Moral Consistency Reasoning Reconsidered

Norbert Paulo[1,2]

Abstract

Many contemporary ethicists use case-based reasoning to reach consistent beliefs about ethical matters. The idea is that particular cases elicit moral intuitions, which provide defeasible reasons to believe in their content. However, most proponents of case-based moral reasoning are not very explicit about how they resolve inconsistencies and how they abstract principles from judgments about particular cases. The aim of this article is to outline a methodology—called Consistency Reasoning Casuistry—for case-based reasoning in ethics. This methodology draws on Richmond Campbell and Victor Kumar's naturalistic model for the resolution of inconsistencies between the content of intuitions about particular cases. I argue that reasons similar to those that motivate their model also support a more abstract form of moral reasoning that goes beyond mere resolutions of inconsistencies between case judgments and demands the formulation of more abstract moral norms. Consistency Reasoning Casuistry, it is argued, is a good candidate for a methodology for case-based moral reasoning that is in harmony with paradigms of contemporary moral psychology and that can accommodate the methodology implicit in the work of many contemporary ethicists.

Keywords Moral philosophy · Moral psychology · Practical reasoning · Analogy · Case-based reasoning · Thought experiments

In his witty introduction to *On What Matters* Derek Parfit praises Kant for *not* always being consistent: "'Consistency", Kant writes, "is a philosopher's greatest duty". That is not true. Originality and clarity are at least as important. And Kant's greatness chiefly consists in his having many original and fruitful ideas. If Kant had always been consistent, he could not have had all these ideas' (Parfit 2011, 1:xlii). Whether or not this assessment of Kant is accurate, Parfit himself went to great lengths to develop thoughts that are consistent; and it is hardly deniable that nonetheless he also had original and fruitful ideas. He was one among those

✉ Norbert Paulo
Norbert.paulo@sbg.ac.at

1   Department of Social Sciences and Economics, Legal and Social Philosophy, University of Salzburg, Churfürststr. 1, 5020 Salzburg, Austria

2   University of Graz, Institute of Philosophy, Attemsgasse 25, 8010 Graz, Austria

⚙ Springer

contemporary philosophers who try to reach consistent beliefs about ethical matters through reasoning with real-life or, more often, imaginary cases.

As I will demonstrate in turn, case-based moral reasoning is ubiquitous, but its proponents seem to lack a clear methodology. It is my aim in this article to suggest a methodology for case-based moral reasoning that emphasizes both consistency and creativity. Building on Richmond Campbell and Victor Kumar's idea of Moral Consistency Reasoning, I will propose a methodology that safeguards consistency and allows for original and fruitful ideas.

Let's start with the ubiquity of case-based reasoning. Everyday moral reasoning seems to work primarily by making comparisons between particular cases. Even where abstract moral principles are invoked, the justification or assessment of these principles usually proceeds by discussions of their implications for particular cases. This is also true of much of academic moral reasoning. Especially wide areas of applied ethics are oriented to cases (Arras 1991). There are several reasons for this: Case-based reasoning can help to bridge the ideological differences between different theory camps (e.g., deontologists vs. consequentialists) and thus appeal to a broader audience that would otherwise not bother to listen to arguments framed within one of the rival theories. Ethicists might themselves be uncertain about which, if any, abstract theory is the correct one. One method of finding reasons in support of some principles and against others is to search for moral principles that best cohere with a range of intuitions about cases, a process that will include refinements and revisions. Moreover, cases are often more illustrative than abstract principle-oriented reasoning, and the points being made appear more accessible. Some even speak of a 'primacy' (Weinberg 2016, 296) or 'priority' (Kagan 2001, 45) of cases in the sense that, given a conflict between an intuition about a rather abstract principle or theory and an intuition about a more concrete case, priority will almost always be given to the intuition about the particular case.

The idea behind case-based moral reasoning is that particular cases elicit moral intuitions, and that these intuitions provide defeasible reasons to believe in their content. Ethicists often seek to justify certain moral principles by reference to such intuitions about cases. They search for principles that can explain and ground the content of these intuitions (Kahane 2013). Proponents of Rawlsian Reflective Equilibrium, for example, search for moral principles that cohere with the content of certain intuitions about cases, which entails the requirement that the pre-selected intuitions themselves must not conflict. Were this not the case one would not know what exactly the principles must cohere with.

However, despite its ubiquity, most proponents of case-based moral reasoning are not very explicit about their methodology. Here is, for instance, Jeff McMahan about his approach to ethics, writing about the badness of different kinds of deaths:

> [I]n moral theory I take intuitions that are deeply and pervasively held to be presumptively reliable, so in thinking about the badness of death I will take strongly held intuitions … to be starting points for inquiry that are not to be lightly abandoned. My aim is in part to discover the deep values and assumptions that underlie our intuitive discriminations among different deaths; but if I am able to identify a unified foundation for our intuitions, I will also try to determine the extent to which it is rationally defensible and how it might be made more rigorous and coherent (McMahan 2002, 104, see also McMahan 2013).

And this is how Frances Kamm describes her methodology:

> Consider as many case-based judgments of yours as prove necessary. Do not ignore some case-based judgments, assuming they are errors, just because they conflict with simple or intuitively plausible principles that account for some subset of your case-based judgments. Work on the assumption that a different principle can account for all of the judgments. Be prepared to be surprised at what this principle is. Remember that this principle can be simple, even though it is discovered by considering many complex cases. (Kamm 2007, 5; see also Kamm 1992, 6–11)

In a series of papers Richmond Campbell and Victor Kumar have recently developed and defended Moral Consistency Reasoning (MCR) as a naturalistic model of moral reasoning with cases that explains both how we *actually* resolve inconsistencies and how we *should* resolve them (Campbell and Kumar 2012; Kumar and Campbell 2012; Campbell 2014; Campbell and Kumar 2013; Kumar and Campbell 2016; Campbell 2017). The basic idea is that, empirically, moral agents have a strong aversion against inconsistent responses to particular cases they regard as similar in morally relevant respects; they *feel* that they should treat like cases alike. MCR can also serve as part of a method to *justify* normative statements or principles and thus makes an important part of the methodology behind case-based moral reasoning more explicit. It is thus somewhat similar to Rawls's *Decision Procedure for Ethics* (1951) and to his idea of Reflective Equilibrium (Cath 2016; Scanlon 2002). Both methodologies aim at capturing—more or less accurately—how people think about what to believe and what to do; but they also provide frameworks for how one should think about these questions. None of these methodologies is tied to particular moral outlooks; they are open for various normative frameworks.

