a desirable feature of a moral theory by reference to the value of autonomy, understood as self-governance, the pragmatic reasons for believing a moral theory will support only a theory which is also supported by evidential reasons. This is so because arguably, the value of autonomy (understood as self-governance) is conditional on the truth of the beliefs by which an agent guides her actions. Third, even if the value of autonomy (understood as self-governance) is not conditional on the truth of the beliefs by which an agent guides her actions, the pragmatic reasons for believing a practically useful moral theory are likely to be outweighed by the pragmatic reasons for believing the true moral theory, because it is more valuable to act rightly than to be self-governing.

In light of these three claims, the prospects for what I dubbed ‘the conciliatory rejoinder’ look dim: while all parties to the debate might indeed be willing to agree that it would be nice if we end up with a practically useful moral theory, this is not likely to be of much, if any, significance with respect to which theory we overall have most reason to believe. For all I have said in this paper, the failure to be practically useful may be a flaw in a moral theory, and perhaps this flaw *is* one that it is worth being concerned about. The only way it can be shown that this is so, though, would involve establishing that what we morally ought to do is something that cannot be epistemically inaccessible to us. And off hand, it seems unlikely that the case for that conclusion will rest on considerations about the value of having a practically useful moral theory. Either way, if the argument of this paper has been successful, conciliatory views about the significance practical usefulness can be set aside.

Acknowledgements Thanks to Govert den Hartogh, Emily Given, and several anonymous referees for their comments on versions of this paper.

References

- Bales, R. E. (1971). ‘Act-Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure?’. *American Philosophical Quarterly* 8: 257-265.
- Carlson, E. (1999). ‘The Oughts and Cans of Objective Consequentialism’. *Utilitas* 11: 91-96.
- Dorsey, D. (2012). ‘Consequentialism, Metaphysical Realism and the Argument from Cluelessness’. *Philosophical Quarterly* 62: 48-70.
- Feldman, F. (2006). ‘Actual Utility, the Objection from Impracticality, and the Move to Expected Utility’. *Philosophical Studies* 129: 49-79.
- Hare, C. (2011). ‘Obligation and Regret When There is No Fact of the Matter About What Would Have Happened if You Had Not Done What You Did’. *Noûs* 45: 190-206.
- Howard-Snyder, F. (1997). ‘The Rejection of Objective Consequentialism’. *Utilitas* 9: pp. 241-248.
- Hursthouse, R. (1999). *On Virtue Ethics*. Oxford: Oxford University Press.

- Parfit, D. (2011). *On What Matters*. Oxford: Oxford University Press.
- Reisner, A. (2008). 'Weighing Pragmatic and Evidential Reasons for Belief'. *Philosophical Studies* 138: 17-27.
- Reisner, A. (2009). 'The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem'. *Philosophical Studies* 145: 257-72.
- Schroeder, M. (2012). 'The Ubiquity of State-Given Reasons'. *Ethics* 122: 457-88.
- Sidgwick, H. (1980 [1907]). *The Methods of Ethics*, 7th edition. Indianapolis: Hackett.
- Smith, H. M. (1988). 'Making Moral Decisions'. *Noûs* 22: 89-108.
- Smith, H. M. (2010). 'Subjective Rightness'. *Social Philosophy & Policy* 27: 64-110.
- Väyrynen, P. (2006). 'Ethical Theories and Moral Guidance'. *Utilitas* 18: 291-309.