

Evaluative Uncertainty and Permissible Preference

Joe Horton and Jacob Ross

Forthcoming in *The Philosophical Review*

Abstract

There has recently been an explosion of interest in rational and moral choice under *evaluative uncertainty*—uncertainty about values or reasons. However, the dominant views on such choice have at least three major problems: they are overly demanding, they are incompatible with supererogation, and they cannot be applied to agents with credence in indeterminate evaluative theories. We here propose a unified view that solves all these problems. According to this view, permissible options maximise expected utility relative to *permissible preferences*, and different kinds of permissibility for options correspond to different kinds of permissibility for preferences. Thus, *rationally* permissible options maximize expected utility relative to *rationally* permissible preferences, and *morally* permissible options maximize expected utility relative to *morally* permissible preferences. We argue that this view has more plausible implications than its rivals not only under evaluative uncertainty but also under ordinary uncertainty.

1. Introduction

There has recently been an explosion of interest in rational and moral choice under *evaluative uncertainty*—uncertainty about values or reasons. We here propose a unified view about such choice and argue that this view has major advantages over its rivals.

The order of our names above is alphabetical and does not indicate any difference in contribution. For helpful discussion of this paper, we are grateful to Zach Barnett, Alan Hájek, Anna Mahtani, Abelard Podgorski, Katie Steele, Patrick Tomlin, and two anonymous reviewers for this journal. We also benefited from presenting parts of this paper at the Australian National University, the Bled Ethics Conference, the Dianoia Institute of Philosophy, the University of Fribourg, the LMU Munich, the National University of Singapore, Princeton University, the University of Southampton, and the University of St Andrews.

According to the standard view about rational choice, rationally permissible options maximise expected *utility*, where utility is a measure of the satisfaction of your preferences. A number of philosophers have argued that, to account for the sensitivity of rational choice to evaluative uncertainty, we should instead hold that rationally permissible options maximise expected *choiceworthiness*, where choiceworthiness is a measure of value or the satisfaction of reasons.¹ It has also been argued that, to account for the sensitivity of *moral* choice to evaluative uncertainty, we should accept a moral analogue of this view.²

We think the standard view about rational choice is broadly correct, and that expected utility is also central to moral choice. According to our unified view, permissible options maximise expected utility relative to *permissible preferences*, and different kinds of permissibility for options correspond to different kinds of permissibility for preferences. Thus, *rationally* permissible options maximize expected utility relative to *rationally* permissible preferences, and *morally* permissible options maximize expected utility relative to *morally* permissible preferences.

We will refer to the choiceworthiness view as *Maximize Expected Choiceworthiness* (MEC), and to our unified view as the *Permissible Preference View* (PPV). While we present PPV as a unified account of both the rational permissibility and the moral permissibility of options, it is worth emphasizing that these two components of the view are separable. It is possible to accept the proposed account of moral permissibility while accepting an alternative account of rational permissibility, or vice versa.

We proceed as follows. In §2, we present examples that suggest that rational and moral choice are sensitive to evaluative uncertainty, and we explain how this sensitivity can be accounted for by MEC and its moral analogues. In §3, we argue that PPV can also account for

¹ See Ross 2006; Sepielli 2009; MacAskill and Ord 2020; and MacAskill, Bykvist, and Ord 2020.

² See Lockhart 2000; Zimmerman 2008, Chapter 1; Moller 2011; and Barry and Tomlin 2016.

this sensitivity. We do this by arguing that there are rational and moral constraints on what you can prefer conditional on *evaluative theories*—theories that ascribe values or reasons.³ These constraints require that your preferences conditional on evaluative theories be adequately responsive to the values or reasons these theories ascribe. In §4, we argue that PPV has major advantages over MEC and its moral analogues. These advantages include its compatibility with evaluative theories that are indeterminate, its non-demandingness, and its compatibility with supererogation. In §5, we argue that PPV has similar advantages even under ordinary, non-evaluative uncertainty. In §6, we address two potential objections, each of which concerns a way in which PPV might seem overly permissive. In §7, we conclude.

2. Evaluative Uncertainty

In this section, we present examples that suggest that rational and moral choice are sensitive to evaluative uncertainty, and we explain how this sensitivity can be accounted for by MEC and its moral analogues. We also address an influential challenge to these views known as *the problem of intertheoretic value comparison*, and we explain the relevance of a distinction between subjective probabilities and evidential probabilities.

In considering the various examples in this section, we will not be presupposing any theoretical account of rational or moral permissibility. We will instead be relying on our pretheoretic understanding of these notions. The judgments to which we appeal might initially seem incompatible with some familiar ways of theorizing about rational or moral choice, but we will later offer a theoretical framework that makes sense of them.

³ As we are understanding them, evaluative theories need not be comprehensive, in the sense of specifying all values or reasons. They can instead be partly or even largely incomplete. While an evaluative theory will typically involve a set of claims about values or reasons, a single such claim can be seen as a limiting case of an evaluative theory.

2.1. Delimiting Evaluative Uncertainty

We said that *evaluative uncertainty* is uncertainty about values or reasons. More specifically, evaluative uncertainty is first-order uncertainty about objective values or objective reasons. It is thus distinct from higher-order uncertainty, such as uncertainty about the correct view of rational or moral choice under first-order uncertainty. Similarly, by *evaluative theories* we mean first-order theories that ascribe objective values or objective reasons.

It might be that rational or moral choice is also sensitive, in some special way, to higher-order uncertainty. However, we are skeptical about this possibility, which generates various puzzles, and it is not our focus here.⁴

2.2. Evaluative Uncertainty and Rational Choice

Here is an example that suggests that rational choice is sensitive to evaluative uncertainty.

Life Choices: You must choose between a life in the real world and a life in a simulation. These lives would seem identical to you, except that the life in the simulation would involve minutely more pleasure. You have credence 0.51 in a hedonistic theory, according to which pleasure is the only thing that matters, and credence 0.49 in a pluralist theory, according to which many other things matter greatly, including real relationships and real achievements.

It seems irrational for you to choose the life in the simulation, given your credences. For you are almost as confident in the pluralist theory as in the hedonistic theory, and if the pluralist theory is correct, then there is *much* stronger reason to choose the life in the real world, whereas

⁴ For discussion of this possibility, see Macaskill 2016b and Russell n.d.

if the hedonistic theory is correct, then there is only *slightly* stronger reason to choose the life in the simulation.⁵

MEC both captures this judgment and precisifies this explanation. For this explanation entails that, given your credences, choosing the life in the simulation fails to maximize expected choiceworthiness.

2.3. *Evaluative Uncertainty and Moral Choice*

Here is an example that suggests that moral choice is sensitive to evaluative uncertainty.

Dinner Choices: You must choose between a local steak and a vegan wrap. You have credence 0.51 in an anthropocentric moral theory, according to which, because ordering the steak would help local farmers, and because animals have no moral status, there is weakly decisive moral reason to order the steak. You have credence 0.49 in a utilitarian moral theory, according to which, because cows have exactly the same moral status as humans, there is strongly decisive moral reason to order the wrap.