In this article, I will argue that MCR significantly advances the debate on the methodology behind case-based moral reasoning. However, it is not the final word on the matter: I will argue that MCR only covers a limited range of moral problems and that reasons similar to those which motivate MCR also support a broader methodology of case-based moral reasoning that goes beyond mere resolutions of inconsistencies between case judgments. I will point out some features this more comprehensive methodology should possess, explain how it advances MCR, and which role MCR plays in it.

This methodology is a form of ethical casuistry. Since it builds on casuistry and MCR, I shall call it Consistency Reasoning Casuistry (CRC). The article proceeds in four parts. I first introduce MCR and present reasons to accept it as a largely correct model of moral reasoning with similar cases (Section 1), before discussing reasons moral agents have for being practically consistent within the realm of MCR (Section 2). In Section 3 I argue that similar reasons speak for moral reasoning beyond inconsistencies. A natural candidate for this task is analogical reasoning; in Section 4 I discuss its relation to MCR. At the end I outline CRC and the respective roles of MCR and analogical reasoning in it (Section 5).

1 Moral Consistency Reasoning

Let us consider the real-life example Campbell and Kumar use to illustrate their model of Moral Consistency Reasoning (MCR):

Jan Baalsrud, a Norwegian resistance fighter whose mission has gone awry and who, pursued by the Nazis through mountainous terrain, collapses snow blind and near death at the door of a Norwegian family. The family begins to nurse him back to health but soon realize that it will be very hard to keep his presence a secret and secure his escape under the nose of the occupying forces. The mother of the family feels sympathy for her compatriot and admires his heroic efforts in the resistance but also feels grave concern about the threat that he poses to her family and the community. Without doubt they will be murdered if their rescue effort is uncovered. With great reluctance she tells her family that it would be moral madness not to surrender him to the Nazis, given the almost certain horrific consequences of doing otherwise. However, her son Marius, who has formed a bond with Jan, asks his mother what she would have a family do if Marius were himself in a parallel situation in Oslo. Though she might have replied that, despite her deep love for her son, they would have to do exactly what she is proposing for Jan, she and her family reacted differently. When they pictured Marius in the same situation they found themselves morally repulsed by the idea of sacrificing Marius to appease the Nazis and began to see Jan as a boy like Marius rather than as a stranger. In fact, they could see no morally relevant difference between the cases. At this point their moral understanding of their circumstances changed. They felt morally bound to do all they could to rescue Jan, despite the enormous risks. (Campbell 2014, 437 f.)

This line of interpersonal reasoning between Marius and his mother is an instance of MCR, because the mother's initial response to surrender Jan is inconsistent with her judgment that, in a morally similar case, Marius should not be surrendered. Faced with this inconsistency between two responses about particular cases the mother changes her mind and reverses her former judgment about Jan. As exemplified in this case, the MCR model concerns a practical inconsistency between judgments in particular cases as perceived by the mother.

Practical (or moral) inconsistency is not to be confused with formal (or logical) inconsistency. The case of Jan can be explained in terms of *formal inconsistency*: Initially, the mother came to believe that Jan should be sacrificed. Then she came to believe that, in the hypothetical Oslo scenario, Marius should not be sacrificed. Then she formed the belief that there is no morally relevant difference between the two cases. Finally, she formed the belief that, when there are no morally relevant differences, then both cases should be judged alike. And because she is much more certain about the Oslo scenario, she comes to believe that Jan should also not be sacrificed. There is a formal inconsistency in the mother's beliefs that she resolves by reference to further normative background criteria.

However, the case is more naturally explained in terms of *practical inconsistency*: The mother's initial response to the case of Jan is emotionally and motivationally opposite to her response to the Oslo scenario involving her son. These are responses to different cases. Were they responses to the same case, however, they would motivate actions that exclude each other. Even though this is also true when the responses are beliefs rather than emotions and motivations, the responses would be practically but not formally inconsistent because they are about different cases. As Campbell puts it, 'The pattern illustrates the related features of *practical inconsistency* and being *emotion-motivation responsive*, understanding that practical inconsistency does not in general exclude formal inconsistency and being responsive to emotion and motivation does not exclude the presence of belief.' (Campbell 2014, 440 italics his).

Summing up, two moral responses (beliefs, emotions or motivations) to similar cases are practically inconsistent if they would lead to actions that exclude each other when the cases were not only similar but the same. Campbell and Kumar use practical and moral inconsistency interchangeably—as do I in this paper.

There are several reasons to accept MCR as a largely correct model for moral reasoning with similar cases (even if not for all case-based moral reasoning as Campbell and Kumar claim). First of all, MCR comes with some empirical support. As they repeat time and time again, the empirical support they provide for MCR is of course tentative. Future evidence can support their view; but it can also favor alternative approaches to moral reasoning. They even suggest possible ways to test their theory (Campbell and Kumar 2012, 298–99). As the purpose of this paper is philosophical rather than psychological, I only mention some of the empirical evidence Campbell and Kumar present (for more evidence see Campbell and Kumar 2012, 276–89; for even more details Campbell 2017).

MCR seems to fulfil what Owen Flanagan has called the 'Principle of Minimal Psychological Realism (PMPR): Make sure when constructing a moral theory or projecting a moral ideal that the character, decision processing, and behavior prescribed are possible, or are perceived to be possible, for creatures like us.' (Flanagan 1991, p. 32) This principle, which overlaps with the better-known "ought implies can" principle, has been reformulated to apply to (second-order) moral epistemic decision procedures (such as MCR, Rawls's Decision Procedure, and his Reflective Equilibrium) which are used to determine the proper contents of first-order morality (BLINDED) Paulo (forthcoming). One element of this is that a methodology that is not at least minimally psychologically realistic is also not warranted philosophically, because it would demand what is impossible 'for creatures like us.' Another element of the naturalistic idea of psychological realism is, roughly and all else equal, that the better a methodology fits with the state of the art in the cognitive sciences, the higher its descriptive accuracy; and the better the descriptive accuracy, the higher the likelihood of the methodology being successfully applied by human beings.

Campbell and Kumar introduce a 'minimalist moral dual process' model (or MMDP) of moral cognition in which two distinct systems guide moral judgment (Campbell and Kumar 2012, 276). System 1 is normally intuitive, fast, automatic, affective and impenetrable; system 2 is, roughly speaking, the reasoning system. It is deliberate, slow and controlled; it is cognitive in the sense that it does not involve any significant emotional or motivational states. This distinction between system 1 and 2 is widely shared by many psychological models of moral reasoning (cf. Haidt 2007). The MMDP behind MCR is a minimalist model of moral reasoning in that it does not make any controversial claims as to whether system 2 is necessarily rule-based, one of the systems supports or undermines particular moral theories, or one is more reliable than the other (for a similar minimalist model of moral judgment see May 2018, chap. 3). Campbell and Kumar spend much effort to outline empirical evidence for the claim that moral change and moral progress largely happen through a long-term process of MCR that ultimately changes moral intuitions and, subsequently, moral beliefs—and not so much through contemplation of abstract moral rules.

What is particularly important for MCR is that, according to MMDP, moral evaluations are guided by both system 1 *and* system 2. Let me explain how system 1 and 2 interact in MCR and exemplify this with Singer's well-known discussion of the moral obligation of citizens in affluent nations to relieve famine in the developing world (Singer 1972; see also Campbell

2014, 441 ff.). How is it possible, Singer asks, that most of us find it unacceptable to allow a child to drown in a nearby pond, even if rescuing the child would ruin a new suit, when many of us also believe that it is acceptable to spend one's disposable income on luxury products, even if this money could effectively feed a starving child in the developing world? In both cases, help is needed and can be offered with only a minor sacrifice. In the terminology of Campbell and Kumar, the drowning child scenario serves as the *base situation*, the famine scenario is the *target*; the process has five steps.

When faced with these cases, system 1 issues fast responses (*step 1*). The responses of many will be that we have a moral duty to save the drowning child, but none to save the starving child. Next, system 2 recognizes that the cases are similar, but the responses different (*step 2*); system 2 might, for example, suggest spatial proximity as a morally relevant difference between the cases. Proximity would then be "checked" in system 1. That is, system 1 provides us with an intuitive response to the question whether or not proximity is a morally relevant factor. If system 1 has no normative resources to react to proximity, it automatically yields a negative affective response to the inconsistency. Only now is the inconsistency *felt* and thereby really understood. Suppose that spatial proximity in itself is indeed not intuitively regarded as being morally relevant. System 1 then develops a motivation to resolve the inconsistency that is felt (*step 3*) between the responses to the base and the target situations. Two possibilities to resolve the inconsistency are given in *step 4*. The first is to deliberate consciously (system 2) about the two responses and to revise the one that is less tenable (*4a*). The process of consistency reasoning ends here. When presenting the two scenarios, Singer's idea was that readers would find the judgment about the target situation (the starving child) less tenable than the one about the base situation and would thus revise their judgment about the target.

The second possibility (*4b*) is to contemplate consciously (system 2) and to come up with a new tentative morally relevant difference between the cases. When pursuing this possibility, new cases could be described in *step 5* instead of revising the response to at least one of the cases, which lack the new allegedly relevant difference but which are nevertheless inconsistent. The process would then start again; MCR is thus partly recursive. For instance, Singer anticipated the critique that the two cases are relevantly different because the agent is the only person who can help; the starving child could, in contrast, be saved by many other persons. Singer answered this objection by changing the base situation by postulating that other people are also standing around the pond and are doing nothing to help the drowning child. MCR starts again from step 1. Does the fact that there are other (inactive) bystanders affect one's own obligation to help the child? Singer argues that the mere fact that others are not fulfilling their moral obligation is not morally relevant for one's own moral obligation.

MCR's fittingness with contemporary moral psychology, its integration of system 1 and system 2 processing and its power to making explicit the reasoning in real-life scenarios (the case of Jan) as well as in thought experiments (Singer's case) are but three reasons to— tentatively—accept MCR as a largely correct model of moral reasoning with similar cases. MCR ultimately aims to yield consistent reactions to morally similar cases.

In the following sections I provide reasons for being consistent in moral reasoning and argue that these reasons support MCR. I will also argue, however, that MCR only covers a limited range of moral problems and that reasons very similar to those behind MCR also suggest a broader methodology of case-based moral reasoning that goes beyond mere resolutions of inconsistencies between case judgments.

2 Reasons for Consistency

Why should we be consistent in the first place? One reason for being consistent is that formal consistency is a minimum condition for truth. Of any two formally inconsistent moral intuitions, at least one does not track the moral truth. Given that practical inconsistency and formal inconsistency often go hand in hand, as explained above, practical inconsistency often also implies that at least one response does not track the moral truth. This *epistemic* reason for consistency might stand behind the two motivating reasons for consistency referred to by Campbell and Kumar: *Personally*, we long to have a personal sense of being reliable, authentic, having integrity and so forth. We feel unrest and discomfort when we lack that sense. While practical inconsistency is not the only possible cause for lacking that sense, it is a very powerful one (Horne, Powell, and Hummel 2015).

Socially, Campbell and Kumar say consistency matters because 'Not to care about being consistent in one's moral attitudes and feelings … would undermine one's credibility as a moral agent, not to mention as a trustworthy and responsible parent; one's moral responses would be unpredictable and one's character unreliable'(Campbell and Kumar 2012, 274). Being unreliable and untrustworthy is problematic, especially in groups whose members regard themselves as being strongly interdependent, for instance because they must coordinate their activities. Whether or not a perceived inconsistency will reduce one's social credibility will most likely depend on a complex web of background information. Returning to the example of Jan, the issue and its outcome will depend upon whether the respective social group draws strong distinctions between the in-group and the out-group and whether Jan is considered an out-group or in-group member. In a context of national identity, Jan will be considered 'one of us' (a fellow Norwegian) and hence the inconsistency of treating him in a different way than Marius will be troubling. In the absence of this method for sorting Jan out as 'one of us,' consistency reasoning may not yield the result described above. However, the same kind of consistency reasoning reiterates at this level, since one may well ask here whether it is a morally appropriate response not to treat Jan as an in-group member.

3 Why Consistency Is Not Enough

Similar epistemic, personal and social reasons also speak for ethical theorizing beyond the mere resolution of inconsistencies. I take ethical theorizing, broadly understood, to be the systematization of normative ethical thought. Such theorizing can gravitate towards one organizing and guiding principle, towards sets of rules or principles, or towards certain virtues which are to be promoted. All these forms of ethical theorizing have two things in common. First, they search for patterns to structure and, where possible, to reduce the complexity of moral phenomena. In other words, they try to establish an epistemic framework for understanding these phenomena and to distinguish them from related spheres such as custom, religion or law. Second, they attempt to establish some means for guiding decisions in moral questions.