It seems morally reckless, and hence morally wrong, for you to order the steak, given your credences. For you are almost as confident in the utilitarian theory as in the anthropocentric theory, and if the utilitarian theory is correct, then there is *strongly* decisive moral reason to order the wrap, whereas if the anthropocentric theory is correct, then there is only *weakly* decisive moral reason to order the steak.

It has been argued that, to account for the sensitivity of moral choice to evaluative uncertainty, we should accept a moral analogue of MEC. According to the most direct moral

⁵ It might seem that this judgment conflicts with the broadly Humean view of rationality. We address this worry below, in Section 3.2.

analogue, morally permissible options maximize expected *moral choiceworthiness*, where moral choiceworthiness is a measure of moral value or the satisfaction of moral reasons.⁶ According to a natural alternative, morally permissible options *minimize* expected *objective moral wrongness*, where an option is objectively morally wrong if and only if, and to the extent that, there is decisive moral reason not to choose it.⁷ These views both capture, and in different ways explain, the judgment that it is wrong for you to order the steak, given your credences.

We will refer to the former view as *Maximize Expected Moral Choiceworthiness* (MEMC), and to the latter view as *Minimize Expected Objective Wrongness* (MEOW).

2.4. *The Problem of Intertheoretic Value Comparison*

It is sometimes claimed that it makes no sense to compare differences in value or strengths of reasons across evaluative theories, on the ground that there is no common scale on which they can be compared.⁸ If this were correct, our judgments about examples like Life Choices and Dinner Choices would be groundless.

However, we think that such intertheoretic comparisons are indeed intelligible. And this intelligibility can be defended in three steps.⁹ First, as we discuss further in sections 3.2 and 3.3, evaluative theories seem to have implications for the fittingness of preferences. Thus, if an evaluative theory holds that A is more valuable than B, or that you have more reason to choose A than to choose B, then this theory seems to imply that you should prefer A to B, with the strength of preference corresponding to the difference in value or strength of reason. Second,

⁶ A closely related view is defended in Zimmerman 2008, Chapter 1. Harman 2015 takes this to be the dominant approach to making moral choice sensitive to evaluative uncertainty, and argues against it. See also Harman 2011.

⁷ See Lockhart 2000 and Moller 2011. For a critical discussion of a closely related view, see Weatherson 2014.

⁸ This problem was introduced in Hudson 1989. For detailed discussion, see Gracely 1996; Ross 2006, 761–65; Sepielli 2013; Gustafsson and Torpman 2014; Hedden 2016; Hicks 2018; Sepielli 2019; MacAskill, Bykvist, and Ord 2020, Chapter 5; and Carr 2020.

⁹ We here develop a suggestion made in Ross 2006, 765.

preferences seem to involve not only dispositions to *choose* but also dispositions to *feel*. Thus, if you are entirely unsure whether P, then it seems that a preference for P over not-P will dispose you to feel pleased upon learning that P and disappointed upon learning that not-P, with the intensity of pleasure or disappointment corresponding to the strength of preference. Third, affective states such as pleasure and disappointment seem to be comparable both across people and across time, and so they seem to be measurable in absolute terms. Thus, it seems that your disappointment today at learning that your marriage is ending can be greater than my disappointment yesterday at learning that my toast was burning. These observations support intertheoretic comparability because, if differences in value or strengths of reasons can be measured in terms of the strengths of the preferences that they make fitting, and if the strengths of preferences can be measured in terms of the intensities of pleasure and disappointment to which they normally give rise, and if intensities of pleasure and disappointment can be measured in absolute terms, then differences in value or strengths of reasons can likewise be measured in absolute terms. And so it follows that it is possible, at least in principle, to compare differences in value or strengths of reasons across evaluative theories.

This argument might be challenged on the ground that many evaluative theories are indeterminate about the absolute magnitudes of the values or reasons they ascribe, and hence about the absolute strengths of the preferences and affective states that it is fitting to have.¹⁰ For example, utilitarianism implies that your reason to save a hundred lives is much stronger than your reason to save only one life, but, at least as it is standardly formulated, it is indeterminate about whether these reasons are both very strong in absolute terms, both very weak in absolute terms, or somewhere in between.

However, the fact that these evaluative theories are indeterminate about the absolute magnitudes of the values or reasons they ascribe, and hence about the absolute strengths of the

¹⁰ See Hedden 2016, 107–115.

preferences and affective states that it is fitting to have, shows only that they are incomplete. We typically have credence not only in the incomplete formulations of these theories, but also across a range of their different precisifications, with more credence in the precisifications that have more plausible implications about the absolute strengths of fitting attitudes. Thus, we typically have more credence in the precisification of utilitarianism on which it is fitting to be *very* pleased, in absolute terms, to discover that you saved a hundred lives rather than only one life, than we have in the precisification on which it is fitting to be only *mildly* pleased, in absolute terms, at this discovery.

2.5. Subjective Probabilities and Evidential Probabilities

MEC, MEMC, and MEOW all hold that permissible options maximize or minimize a particular expectation. We can distinguish between different versions of these views depending on the probabilities relative to which this expectation is calculated. The most natural candidates seem to be the *subjective* probabilities, which are given by your current credences, and the *evidential* or *epistemic* probabilities, which can be thought of as the credences that are most rational given the evidence available to you.

We will refer to views that focus on the subjective probabilities as *belief-relative*, and to views that focus on the evidential probabilities as *evidence-relative*. It is the belief-relative versions of MEC, MEMC, and MEOW that have received most attention in the literature on evaluative uncertainty. But the evidence-relative versions seem to us more plausible, for two reasons.

First, our judgments about examples like Life Choices and Dinner Choices are strongest when we assume that your credences in the various evaluative theories are rational given the evidence available to you. For if we instead stipulate that these credences are irrational, then it

becomes natural to hold that, rather than acting on these irrational credences, you should revise your credences and then act accordingly.

Second, the belief-relative versions have counterintuitive implications when you have high credence in absurd evaluative theories.¹¹ For example, if a bout of insanity makes you certain that you have most reason, or decisive moral reason, to torture kittens, then the belief-relative versions imply that torturing kittens is your only permissible option. The evidence-relative versions avoid such implications, assuming the credences in question are not rational given your evidence.

However, while these reasons move us to favor the evidence-relative versions of these views, we do not take them to be decisive.¹² We expect that others might still favor the belief-relative versions, or they might hold that both versions are important for answering different kinds of normative question.¹³ The arguments that follow are compatible with each of these positions.¹⁴

3. Permissible Preference

MEC and its moral analogues are not the only way to account for the sensitivity of rational and moral choice to evaluative uncertainty. In this section, we argue that the Permissible Preference View can also account for this sensitivity. We do this by arguing that there are rational and moral constraints on what you can prefer conditional on evaluative theories.

¹¹ This criticism is developed in detail in Harman 2015. See also Hedden 2016, 125–126; and Staffel 2021. For related discussion concerning the relationship between evaluative belief and morally worthy choice, see Arpaly 2002; Markovits 2010; Harman 2011; Arpaly 2015; Sliwa 2016; Johnson King 2020a; and Johnson King 2020b.