There are many reasons for ethical theorizing in this sense. Among these reasons are epistemic, personal and social reasons similar to those that speak for consistency. *Personally*, we long to gain a certain grasp of the moral landscape, in order to have at least some understanding of what it is that makes actions good or character traits praiseworthy. Most of us also wish to be in accord with this understanding. This kind of understanding is conducive,

if not necessary, for being a reliable moral agent that shows integrity. Our moral judgments are closely linked to our emotional dispositions and responses to moral or immoral acts (Prinz and Nichols 2010). Yet how can one be a person that judges and acts reliably according to such an understanding? This will require some serious thinking about morality. One will need to contemplate different options of how to organize moral phenomena. We must make up our minds about what kind of person we wish to be, how we want to live and act and how we wish to behave in regard to ourselves, other persons, animals and the environment. This is not to say that everyone needs to be a moral philosopher, but some limited understanding and guidance is warranted to negotiate the social world. This kind of thinking about morality, together with a certain level of experience, also seems to be important for one's moral authenticity. Challenging and contemplating one's moral views usually enhances the value of authenticity. The kind of reflection and theorizing, however mundane, necessary for valuable forms of moral authenticity, integrity and reliability likely exceeds the mere resolution of inconsistencies between morally similar cases.

More importantly, large areas of moral thought are not adequately covered by MCR. Moral problems usually occur one at a time and not in the form of inconsistencies or dilemmas. MCR alone can only very indirectly account for these situations. What MCR advises in normal situations—where single moral issues arise that do not come in pairs of cases—is to deal with the respective moral issue by considering other cases that one was confronted with in the past, or to check one's response to comparable hypothetical cases. So MCR does not allow for a direct answer to a moral question. One might well have an intuition about the case in need of resolution, but MCR always asks for a comparison of this intuition to intuitions about relevantly similar cases. This not only requires an independent understanding of the more abstract question which features are relevantly similar, but also considerable effort und creativity. MCR is not unable to account for most normal moral problems, namely those that do not occur in the form of inconsistencies or dilemmas. But it can account for these problems only in a somewhat indirect way, namely through the comparison with another case one has to remember or imagine.

Moreover, some moral dilemmas are not cases of practical inconsistency. Consider the example of the student of Jean-Paul Sartre's who came to see him in occupied Paris during World War II. He asked his teacher whether he should stay home and care for his mother or join the Free French resistance forces instead to fight the Nazis. Let us assume that both causes were morally good and that the student could not have pursued both. We can assume that the choice between the two options has a heavy emotional charge and that the student felt strongly for both his mother *and* for his country. The young man approached Sartre in order to obtain practical advice. He was thus, in my supposition, open to moral reasoning (although Sartre makes the point that this is a case where reasoning is doomed to fail). Dilemmatic situations of this kind are somewhat similar to MCR. They are, however, still not covered by MCR, because the agent does not face a pair of practically inconsistent responses.

To see why, recall Campbell and Kumar's understanding of practical inconsistency: two moral responses (beliefs, emotions or motivations) to similar cases are practically inconsistent if they would lead to actions that exclude each other when the cases were not only similar but the same. The Sartre scenario is not a case of practical inconsistency although the two possible actions are mutually exclusive. This is because the cases are not relevantly similar. There is no inconsistency of beliefs, emotions or motivations in the sense required for MCR, namely such that one feels one should change one moral response in order to be an authentic moral agent or in order to show integrity. It is the very point of dilemmas such as this one that both responses

are understood to be appropriate, which makes the question how to act so hard to answer. MCR cannot account for problems of this kind. The problem MRC is meant to work with is that one has two responses to different cases and these 'responses together seem morally indefensible because we see no morally relevant difference between them.' (Campbell and Kumar 2012, 296).[1]

Sartre's student may well have felt emotional discomfort; but this discomfort is different from that felt in MCR. The former is external (i.e., the reason lies in the tragic circumstances), the latter is internal (i.e., the reason lies in the agent). Moreover, even if the student compared his responses to the two options he faced with previous situations or to hypothetical cases, this would not have helped him in adequately addressing the moral issue he faced. The point here is not only that MCR cannot *resolve* Sartre's student's moral quandary. It is quite possible that there is no good solution. The point is that MCR alone does not even have the means to adequately *address* the problem. A more comprehensive methodology of case-based moral reasoning should at least have the means to capture what is problematic about such dilemmas.

The need to employ modes of reasoning beyond MCR is also an *epistemic* reason for theorizing. Otherwise one would not have the means to investigate the truth about all moral questions (but only about cases triggering inconsistent responses). It is also at least possible that some generalizations in the moral realm are true (and not only judgments about particular cases). A full-fledged moral methodology should have the means to explore this possibility.

The difference between MCR and the resolution of other dilemmatic situations can also be viewed from an outside perspective: When the student explained his dilemma, Sartre probably felt compassion for him. The student was after all not responsible for the dilemma. In cases of inconsistent responses to morally relevant similar cases, by contrast, the agent herself is responsible for the inconsistency. When someone exhibits a moral inconsistency one automatically feels a disapprobation toward the agent (Campbell and Kumar 2012, 290). This is one of the *social* reasons for MCR. There are similar social reasons for ethical theorizing beyond inconsistencies between case judgments. Every form of social organization—be it a family, a company, or a country—depends on the reliability and credibility of individual moral agents. In order to organize and coordinate social life we need to be able to predict, at least in rough terms, how others will think and act. All these kinds of social organization depend on ethical theorizing of some sort—and the mere resolution of inconsistencies between similar cases would not suffice for any of these.