¹² For some other considerations that support the evidence-relative versions over the belief-relative versions, see Zimmerman 2008, Chapter 1; and Kiesewetter 2016.

¹³ For relevant discussion, see Parfit 2011, Chapter 7; Pittard and Worsnip 2017; and Barnett 2021.

¹⁴ We are grateful to an anonymous reviewer for helpful discussion here.

3.1. *The Permissible Preference View*

PPV holds that permissible options maximise expected utility relative to permissible preferences, and that different kinds of permissibility for options correspond to different kinds of permissibility for preferences. Thus, *rationally* permissible options maximize expected utility relative to *rationally* permissible preferences, and *morally* permissible options maximize expected utility relative to *morally* permissible preferences.

We are using the term *options* broadly. An option can be an individual action, but it can also be a set of simultaneous actions or a sequence of actions extended over time. Thus, where an option consists of a set or sequence of actions, PPV holds that the option is permissible just in case the entire set or sequence maximizes expected utility relative to *a single permissible set* of preferences.

We can again distinguish between different versions of PPV depending on the probabilities relative to which expected utility is calculated. The belief-relative version uses the subjective probabilities. The evidence-relative version uses the evidential probabilities.

PPV has determinate implications only when combined with an account of the rational and moral permissibility of preferences. We will now argue that there are rational and moral constraints on what you can prefer conditional on evaluative theories, and that when PPV is combined with these constraints, it accounts for the sensitivity of rational and moral choice to evaluative uncertainty.

3.2. *The Rational Conditional Preference Principle*

Let T1 be an evaluative theory according to which you have more reason to choose A than to choose B. Suppose you are certain in T1, and yet you prefer B to A. It seems you are rationally criticisable, for you prefer an option that you are certain you have less reason to choose.

Let T2 be an evaluative theory according to which you have *much* more reason to choose A than to choose B. Suppose you are certain in T2, and you prefer A to B, but this preference is weak. It again seems you are rationally criticisable, for you only *weakly* prefer an option that you are certain you have *much* more reason to choose.

If there are rational constraints on what you can prefer while being certain in T1 and T2, then it seems there will be corresponding rational constraints on what you can prefer *conditional on* T1 and T2. For it seems that your preferences upon learning T1 or T2 should correspond to your prior preferences conditional on T1 or T2, respectively. So, it seems that you are rationally required to prefer A to B conditional on T1, and that you are rationally required to *strongly* prefer A to B conditional on T2.

These observations support the following general principle.

The Rational Conditional Preference Principle (RCPP): For any evaluative theory T and any preference P, having P conditional on T is rationally permissible only if P is fitting relative to, or does not conflict with, the values or reasons ascribed by T.¹⁵

It might seem that this principle is incompatible with the broadly Humean view that there are no *substantive* rational constraints on preferences.¹⁶ However, even the Humean view allows for *coherence* constraints on preferences, such as a constraint requiring that preferences

¹⁵ This principle applies when your uncertainty is purely evaluative. If your uncertainty is also non-evaluative (perhaps because you are uncertain about the consequences of your actions), then the principle needs to be reformulated so as to range not over evaluative theories but rather over *maximally relevantly specific situations*, where each such situation is a complete specification of all the facts, both evaluative and non-evaluative, that are not under your control and that are relevant to the evaluation of your options. For discussion of such mixed uncertainty, and the problems it might present for standard formulations of MEC, MEMC, and MEOU, see Podgorski 2020 and Robinson 2022.

¹⁶ We are grateful to an anonymous reviewer for encouraging us to address this issue.

be transitive. It can also allow for coherence constraints on the relationship between preferences and other attitudes, such as a prohibition on intending to choose A over B while preferring B to A. RCPP is best seen as another kind of coherence constraint. More specifically, RCPP is best seen as a conditionalized version of an *anti-akratic principle*, according to which it is irrational to have a preference while believing that this preference is incompatible with your reasons. For if your preferences violate RCPP, then there are evaluative theories such that, conditional on these theories, you are akratic.¹⁷

Return to Life Choices. RCPP implies that you are rationally required to *weakly* prefer the life in the simulation conditional on the hedonistic theory and to *strongly* prefer the life in the real world conditional on the pluralist theory. If you have these preferences, then because your credences over the hedonistic theory and the pluralist theory are divided roughly evenly, choosing the simulation will fail to maximize expected utility. So, given RCPP, the belief-relative version of PPV implies that it is irrational for you to choose the life in the simulation. And if we assume that your credences are rational given the evidence available to you, then the evidence-relative version of PPV implies the same.

More generally, in a range of contexts, the combination of PPV and RCPP mimics MEC, in a way that allows it to account for the sensitivity of rational choice to evaluative uncertainty.

¹⁷ Many philosophers hold that an anti-akratic principle applies to other attitudes as well. Thus, it is widely held that it is irrational to believe something while believing that you lack sufficient reason to believe it, to fear something while believing that you lack sufficient reason to fear it, and so on, and, more generally, that it is irrational to have an attitude while believing that you lack sufficient reason to have it. Scanlon 1998 (page 25) even suggests that all irrationality can ultimately be understood in terms of such akratic attitudes. For further discussion, see Broome 2013a; Broome 2013b, 170–75 and 291–94; and Way and Whiting 2016.

3.3. *The Moral Conditional Preference Principle*

Let T3 be a moral theory according to which, given a choice between only A and B, you have decisive moral reason to choose A. Suppose you are certain in T3, and yet you prefer B to A. It seems you are morally criticisable, for you prefer an option that you are certain is morally wrong.

Let T4 be a moral theory according to which, given a choice between only A and B, you have *strongly* decisive moral reason to choose A. Suppose you are certain in T4, and you prefer A to B, but this preference is weak. It again seems you are morally criticisable, for you only *weakly* disprefer an option that you are certain is *seriously* morally wrong.

If there are moral constraints on what you can prefer while being certain in T3 and T4, then it seems there will be corresponding moral constraints on what you can prefer *conditional on* T3 and T4. For it seems that your preferences upon learning T3 or T4 should correspond to your prior preferences conditional on T3 or T4, respectively. So, it seems that you are morally required to prefer A to B conditional on T3, and that you are morally required to *strongly* prefer A to B conditional on T4.

These observations support the following general principle.

The Moral Conditional Preference Principle (MCP): For any evaluative theory T and any preference P, having P conditional on T is *morally* permissible only if P is fitting relative to, or does not conflict with, the *decisive moral* reasons ascribed by T.¹⁸

Return to Dinner Choices. According to the utilitarian theory, cows have exactly the same moral status as humans, and so there is strongly decisive moral reason to order the wrap. MCP thus implies that you are morally required to *strongly* prefer the wrap conditional on the

¹⁸ This principle is subject to the same caveat as applied to RCPP in footnote 15.

utilitarian theory. According to the anthropocentric theory, because ordering the steak would support local farmers, and because animals have no moral status, you have weakly decisive moral reason to order the steak. If this were all this theory held, then this theory would be compatible with any sufficiently strong preference for the steak. But this theory presumably also holds that it would be morally wrong to choose the steak over the wrap when doing so would involve a non-trivial moral cost. For example, if choosing the steak over the wrap would result in your breaking a promise, then this theory will imply that you should not do so. MCPP thus implies that you are morally required to have a relatively *weak* preference for the steak conditional on the anthropocentric theory.

If you have these preferences, then because your credences over the utilitarian theory and the anthropocentric theory are divided roughly evenly, ordering the steak will fail to maximize expected utility. So, given MCPP, the belief-relative version of PPV implies that it is morally wrong for you to order the steak. And if we assume that your credences are rational given the evidence available to you, then the evidence-relative version of PPV implies the same.

More generally, in a range of contexts, the combination of PPV and MCPP mimics MEMC and MEOW, in a way that allows it to account for the sensitivity of moral choice to evaluative uncertainty.

3.4. The Permissibility of Unconditional Preferences

RCPP and MCPP directly constrain the permissibility of preferences conditional on evaluative theories. But in so doing they also indirectly constrain the permissibility of *unconditional* preferences. For completeness, we will now precisify this indirect constraint.¹⁹

¹⁹ We are grateful to an anonymous reviewer for helpful discussion here.

For any agent S and any unconditional preference P regarding a pair of outcomes O1 and O2, P is permissible for S only if, for some permissible set of conditional preferences regarding O1 and O2 that is complete relative to the relevant probability distribution over evaluative theories, the strength of P is the sum of the strengths of these conditional preferences when each is weighted by the probability of the evaluative theory on which it is conditional.²⁰ P is *rationally* permissible for S only if this constraint is satisfied with respect to *rationally* permissible conditional preferences, and P is *morally* permissible for S only if this constraint is satisfied with respect to *morally* permissible conditional preferences. We can again distinguish between belief-relative and evidence-relative versions of this constraint, with the former using the subjective probabilities and the latter using the evidential probabilities.

If the only fundamental constraints on permissible preferences are RCPP and MCPP together with the standard coherence requirements, then the preceding constraint will be not only a necessary but also a sufficient condition for the permissibility of unconditional preferences. We return to the question of whether there are further constraints on the permissibility of preferences in Section 6.2.

4. The Advantages of PPV

In this section, we argue that PPV has major advantages over MEC and its moral analogues. These advantages include its compatibility with indeterminate evaluative theories, its non-demandingness, and its compatibility with supererogation.

²⁰ Note that this constraint on the permissibility of unconditional preferences presupposes a solution to the problem of intertheoretic value comparison—either the solution we provided in Section 2.4, or some other solution. For if there is no way to compare differences in value or strengths of reasons across evaluative theories, then there will be no way to compare the strengths of the preferences that are permissible conditional on different evaluative theories. And so there will be no way to sum the strengths of these preferences.

4.1. Indeterminacy

We argued earlier that when evaluative theories fail to ascribe absolute magnitudes to values or reasons, they can typically be precisified so that they do. However, there are some evaluative theories that cannot be so precisified, because it is essential to them that values or reasons do not have absolute magnitudes that can be precisely quantified. These include theories on which options can be ranked only ordinally, and theories on which options can be given only a partial ordering.²¹

These *essentially indeterminate* theories are a serious problem for MEC and its moral analogues. For when you have any credence in any of these theories, some of the inputs needed to calculate the relevant expected value—such as choiceworthiness, moral choiceworthiness, or objective moral wrongness—are undefined, and so these views imply that the rational and moral permissibility of your options is itself undefined. To see this, consider an example.

Afternoon Alternatives: You must choose between an afternoon at the beach and an afternoon at a philosophy lecture. Your credence is evenly divided between a fully determinate hedonistic theory and an essentially indeterminate pluralist theory. According to the hedonistic theory, going to the beach is much better in absolute terms. According to the pluralist theory, going to the lecture is better, but your options can be ranked only ordinally.

Since the pluralist theory holds that value differences between options cannot be quantified, there is no way of comparing the difference in value ascribed by the pluralist theory with the

²¹ For discussion of these theories and the problems they present for views like MEC, MEMC, and MEOW, see MacAskill 2016a; Tarsney 2019; Cotton-Barratt, MacAskill, and Ord 2020; MacAskill, Bykvist, and Ord 2020, Chapter 3, Chapter 4, and Chapter 6; Tarsney 2021; Carr 2022; and Cotton-Barratt and Greaves 2024. For discussion of a closely related problem, see MacAskill 2013.

difference in value ascribed by the hedonistic theory. MEC thus implies that the rational permissibility of your options is undefined.

This problem is especially serious if, as seems plausible, it is irrational not to have at least a smidgeon of credence in at least some essentially indeterminate theories. For then MEC and its moral analogues imply that rational and moral permissibility is undefined for every rational agent.

However, essentially indeterminate theories are not a problem for RCPP and MCPP, and hence they are not a problem for PPV. For even when the values or reasons ascribed by a theory are indeterminate, your preferences can still be consistent or inconsistent with these values or reasons, and hence your preferences conditional on this theory can still be permissible or impermissible. For example, since the pluralist theory holds that going to the lecture is better and that your options can be ranked only ordinally, RCPP forbids all and only those sets of preferences that fail to rank going to the lecture above going to the beach.

4.2. Demandingness

The moral analogues of MEC face another problem. Since they ignore your credences about non-moral reasons, and in particular about prudential reasons, they are counterintuitively demanding.²² To see this, consider another example.

Sacrifice: You must decide whether to sacrifice your life to save two strangers. The overwhelming majority of your credence is in a deontological moral theory, according to which not saving the strangers is objectively morally permissible. But you have a

²² This kind of problem is pressed against moral analogues of MEC in Weatherson 2002; Barry and Tomlin 2016; Hedden 2016, 115–116; Rosenthal 2021; and Sung n.d.

sliver of credence in a utilitarian moral theory, according to which not saving the strangers is objectively morally wrong.

Since you are certain that there is at least as much moral reason to save the strangers as to not save them, and you have a sliver of credence that not saving them is objectively morally wrong, saving them both maximizes expected moral choiceworthiness and minimizes expected objective moral wrongness. So, MEMC and MEOW both imply that you are morally required to save the strangers, even though you are almost certain that not saving them is objectively morally permissible.

There is a dilemma here for all moral analogues of MEC. If they ignore prudential reasons, then they will imply that options like sacrificing yourself to save the strangers are morally *required*, because these will be the only options that maximize or minimize the relevant expectation. But if they instead take prudential reasons into account, and they give them sufficient weight to avoid overly demanding implications, then they will imply that, in Sacrifice or in similar examples, such options are morally *wrong*, because these options will fail to maximize or minimize the relevant expectation. These views seem unable to capture the judgment that such options are neither morally required nor morally wrong.