The upshot is that epistemic, person and social reasons similar to those that speak for practical consistency in MCR also speak in favor of moral theorizing beyond the resolution of inconsistencies between judgments in particular cases. What is at stake here is an understanding of how to be a reliable moral agent that is truthful to her basic moral convictions while also recognizing the complexity of morality and the need to revise one's convictions and to adapt to

---

[1] At one point, Campbell and Kumar talk about situations such as Sartre's—where responses to two cases are mutually exclusive without being morally inconsistent—in terms of inconsistency (Campbell and Kumar 2012, 301–3). However, the discussion there does not fit the general model of MCR. As I explained in the main text, the two situations are, in fact, relevantly different; and Campbell and Kumar have introduced a specific understanding of 'moral inconsistency', and MCR as a method for the resolution of such inconsistencies, such that MCR does not apply to responses to cases which are mutually exclusive without being inconsistent. In the respective section, Campbell and Kumar sometimes talk about 'incompatibility' instead of inconsistency, which might be taken as a hint that they meant to distinguish between the two notions (but didn't). At least, they do not explain how their notion of practical inconsistency is meant to apply to cases such as Sartre's.

new circumstances. The challenge is thus to retain flexibility in one's moral thinking while also being reliable and predictable.

4 Analogical Reasoning

A rather obvious way to expand case-based moral theorizing beyond the scope of MCR is analogical reasoning. Intuitively, analogical reasoning informs our judgment of one thing using what we know about another thing when the two share some but not all features. Imagine that Cass has a German shepherd dog that is gentle with children. Whenever Cass sees another German shepherd dog, he assumes that this dog is also gentle with children (the example is taken from Sunstein 1993, 743). It is in this sense that analogical reasoning with cases provides positive, if tentative, reasons for future decisions in a way that MCR does not. The problem is that knowing the (gentle) behavior of only one dog is not a sufficient basis for making judgments about other dogs. But although analogies might not yield *conclusive* reasons to accept a judgment, this is no reason to *reject* the judgment reached by analogical reasoning. We will frequently not have more certainty about more abstract principles or theories than we have about the reliability of analogies (Kamm 2013, 553 f.).

In the case of morality the idea is thus that cases in which we are more certain about our moral response (the base situation) can inform our judgment in similar cases to which our response is less certain (the target situation). Analogical reasoning can complement MCR because it is creative and informative in a sense that MCR is not. Both are means to follow the precept to treat like cases alike and different cases differently. But analogical reasoning is often creative in the sense that it provides information about what to do or to believe. In this respect, analogical reasoning is like the Golden Rule—to treat others as one would wish to be treated by others. MCR, in contrast, is an exercise in exclusionary thinking. It merely tells you what not to do or believe. In this respect, MCR is similar to the Categorical Imperative—to act only according to that maxim whereby you can at the same time will that it should become a universal law.

One might think that MCR just is a version of analogical reasoning, or vice versa. Campbell and Kumar argue against this suspicion by pointing to an assumption that seems to underlie analogical reasoning, which is its domain-generality: The more similar two cases are, the more likely it is that the judgment from the base situation is also warranted in the target situation. MCR by contrast, is meant to be domain-specific. That is, it is hypothesized that, empirically, it only works with specifically moral similarities and is largely indifferent to general (i.e., non-moral) similarities between cases (Campbell and Kumar 2012, 297 f.). The idea is that people are much better to detect inconsistencies in the moral domain than in other domains; and that this is not the case for analogical reasoning. Whether or not this is the case remains to be seen; this is a question than requires empirical inquiry.

Be this as it may, I doubt that this alleged difference between MCR and analogical reasoning really is relevant when we talk about moral methodologies. Recall Singer's drowning child scenario. It is in no way similar in domain-generic terms to the famine relief scenario; these two are only similar when considered in a domain-specific perspective, e.g., from the point of view of a moral duty to help those in need. Domain-generality in this sense is not very useful for moral reasoning. That a base and a target situation share 10, 100 or 1000 features does not make it any more reasonable to assume that the moral judgment in the base situation is also warranted in the target situation when these features are not relevant in the relevant

(moral) domain. Analogical reasoning does not seem to be necessarily domain-general in that sense (cf. Spielthenner 2014; Bartha 2010). One might argue that people are better in detecting analogies (and disanalogies) in the moral domain as compared to other domains. If Campbell and Kumar want to distinguish between analogical reasoning and MCR by pointing to their domain-generality or domain-specificity, what they should argue is that, in the moral domain, people are better to detect inconsistencies than they are in detecting analogies (and disanalogies).

Nonetheless, I do not regard MCR as a mere version of analogical reasoning. Rather, MCR might be understood as covering some elements of analogical reasoning; it might be a particularly helpful method for reasoning with disanalogies—a naturalistic method to figure out which cases are to be treated differently. Just as successful disanalogical reasoning, MCR tells people what *not* to do. However, as we shall see in Section 5, MCR also plays roles within a broader methodology that analogical reasoning does not.

In Section 3, I argued that reasons similar to those that speak for consistency in MCR also speak in favor of moral theorizing beyond the mere resolution of inconsistencies. A methodology should help to retain flexibility in one's moral thinking while also allowing for reliability and predictability. While analogical reasoning is a first step for case-based moral theorizing beyond MCR—mainly because it adds flexibility and allows to assess moral situations that do not come in the form of inconsistencies—, it hardly suffices to allow for reliability and predictability. In the remainder of this article I outline a more comprehensive methodology for case-based moral reasoning that fulfills all of the tasks mentioned; I will highlight the contributions of MCR to this methodology.

5 Consistency Reasoning Casuistry

Let's call this methodology Consistency Reasoning Casuistry (CRC). CRC builds on MCR and casuistry, and it borrows from legal theory. Many authors recognized the close relation between casuistry and the reasoning in legal systems of the common-law tradition (for instance, Kamm 2013; Arras 1991). CRC draws on a common-law inspired version of ethical casuistry that I defended elsewhere (BLINDED) Paulo 2015. This version of casuistry complements MCR rather naturally, which enables it to account for the shortcomings of MCR. As we have seen, MCR stops after an inconsistency is resolved; it does not abstract from the particular cases to more abstract rules. In CRC, in contrast, the rules established in analogical reasoning subsequently function as an epistemic framework for the understanding of new moral phenomena. Let me explain.

In common law, a case resolution usually consists of the following elements: the case description; the (legal) question at stake (i.e., what triggered the examination of the case); the reasoning about this question; the ruling on the (legal) question, i.e., the *ratio decidendi*; and the result for the particular case, which follows from the ruling (Lamond 2014; for discussion see Stevens 2018). Note that—although the general legal methodology is case-based—the results are ultimately derived from the rule (i.e., the answer to the legal question).