PPV combines with MCPP in a way that avoids this dilemma. For while the decisive moral reasons ascribed by the utilitarian theory are compatible only with a strong preference for saving the strangers, the decisive moral reasons ascribed by the deontological theory are compatible with a wide range of preferences, including both a strong preference for saving the strangers and a strong preference for not saving them. MCPP thus implies that, given your credences, there are *some* morally permissible sets of preferences relative to which saving the strangers maximizes expected utility, and *other* morally permissible sets of preferences relative

to which *not* saving the strangers maximizes expected utility. So, PPV implies that both options are morally permissible.

4.3. *Supererogation*

The moral analogues of MEC also face a problem regarding supererogation. To see this, consider another example.

Stamp: You have acquired tickets to the World Series and are considering sending them to an acquaintance who loves baseball. However, to send the tickets in time, you would need to take a postage stamp from the supplies at your work. And while you are almost certain that taking the stamp to send the tickets is both morally permissible and supererogatory, you have a sliver of credence that taking the stamp is very slightly morally wrong.

MEMC and MEOW both have the wrong implication here. MEOW implies that taking the stamp to send the tickets is morally *wrong*, even though you are almost certain that doing so is both morally permissible and supererogatory. This is because you have a sliver of credence that taking the stamp to send the tickets is very slightly objectively morally wrong, whereas you are certain that not doing so is objectively morally permissible. And MEMC implies that, when you have sufficiently low credence that it is wrong to take the stamp to send the tickets, then doing so is morally *required*, even though you are certain that not doing so is objectively morally permissible. This is because MEMC is sensitive not just to the moral reasons that make options objectively morally wrong, but also to the moral reasons that make options supererogatory.

This illustrates another dilemma for all moral analogues of MEC. If they ignore moral reasons for supererogation, then they will imply that options like taking the stamp to send the tickets are morally *wrong*, because these options will fail to maximize or minimize the relevant expectation. But if instead they take these moral reasons into account, then they will imply that such options are morally *required*, because these will be the only options that maximize or minimize the relevant expectation. These views seem unable to capture the judgment that such options are morally permissible but not morally required.

PPV again combines with MCPP in a way that avoids this dilemma. If taking the stamp to send the tickets is supererogatory, then the decisive moral reasons are compatible with a range of preferences, including both a strong preference for doing so and a strong preference against doing so. MCPP thus implies that, given your credences, there are *some* morally permissible sets of preferences relative to which taking the stamp to send the tickets maximizes expected utility, and *other* morally permissible sets of preferences relative to which not doing so maximizes expected utility. So, PPV implies that taking the stamp to send the tickets is morally permissible but not morally required.

5. PPV without Evaluative Uncertainty

We have argued that, when applied under evaluative uncertainty, PPV has major advantages over MEC and its moral analogues. In this section, we argue that PPV has the same advantages even when applied under ordinary, non-evaluative uncertainty. We also argue that PPV has two further advantages relating to *discretionary beneficence* and *robust permissibility*.

5.1. Indeterminacy without Evaluative Uncertainty

We have seen that expected value views, like MEC and its moral analogues, can be applied under evaluative uncertainty only when all the evaluative theories in which you have credence

are fully determinate. But indeterminate evaluative theories can create similar problems for these views even under ordinary uncertainty. To see this, consider an example.

Art Exhibit: You are deciding between going to the museum to see a new art exhibit and staying home to play video games. However, your credence that the museum is open today is only 0.5. You are certain both that seeing the art exhibit is more valuable than playing video games, and that playing video games is more valuable than going to a closed museum. But you are also certain that these alternatives can be ranked only ordinally.

Since you have credence 0.5 that the museum is open, going to the museum has greater expected value than staying home if and only if the difference in value between seeing the art exhibit and playing video games is greater than the difference in value between playing video games and going to a closed museum. But you are certain that these alternatives can be ranked only ordinally, and hence that these value differences cannot be compared. Thus, the expected value of your options is undefined, and so expected value views imply that the permissibility of your options is likewise undefined.

PPV avoids this indeterminacy. If your options can be ranked only ordinally, then your preference for seeing the art exhibit over playing video games can fittingly be either stronger or weaker than your preference for playing video games over going to a closed museum. Thus, both options maximize expected utility relative to permissible preferences, and so PPV implies that both options are permissible.

5.2. Demandingness without Evaluative Uncertainty

We have also seen that, when applied under evaluative uncertainty, expected value views face a dilemma regarding the moral status of self-sacrifice. Again, the same problem arises for these views even under ordinary uncertainty.²³ To see this, consider another example.

Grenade: You are confined to a small area where there is a live grenade. The only way to prevent the grenade from killing you is to lob it over a wall. However, you have a small amount of credence that there is an innocent person on the other side of the wall who will be killed if you lob the grenade. You are certain that lobbing the grenade is objectively morally wrong if there is someone on the other side of the wall, and objectively morally permissible otherwise.

If expected value views ignore prudential considerations, then they will imply that lobbing the grenade is morally *wrong*, even when your credence that there is someone on the other side of the wall is minuscule. But if instead they take prudential considerations into account, then they will imply that, when this credence is sufficiently low, lobbing the grenade is morally *required*. These views seem unable to capture the judgment that, when this credence is sufficiently low, lobbing the grenade is neither morally wrong nor morally required.

By contrast, PPV can easily capture this judgment, along with the judgment that you must refrain from lobbing the grenade when your credence that there is someone on the other side of the wall is sufficiently high. To explain these judgments, we need assume only that you are morally required to have a very strong preference against killing innocent people, and that your preference for your own survival can permissibly range anywhere from zero to quite strong.

²³ Lazar 2017 discusses this kind of problem for MEMC and MEOW. See also Zimmerman 2008, 55–56.

5.3. Supererogation without Evaluative Uncertainty

We have also seen that, when applied under evaluative uncertainty, expected value views face a dilemma with respect to supererogation. Again, the same problem arises for these views even under ordinary uncertainty. To see this, consider the following example.

Stamp Collector: You are again deciding whether to send your World Series tickets to a baseball-loving acquaintance. This time, there is only one available stamp that you can use to mail the tickets, and you are unsure whether it is an ordinary stamp belonging to you or an irreplaceable stamp belonging to your stamp-collecting colleague. You are certain that, if the stamp belongs to your colleague, then taking the stamp to send the tickets is objectively morally wrong, but if the stamp belongs to you, then doing so is supererogatory.

If expected value views ignore reasons for supererogation, then they will imply that taking the stamp to send the tickets is morally *wrong*, even when your credence that the stamp belongs to your colleague is extremely low. But if they instead take these reasons into account, then they will imply that, when this credence is sufficiently low, taking the stamp to send the tickets is morally *required*. These views seem unable to capture the judgment that, when this credence is sufficiently low, taking the stamp to send the tickets is neither morally wrong nor morally required.

By contrast, PPV can easily capture this judgment, along with the judgment that you must refrain from taking the stamp when your credence that it belongs to your colleague is sufficiently high. To explain these judgments, we need assume only that you are morally required to have a strong preference against taking an irreplaceable stamp belonging to your

colleague, and that, conditional on the stamp belonging to you, your preference for sending the tickets can permissibly range anywhere from zero to moderately strong.

5.4. Discretionary Beneficence

Most people think that we have considerable latitude with respect to how we choose to benefit the world. We can permissibly choose to devote our resources to feeding the hungry, curing the sick, protecting the environment, promoting peace, and so on. However, most people also think that, other things being equal, we should choose to promote any particular cause in a more effective way rather than in a less effective way.²⁴ Thus, if we are choosing between feeding more people and feeding fewer people, then, other things being equal, it would be wrong to choose to feed fewer people. And if we are choosing between a greater chance of averting a tragedy and a lower chance of averting this same tragedy, then, other things being equal, it would be wrong choose the lower chance.

But expected value views cannot account for both these judgments simultaneously. To see this, consider the following example.

Donation: You have two million dollars to donate and four charities to choose among. For legal reasons, each of these charities can accept only one million dollars from any single donor. Charity A and Charity B would both use the money to reduce the risk of the same civil war, but while A would use it to reduce this risk by 10%, B would use it to reduce this risk by 11%. Charity C and Charity D would both use the money to reduce the risk of the same species going extinct, but while C would use it to reduce this risk by 10%, D would use it to reduce this risk by 11%.

²⁴ For discussion and defense of this judgment, see MacAskill 2015; Pummer 2016; Horton 2017; and Sinclair 2018.

If we take seriously the idea that we have considerable latitude with respect to how we choose to benefit the world, then we should allow that giving all the money to A and B could be morally permissible, that giving all the money to C and D could be morally permissible, and that giving all the money to B and D could be morally permissible. And if we take seriously the idea that we should choose to promote any particular cause in a more effective way rather than in a less effective way, then we should maintain that it would be morally wrong to choose A over B, or to choose C over D, and hence that it would be morally wrong to give all the money to A and C.

But expected value views cannot account for all these judgments. For these views will imply that it is morally wrong to choose A over B, and to choose C over D, only if they assign disvalue to the war and to the extinction. But if the war and the extinction both have disvalue, and if the expected disvalue of the war is *at least as great as* the expected disvalue of the extinction, then giving to B and D will have greater expected value than giving to C and D, and so expected value views will imply that giving to C and D is morally wrong. And if, instead, the expected disvalue of the war is *less than* the expected disvalue of the extinction, then giving to C and D will have greater expected value than giving to A and B, and so expected value views will imply that giving to A and B is morally wrong.

PPV again gets the right results in this example. We need assume only that, other things being equal, you are morally required to prefer preventing the war to allowing it, you are morally required to prefer preventing the extinction to allowing it, and it is morally permissible for you to care somewhat more about the war than about the extinction or vice versa.

5.5. Robust Permissibility

Many of these problems for expected value views are instances of a more general problem. It seems common for two options to be not only permissible but also *robustly* permissible, in the

sense that both would remain permissible if either were slightly improved. This is illustrated by all the examples considered above. For example, you have robustly permissible options in Grenade, because lobbing the grenade and not lobbing the grenade are both permissible when your credence that there is someone on the other side of the wall is very low, and both options would remain permissible if this credence were slightly lower or higher. And you have robustly permissible options in Stamp Collector, because taking the stamp and not taking the stamp are both permissible when your credence that the stamp belongs to your colleague is very low, and both would remain permissible if this credence were slightly lower or higher.

However, in general, expected value views have trouble accounting for robust permissibility. For they hold that two options are both permissible only if both have exactly the same expected value. And if two options have exactly the same expected value, then they would have different expected values if one of them were slightly improved.

PPV captures robust permissibility straightforwardly. For two options to be permissible according to this view, they need not be ranked equally by any permissible set of preferences. It could instead be that one option is ranked higher by *one* permissible set of preferences, and the other option is ranked higher by *another* permissible set of preferences. And this condition can be robust, in the sense that it would remain satisfied even if either option were slightly improved.

5.6. *Why Not Satisfice?*

A natural way to revise expected value views to capture robust permissibility, and hence to do better with respect to demandingness, supererogation, and discretionary beneficence, is by having them require not that you *maximize* the relevant expectation, but rather that you *come sufficiently close to maximizing* this expectation.²⁵ We will refer to this view as *Satisfice*.

²⁵ This kind of view is advocated in Slote and Pettit 1984; and Slote 1985, 45–47.

Satisfice has received significant attention as a response to the kinds of problems we have been discussing.²⁶ But it does not have the same explanatory power as PPV. For example, as we have seen, one of the virtues of PPV is that it can allow for the right kind of discretionary beneficence without allowing for the wrong kind. That is, given plausible assumptions about permissible preferences, it will allow you to choose *which cause to promote*, but it will not allow you to choose promoting a cause in a *less effective way* over promoting the same cause in a *more effective way*, other things being equal. Satisfice does not have this virtue, as it allows for both kinds of discretion equally.

This failure to account for our judgments about discretionary beneficence is a symptom of a deeper problem. Satisfice allows for what we might call *senseless suboptimality*. For example, suppose you must choose the probability that an innocent person is tortured and killed, and you can choose any probability from 0% to 100% inclusive. Satisfice implies that it is permissible for you to choose a probability greater than zero, so long as this probability is *close enough* to zero. But this is highly counterintuitive—inflicting a risk of torture and death on an innocent person for no reason seems clearly wrong.²⁷ PPV accounts for this judgment, given the plausible assumption that you are morally required to prefer that innocent people not be tortured and killed.

Moreover, Satisfice faces a dilemma. For the greater the margin by which it allows you to fall short of optimality, the worse it does with respect to senseless suboptimality. But the smaller this margin, the worse it does with respect to demandingness, supererogation, discretionary beneficence, and robust permissibility.

²⁶ See the discussions of supererogation and discretionary beneficence in Slote and Pettit 1984.

²⁷ For detailed criticism of Satisfice, see Bradley 2006.

6. Problems of Permissiveness

In this section, we explain and respond to two potential objections, each of which concerns a way in which PPV might seem overly permissive.

6.1. *Choosing Against Preference*

For any two options A and B, it seems irrational for you to choose A over B while preferring B to A, regardless of your evaluative credences. For example, it seems irrational for you to choose apples over oranges while preferring oranges to apples. But PPV cannot account for this. For you could prefer oranges to apples while being certain in an evaluative theory according to which a preference for apples over oranges and a preference for oranges over apples are both perfectly fitting. And PPV would then imply that choosing apples over oranges is rationally permissible.

However, we can capture the judgment that it is irrational to choose against your preferences by combining PPV with a widely accepted coherence constraint on rational choice. According to this constraint, it is irrational to choose an option while having credences and preferences relative to which this option fails to maximize expected utility.²⁸ We will refer to this constraint as *Wide-Scope Rational Coherence* (WSRC).

WSRC captures the judgment that it is irrational for you to choose A over B while preferring B to A, regardless of your evaluative credences. And it is fully compatible with PPV. For it can be true both that an option is rationally permissible and that it is irrational to choose this option while having particular credences and preferences. For example, it can be rationally permissible to put sugar in your tea, but irrational to do so while believing that the sugar is poisonous.