I suppose that this is also true for ethical casuistry, where (similar to the *ratio decidendi* in legal cases) maxims and paradigms are understood as rules that are articulated within a set of circumstances (Jonsen 1991, 298 f.). Recall the case of the Norwegian resistance fighter Jan. In this case, possible maxims include, "help those who are in danger" and "protect your family." If Jan's case is understood as an instance of the duty to help, then Marius's mother initially

decided wrongly to hand Jan over to the Nazis. But if one understands this as falling under the paradigm of protecting one's family, then handing Jan over to the Nazis might have been the morally correct thing to do. Note that the normative source for future cases is not the particular case resolution but the paradigm as articulated within a certain case. In fact, all of the elements of a case resolution in common law are also elements of ethical casuistry (BLINDED) Paulo 2015.

Common-law reasoning can inform case-based moral reasoning with regard to the understanding of the application of paradigms in new cases and to the refinement of paradigms in light of novel cases. Let me explain this version of casuistry in more detail, step by step, emphasizing the role MCR plays in CRC as well as the ways in which CRC is able to deal with MCR's problems as discussed above.

5.1 Case Resolution

As every moral methodology—including MCR, Rawls's Decision Procedure and Reflective Equilibrium—ethical casuistry has no fixed moral content. It is meant to work with a plurality of moral inputs. Since none of these methodologies are attached to or based on certain substantive normative theories, they have no obvious means to determine the criteria of a situation that together constitute what we call "the case" and formulate as the *case description* (Arras 1991, 39 f.). So, unless one starts from tentative frameworks and paradigms provided by institutions such as religious communities, cultural norms or legal systems one has to start from one's own moral experience. As Kamm explains, 'in order to have a judgment about a case you really have to situate yourself in the case. … You have to imagine [the] case in detail. … [Y]ou have to sink into these details. … Just *situate* yourself in the case' (Kamm 2011, 23 her italics). To reach the moral imagination necessary to understand difficult cases, various means such as moral education, personal experience, or reading novels might be helpful. MCR can explain how to work through alternatively possible case descriptions, each highlighting different aspects of the same situation; it can explain how people do this moral work on the ground, how they test—employing both system 1 and system 2—possible morally differences between cases, which are then, in turn, taken to be the main ingrediencies of the case description.

The *moral question* is triggered by the examination of the case. It points to the moral problem underlying the case. The question can have many forms and varies with the depiction of the case. In the case of Jan, for instance, the question might be whether there is a morally relevant difference between saving a stranger (like Jan) and saving a person close to oneself (like Marius is to his mother) in otherwise similar circumstances. The question can also be if one should sacrifice a stranger when not doing so would put one's family and community in danger. In MCR and in casuistry in general the moral question can have various degrees of abstraction.

But in MCR the question must be one of inconsistency between at least two case judgments. Above I argued that MCR does not adequately cover cases such as Sartre's. Remember that the Sartre scenario is a moral dilemma that is not a case of practical inconsistency. The two actions the student is contemplating—joining the resistance and caring for his mother—are mutually exclusive, but they are not relevantly similar. There is no inconsistency of beliefs, emotions or motivations in the sense required for MCR, namely such that one feels one should change one moral response in order to be an authentic moral agent or in order to show integrity. If one were to use only MCR as a moral methodology one would not even know what the moral question is in cases such as Sartre's. By definition, MRC only

applies to responses to different cases when these 'responses together seem morally indefensible because we see no morally relevant difference between them.' (Campbell and Kumar 2012, 296) In dilemmas such as Sartre's, the student's responses to both alternative actions are morally defensible. This is what makes it a dilemma. They are both appropriate yet mutually exclusive. The moral question is how to resolve such a dilemma between two praiseworthy actions that are mutually exclusive. MCR cannot account for such moral questions. CRC does not have this kind of restriction.

Depending on the case, the moral question and the background theories referred to, the *reasoning element* can take various forms. The most important forms of reasoning are direct rule application and creative analogical reasoning (for sophisticated models of analogical reasoning see Brewer 1996; Bartha 2010). Another form of reasoning is distinguishing, which I will discuss in turn. MCR is yet another mode of reasoning that will often be used to exclude the less tenable response to a case or to come up with new cases for comparison (as illustrated with the case of Marius' mother). Note that the moral question and the reasoning element are abstractions from the particulars of the case under scrutiny. This is also reflected in MCR. Recall the five-step procedure illustrated with Singer's drowning child scenario above. In this procedure, steps 2, 3 and 4b also abstract from the particular case in that system 2 tentatively suggests relevant differences and isolates them for examination by system 1 and 2 before the result is then built into particular cases again. The reasoning thus leads to a tentative rule about which features count as morally relevant.

This tentative rule is similar to what is known as the *ratio decidendi* in common law reasoning. Reasoning about Jan's case one should, among other considerations, take into account the general maxim to protect one's family. An example for a tentative rule as the result of the reasoning process is this: in instances of the maxim "protect your family," a family's surrendering of a fellow countryman to an occupation army is morally legitimate. In contrast, this seems to be the general rule Marius' mother arrives at after Marius' intervention: In instances of the maxim "protect your family," exceptions from the general wrong of putting one's family at risk are justified when this is necessary to fulfill one's duty to help someone in an unprovoked situation of danger (as long as there is a considerable chance that one can help without sacrificing one's family).

Results for particular cases are derived from the tentative and defeasible rule that answers the question. Above I said that the case of Jan might pose different questions. Let us stick to the question if one should sacrifice a stranger when not doing so would put one's family and community in danger. When we now apply the quite complex rule "in instances of the maxim "protect your family," exceptions from the general wrong of putting one's family at risk are justified when this is necessary to fulfill one's duty to help someone in an unprovoked situation of danger (as long as there is a considerable chance that one can help without sacrificing one's family)" to the case, then we arrive at the conclusion that Marius' family should not surrender Jan, at least as long as there is a considerable chance that the rescue effort will not be uncovered (which is a factual question).

Having a formal structure for case-based reasoning, such as the one just presented, helps to get the moral issues straight and to prevent misunderstandings; it also allows for pointed criticism and, ultimately, enhances reasoning. Even though it does not provide moral guidance in a substantial sense, it does enhance case-based moral reasoning through the structural constraints it provides. Keeping the issues thus apart clarifies the individual case resolution and makes it easier to use the rule articulated in the case as a normative source for future cases, which is the idea behind case-based moral theorizing. Combining ethical casuistry with MCR

and similar means of reasoning, CRC can yield answers to substantive normative problems. CRC, as suggested here, is thus (a sketch of a) comprehensive methodology for reasoning with moral cases.