²⁸ This coherence constraint is widely accepted in the literature on decision theory.

Furthermore, WSRC might actually *support* PPV, at least insofar as PPV is applied to the rational permissibility of options. For it seems plausible that an option is rationally permissible only if it can be chosen without irrationality. And if it is irrational to choose an option while having credences and preferences relative to which this option fails to maximize expected utility, then an option can be chosen without irrationality only if there are rationally permissible credences and preferences relative to which it maximizes expected utility.²⁹

6.2. Infectious Permissiveness

There are some evaluative theories that are *absolutely permissive*, in the sense that they do not rule out any preferences at all. For example, *nihilism* is plausibly a theory of this kind. PPV together with the conditional preference principles might seem to imply that, if the probability of such a theory is greater than zero, then every option is permissible. For if the probability of an absolutely permissive theory is greater than zero, and if, conditional on such a theory, there is no limit to how strongly you can permissibly prefer any given option to any other, then it seems to follow that, for any given option, there are permissible preferences relative to which this option maximizes expected utility.³⁰

²⁹ As an anonymous reviewer pointed out to us, the inference from WSRC to PPV relies on an assumption about ‘deontic detachment’. For it assumes that, if S is rationally required to be such that (if P, then Q), and S is rationally required to be such that P, then S is rationally required to be such that Q. If this assumption holds, then we can infer that if you are rationally required to be such that (if you have credences and preferences relative to which an option O does not maximize expected utility, then you do not choose O), and if you are also rationally required to have credences and preferences relative to which O does not maximize expected utility, then you are rationally required to not choose O. Moreover, it seems to us that there is reason to accept this assumption about deontic detachment. For, plausibly, to say that S is rationally required to be such that P is to say that, necessarily, if S is fully rational, then P. And given this interpretation of the rational requirement relation, the assumption about deontic detachment follows from the standard laws of modal logic.

³⁰ There are also evaluative theories that are *fanatical*, in the sense that they ascribe infinite values or infinitely strong reasons. These theories pose a related problem for PPV, as well as for MEC and its moral analogues. For discussion of these theories, see Ross 2006, 765–767; and MacAskill, Bykvist, and Ord 2020, 150–155.

To see this problem at its starkest, consider a simplified example.

Life or Death: You must decide whether to kill an innocent person. You are almost certain in a commonsense evaluative theory, according to which you have both most reason and decisive moral reason not to kill this person. But you also have a shred of credence in an absolutely permissive theory.

PPV, together with RCPP and MCPP, might seem to have the implausible implication that it is both rationally and morally permissible for you to kill the innocent person. For while the reasons not to kill this person ascribed by the commonsense theory will be strong, they will not be of unlimited strength, as they will be weaker than certain other reasons, such as the reason not to kill multiple innocent people. And so a sufficiently extreme preference for killing this person conditional on the absolutely permissive theory will outweigh the preference against killing this person that is required conditional on the commonsense theory, in spite of the vast difference between the probabilities of these theories.

This objection to PPV rests on two assumptions. First, it rests on the assumption that, besides RCPP and MCPP, there are no other rational or moral constraints on preferences. For if instead there are other constraints on preferences, then these constraints might rule out the preferences relative to which killing the innocent person maximizes expected utility. Second, it rests on the assumption that, in cases like *Life or Death*, it could be most rational to assign non-zero probability to an absolutely permissive theory. For if instead such a probability assignment is irrational, then at least the evidence-relative version of PPV avoids the implication that killing the innocent person is permissible.³¹

³¹ These two assumptions also underlie the problem of fanaticism mentioned in footnote 30.

It is worth noting that these two assumptions make trouble for everyone, regardless of whether PPV is accepted. For these assumptions together imply that, in Life or Death, it is rationally and morally permissible for you to *overall prefer* killing the innocent person. Since that is implausible, everyone must reject at least one of these assumptions. And as we will now show, our view is compatible with rejecting either.

We can begin by focusing on the assumption that, besides RCPP and MCPP, there are no other rational or moral constraints on preferences. Since we formulated RCPP and MCPP as *necessary* conditions for the permissibility of conditional preferences rather than as *sufficient* conditions, they are compatible with the rejection of this assumption. And there are at least three ways in which we could reject it. First, we could hold that there are additional constraints on preferences *conditional on evaluative theories*. We could thus hold that, in Life or Death, even conditional on the absolutely permissive theory, it is rationally and morally impermissible for you to have a fanatical preference for killing the innocent person. Second, we could hold that, while RCPP and MCPP might be the only constraints on preferences conditional on evaluative theories, there are also irreducible constraints on *unconditional preferences*. We could thus hold that, in Life or Death, an overall preference for killing the innocent person is ruled out by these irreducible constraints.³² Third, we could hold that there are psychological limits on what preferences are *possible* for you and that these give rise to constraints on what preferences are *permissible* for you. For it seems plausible that, just as the *actions* that are permissible for you must lie within the range of actions that are possible for you, so too the *preferences* that are permissible for you must lie within the range of preferences that are

³² If there are irreducible constraints on unconditional preferences, then while they will not rule out any *particular* conditional preferences, they will rule out certain *combinations* of conditional preferences. Thus, if there is an irreducible constraint prohibiting an unconditional preference for killing the innocent person in Life or Death, then it rules out a conditional preference for killing this person unless this preference is more than offset by other conditional preferences against killing this person.

possible for you. And we could hold that the extreme conditional preferences that give rise to the problem of permissiveness lie outside of this range. We could thus hold that, in Life or Death, a preference with the kind of stupendous strength needed to give rise to the problem is not psychologically possible for you, and so it cannot count as permissible for you.

We can next turn to the assumption that, in cases like Life or Death, it could be most rational to assign non-zero probability to an absolutely permissive theory. There seem to be two main ways of rejecting this assumption. First, we could hold that, for purely epistemic reasons, extremely permissive theories are implausible, and the more permissive they are, the more implausible they are, so that at the limit their plausibility approaches zero. Second, we could hold that absolutely permissive theories should be ruled out on practical grounds, because of their inadequacy as a basis for deciding what to do. We could thus hold that, at least in the context of practical reasoning, absolutely permissive theories should be assigned a probability of zero.³³

So, in summary, the existence of absolutely permissive theories shows that either there are constraints on permissible preferences other than RCPP and MCPP, or else certain kinds of evaluative theories must be assigned a probability of zero. We have outlined a number of ways of developing each of these positions, but we leave a more complete investigation for future work.

7. Conclusion

It is common to think of the history of ethics as dominated by three traditions: *consequentialism*, *deontology*, and *virtue ethics*.³⁴ It is also commonly held that virtue ethics is irresolvably unsystematic and imprecise. However, the Permissible Preference View can be

³³ We are grateful to an anonymous reviewer for suggesting a view of this kind.