5.2 Application of Paradigms: Reliability

As mentioned earlier, the normative source for future cases is not the particular case resolution but the paradigm in the form of a tentative rule as articulated in the particular case with all its specific circumstances. It is worth stressing that the rule only incorporates some of the circumstances of the paradigm case; further circumstances of the paradigm might well be important in cases of doubt concerning if or how the paradigm rule should apply to a novel case. It is in this sense that the paradigm rule only provides defeasible reasons for judgments about novel cases. Careful casuists can thus avoid the so-called additive fallacy (Kagan 1988). That is, they can use paradigm rules to inform their judgments in new cases without neglecting the possibility that similar reasons have different roles (excluding, strengthening, favoring etc.) in different contexts.

Thus when a new case has the features of a prior case (paradigm), one can simply apply the rule of the paradigm case, other things being equal. Analogical reasoning—as is characteristic of casuistic moral reasoning—is only needed when there is no rule that applies without ambiguity, when the prior and the new case share only some relevant features, or when the relevant features function differently in the cases (again, e.g., excluding, strengthening, favoring etc.). That is, analogical reasoning is important where a few cases only have been decided through CRC. The more decisions have been made—and are accepted as relevant precedents—the more rules there are to be applied directly (without analogical reasoning).

MCR is useful for testing whether the two cases really are sufficiently similar to apply the old rule to the new case. That is, MCR describes how one can recognize, emotionally, whether the two cases are relevantly similar and should thus be treated alike. The whole structure for case resolutions and the sticking to established rules is a significant step toward reliability and predictability in moral matters, which are two important aspects of the personal and social reasons underlying MCR as well as more abstract moral theorizing.

As highlighted above, MCR does not adequately cover moral problems that do not come in pairs of cases. MCR does not allow for a direct answer to a moral question. One might well have an intuition about the case in need of resolution, but this does not suffice for MCR. The method requires that one asks for a comparison of this intuition to intuitions to relevantly similar cases. This, in turn, requires considerable effort und creativity. So MCR is not unable to account for most normal moral problems, namely those that do not occur in the form of inconsistencies or dilemmas. But it can account for these problems only in a somewhat indirect way, namely through the comparison with another case one has to remember or invent. I regard this as one of the main weaknesses of MCR. But once MCR is understood as being part of a broader methodology such as CRC, the problem vanishes. One might want to stick to the idea of case-based moral reasoning as always working through the comparison of responses to cases. Such a minimalistic methodology might fare well in many difficult moral cases such as Singer's or Jan's. But it fails in other difficult moral cases such as Sartre's and is at least problematic in more mundane everyday moral problems that do not come in pairs. A methodology such as CRC that has a place for moral rules as abstractions from previous paradigm cases is advantageous over MCR because it can allow for simple applications of these rules in clear moral cases.

5.3 Refining Paradigms: Flexibility

But what shall one do when MCR reveals that the two cases are relevantly dissimilar and that the paradigm rule does not seem to fit the new case? So far I have merely explained how CRC establishes stability and continuity. It does so by establishing tentative rules in initial cases which are then applied to future cases that are similar to the initial cases. But there is also a need for flexibility in moral theorizing. CRC has several means for flexibility: The rules established in cases will often be in need of interpretation, which already provides some flexibility. These rules are also tentative and defeasible; one important means to account for this tentativeness and defeasibility is distinguishing as known in the common law tradition.

Distinguishing is the creation of a new rule by narrowing down the paradigm rule under certain constraints. The effect of distinguishing is that one does not have to follow a paradigm rule although it applies to the case at hand. There are two constraints on this kind of narrowing. In Joseph Raz' words: '(1) The modified rule must be the rule laid down in the precedent restricted by the addition of a further condition for its application. (2) The modified rule must be such as to justify the order made in the precedent' (Raz 2009, 186).

I believe this effect of distinguishing captures more precisely what the casuists Jonsen and Toulmin wanted to get at when they said 'openness to novelty and a readiness to rethink the relevance and appropriateness of older rules or principles' are important but that this 'does not mean questioning the *truth* of those principles: they will remain as firm and trustworthy as ever' (Jonsen and Toulmin 1992, 316). That the principles kept in the paradigm remain 'firm and trustworthy' has two aspects: first, as secured by the two constraints on distinguishing, the new (narrower) rule keeps faith with the paradigm in that it must also justify the decision in the paradigm case although it is narrower than the paradigm rule. Second, the paradigm rule remains part of the normative system. It can still be invoked as a normative source for future cases. Moreover, distinguishing explicates more clearly, I suppose, how the inter- and intra-personal reconciliation of inconsistencies is described in MCR in cases 'when two moral norms are internalized and expressed in many instances without conflict until a new situation arises in which the norms cannot be applied without inconsistency' (Campbell and Kumar 2012, 301–3).

Above, in Section 3, I already referred to this passage from Campbell and Kumar. I argued that they there use the term "inconsistency" in a way that is does not fit their definition of practical or moral inconsistency, which might be the reason why they sometimes speak of "incompatibility" instead. The point is that MCR does not apply to responses to cases which are mutually exclusive ("incompatible") without being practically inconsistent. Sartre's is the case I used to illustrate this difference. His student faces two mutually exclusive possible actions, both of which are morally good, which is why they are not practically inconsistent. The student's situation can indeed be described as one in which two internalized moral norms (of family duties and of civic duties, say) that are normally expressed in various instances without conflict suddenly cannot be applied without incompatibility because the situation is one of the rare instances in which the norms conflict. It is precisely such situations that call for distinguishing. The moral question in Sartre's case is how to resolve the dilemma between two praiseworthy actions that are mutually exclusive. MCR cannot account for such moral questions. Successful distinguishing in CRC, however, yields what is asked for, namely how to amend the internalized norms such that it is clear what to do in the particular case, without thereby abandoning the norms for they are still useful and appropriate in most other cases. Raz' two constraining conditions alone do not solve the dilemma, of course. But they

significantly limit the range of possible options for specific norms distinguished from the norms of family duty and civic duty. Moreover, creative analogical reasoning can be used to come up with likely candidates for specific norms, and MCR will be a useful means to test these norms.

6 Conclusion

As a methodology for case-based moral reasoning, CRC does have certain advantages over MCR: MCR stops after an inconsistency is resolved; it does not abstract from the particular cases to more abstract rules. It is moreover primarily a procedure of exclusion: if successful, it only tells you which judgment *not* to endorse. Moreover, it does not adequately cover large areas of moral thought. It cannot directly deal with the normal situation that a moral problem does not come in pairs of cases; neither can it account for Sartre-style dilemmas in which two possible actions are mutually exclusive without being practically inconsistent.

By contrast, CRC is more creative. Its emphasis on analogical reasoning accounts for moral problems beyond inconsistencies. These analogies often require the establishment of tentative and defeasible moral rules. Subsequently, such rules function as an epistemic framework for the understanding of moral phenomena; they also allow for simple applications in normal situations that do not require analogical reasoning or MCR. CRC also has means such as distinguishing to deal with Sartre-style dilemmas.

But MCR is still important as it complements CRC's formal elements in crucial ways. For example, casuistry only functions when one already has some kind of normative framework that determines, however tentatively, the criteria of a situation that together form what we call "the case." One already needs to have an idea of what is morally relevant in order to describe a particular case as a moral problem and to compare different cases. MCR helps with this. It makes use of our affective responses to cases as a first source of reference. Casuistry's formal elements, MCR's emotional response check and background norms make CRC a comprehensive methodology for reasoning with moral cases—one that is both descriptively accurate and normatively guiding.

References

Arras JD (1991) Getting down to cases: the revival of casuistry in bioethics. The Journal of Medicine and Philosophy 16(1):29–51
Bartha P (2010) By parallel reasoning. Oxford University Press, New York

Brewer S (1996) Exemplary reasoning: semantics, pragmatics, and the rational force of legal argument by analogy. Harvard Law Review 109(5):923–1028

Campbell R (2014) Reflective equilibrium and moral consistency reasoning. Australas J Philos 92(3):433–451

Campbell R (2017) Learning from moral inconsistency. Cognition. https://doi.org/10.1016/j.cognition.2017.05.006

Campbell R, Kumar V (2012) Moral reasoning on the ground. Ethics 122(2):273–312

Campbell R, Kumar V (2013) Pragmatic naturalism and moral objectivity. Analysis 73(3):446–455

Cath Y (2016) Reflective Equilibrium. In: Cappelen H, Gendler TS, Hawthorne J (eds) The Oxford handbook of philosophical methodology. Oxford University Press, Oxford, pp 213–230

Flanagan O (1991) Varieties of moral personality. Harvard University Press, Cambridge

Haidt J (2007) The new synthesis in moral psychology. Science 316(5827):998–1002

Horne Z, Powell D, Hummel J (2015) A single counterexample leads to moral belief revision. Cogn Sci 39(8):1950–1964. https://doi.org/10.1111/cogs.12223

Jonsen AR (1991) Casuistry as methodology in clinical ethics. Theor Med 12(4):295–307

Jonsen AR, Toulmin S (1992) The abuse of casuistry: a history of moral reasoning. University of California Press, Berkeley

Kagan S (1988) The additive fallacy. Ethics 99(1):5–31

Kagan S (2001) Thinking about cases. In: Paul EF, Miller FD, Paul J (eds) Moral knowledge. Cambridge University Press, New York, pp 44–63

Kahane G (2013) The armchair and the trolley: an argument for experimental ethics. Philos Stud 162(2):421–445

Kamm FM (1992) *Creation and Abortion: A Study in Moral and Legal Philosophy: An Essay in Moral and Legal Philosophy.* Oxford; New York: Oxford University Press

Kamm FM (2007) *Intricate Ethics: Rights, Responsibilities, and Permissible Harm.* Oxford; New York: Oxford University Press

Kamm FM (2011) "In Search of the Deep Structure of Morality." In *Conversations on Ethics*, edited by Alex Voorhoeve, 15–40. Oxford; New York: Oxford University Press

Kamm FM (2013) Theory and analogy in law and philosophy. In: Bioethical prescriptions. Oxford University Press, New York, pp 551–567

Kumar V, Campbell R (2012) On the normative significance of experimental moral psychology. Philos Psychol 25(3):311–330

Kumar V, Campbell R (2016) Honor and moral revolution. Ethical Theory Moral Pract 19(1):147–159

Lamond G (2014) "Precedent and Analogy in Legal Reasoning." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2014. http://plato.stanford.edu/archives/spr2014/entries/legal-reas-prec/

May J (2018) Regard for reason in the moral mind. Oxford University Press, New York

McMahan, Jeff. 2002. *The Ethics of Killing*. Oxford; New York: Oxford University Press

McMahan J (2013) Moral intuition. In: LaFollette H, Persson I (eds) The Blackwell guide to ethical theory, 2nd edn. Blackwell, Malden, pp 103–120. https://doi.org/10.1111/b.9780631201199.1999.00007.x

Paulo, N. (forthcoming), "A Principle of Minimal Psychological Realism for Decision Procedures in Ethics", Max Bauer & Malte Meyerhuber (eds.), Empirical research and normative theory: Transdisciplinary perspectives on two methodical traditions between separation and interdependence, DeGruyter

Paulo, N. (2015), "Casuistry as Common Law Morality," *Theoretical Medicine and Bioethics* 36: 373-89

Parfit D (2011) *On What Matters*. Vol. 1. 3 vols. Oxford; New York: Oxford University Press

Prinz, Jesse, and Shaun Nichols. 2010. "Moral Emotions." In *The Moral Psychology Handbook*, edited by John Michael Doris, 111–46. Oxford; New York: Oxford University Press

Rawls J (1951) Outline of a decision procedure for ethics. Philos Rev 60(2):177–197

Raz J (2009) "Law and Value in Adjudication." In *The Authority of Law*, 2nd ed., 180–209. Oxford University Press

Scanlon TM (2002) "Rawls on Justification." In *The Cambridge Companion to Rawls*, edited by Samuel Freeman, 139–67. Cambridge; New York: Cambridge University Press

Singer P (1972) Famine, affluence, and morality. Philos Public Aff 1(3):229–243

Spielthenner G (2014) Analogical reasoning in ethics. Ethical Theory Moral Pract 17(5):861–874

Stevens K (2018) Reasoning by precedent - between rules and analogies. Legal Theory:1–39. https://doi.org/10.1017/S1352325218000113

Sunstein CR (1993) On analogical reasoning. Harvard Law Review 106(3):741–791

Weinberg JM (2016) "Intuitions." In *The Oxford Handbook of Philosophical Methodology*, edited by Herman Cappelen, Tamar Gendler, and John Hawthorne, 287–308. Oxford; New York: Oxford University Press