³⁴ For some important discussions of virtue ethics, see Anscombe 1958; Foot 1978; and Hursthouse 1999.

seen as a way of precisifying virtue ethics. For if we identify *permissible* preferences with *virtuous* preferences, then PPV will be equivalent to the view that permissible options are options that maximize expected utility relative to virtuous preferences. And if we understand a *virtuous agent* as a *rational agent with virtuous preferences*, then PPV will be equivalent to the view that permissible options are options that could be chosen by a virtuous agent.

Moreover, PPV could be extended to provide an account of the specific virtues. For example, we could understand *just* actions as actions that maximize expected utility relative to *just* preferences, and we could understand *temperate* actions as actions that maximize expected utility relative to *temperate* preferences, and so on for the other virtues. And we could understand *fully virtuous* actions as actions that maximize expected utility relative to preferences that jointly satisfy all the virtues.³⁵

Of course, in order to accept PPV, we do not need to be virtue ethicists. Nor do we need to take constraints on preferences as primitive. The consequentialist, for example, could take facts about the values of outcomes as primitive, and she could use MCPP to derive constraints on preferences from her views about the values of outcomes. And the deontologist could take facts about moral obligations and moral reasons as primitive, and she could use MCPP to derive constraints on preferences from her views about moral obligations and moral reasons.

However, given that the debate among these three philosophical positions has persisted without resolution for centuries, it may be rational to be uncertain as to which, if any, of these positions is correct. Thus, it may be rational to divide our credence among theories of all three types, and perhaps among some other theories as well. But so long as we are able to derive

³⁵ The resulting view could be described as a kind of *decision-theoretic* virtue ethics. However, it should not be confused with the view defended in Wedgwood n.d. For these views are not virtue-theoretic in the same sense. While the view proposed here is *agent-focused*, in that it understands the normative features of actions in terms of the normative features of the preferences that could motivate these actions, the view defended by Wedgwood is, in his words, “*act-focused*”, in that “it focuses on the virtue-properties—such as justice or prudence or beneficence—that are instantiated by acts themselves (rather than by the agents of those acts).”

implications about fitting preferences from each of these theories, PPV together with MCPP will provide practical guidance.

Perhaps this debate will one day be resolved, and we will come to have full confidence in a single evaluative theory. Even then, however, there will remain ordinary uncertainty about the consequences of our actions. As we have seen, PPV has an important role to play in helping us navigate this ordinary uncertainty. And so, regardless of which evaluative theory prevails, PPV will continue to provide guidance in applying this theory in an uncertain world.

Bibliography

- Anscombe, G. E. M. 1958. "Modern Moral Philosophy." *Philosophy* 33: 1–19.
- Arpaly, Nomy. 2002. "Moral Worth." *The Journal of Philosophy* 99: 223–45.
- Arpaly, Nomy. 2015. "Huckleberry Finn Revisited: Inverse Akrasia and Moral Ignorance." In *The Nature of Moral Responsibility: New Essays*, ed. Randolph Clarke, Michael McKenna, and Angela M. Smith, 141–56. Oxford: Oxford University Press.
- Barnett, Zach. 2021. "Rational Moral Ignorance." *Philosophy and Phenomenological Research* 102: 645–64.
- Barry, Christian, and Patrick Tomlin. 2016. "Moral Uncertainty and Permissibility: Evaluating Option Sets." *Canadian Journal of Philosophy* 46: 898–923.
- Bradley, Ben. 2006. "Against Satisficing Consequentialism." *Utilitas* 18: 97–108.
- Broome, John. 2013a. "Enkrasia." *Organon F* 20: 425–36.
- Broome, John. 2013b. *Rationality Through Reasoning*. Oxford: Wiley Blackwell.
- Carr, Jennifer Rose. 2020. "Normative Uncertainty Without Theories." *Australasian Journal of Philosophy* 98: 747–62.
- Carr, Jennifer Rose. 2022. "The Hard Problem of Intertheoretic Value Comparisons." *Philosophical Studies* 179: 1401–27.
- Cotton-Barratt, Owen, and Hilary Greaves. 2024. "A Bargaining-Theoretic Approach to Moral Uncertainty." *Journal of Moral Philosophy* 21: 127–69.
- Cotton-Barratt, Owen, William MacAskill, and Toby Ord. 2020. "Statistical Normalization Methods in Interpersonal and Intertheoretic Comparisons." *The Journal of Philosophy* 117: 61–95.
- Foot, Philippa. 1978. *Virtues and Vices*. Oxford: Oxford University Press.
- Gracely, Edward J. 1996. "On the Noncomparability of Judgments Made by Different Ethical Theories." *Metaphilosophy* 27: 327–32.
- Gustafsson, Johan E, and Olle Torpman. 2014. "In Defence of My Favourite Theory." *Pacific Philosophical Quarterly* 95: 159–74.
- Harman, Elizabeth. 2011. "Does Moral Ignorance Exculpate?" *Ratio* 24: 443–68.
- Harman, Elizabeth. 2015. "The Irrelevance of Moral Uncertainty." *Oxford Studies in Metaethics* 10: 53–79.

- Hedden, Brian. 2016. "Does MITE Make Right? On Decision Making Under Normative Uncertainty." *Oxford Studies in Metaethics* 11: 102–28.
- Hicks, Amelia. 2018. "Moral Uncertainty and Value Comparison." *Oxford Studies in Metaethics* 13: 161–83.
- Horton, Joe. 2017. "The All or Nothing Problem." *The Journal of Philosophy* 114: 94–104.
- Hudson, James L. 1989. "Subjectivization in Ethics." *American Philosophical Quarterly* 26: 221–29.
- Hursthouse, Rosalind. 1999. *On Virtue Ethics*. Oxford: Oxford University Press.
- Johnson King, Zoë A. 2020a. "Accidentally Doing the Right Thing." *Philosophy and Phenomenological Research* 100: 186–206.
- Johnson King, Zoë A. 2020b. "Praiseworthy Motivations." *Noûs* 54: 408–30.
- Kiesewetter, Benjamin. 2016. "You Ought to Φ Only If You May Believe That You Ought to Φ ." *The Philosophical Quarterly* 66: 760–82.
- Lazar, Seth. 2017. "Deontological Decision Theory and Agent-Centred Options." *Ethics* 127: 579–609.
- Lockhart, Ted. 2000. *Moral Uncertainty and Its Consequences*. Oxford: Oxford University Press.
- MacAskill, William. 2013. "The Infectiousness of Nihilism." *Ethics* 123: 508–20.
- MacAskill, William. 2015. *Doing Good Better*. New York: Avery.
- MacAskill, William. 2016a. "Normative Uncertainty as a Voting Problem." *Mind* 125: 967–1004.
- MacAskill, William. 2016b. "Smokers, Psychos, and Decision-Theoretic Uncertainty." *The Journal of Philosophy* 113: 425–45.
- MacAskill, William, Krister Bykvist, and Toby Ord. 2020. *Moral Uncertainty*. Oxford: Oxford University Press.
- MacAskill, William, and Toby Ord. 2020. "Why Maximize Expected Choiceworthiness?" *Noûs* 54: 327–53.
- Markovits, Julia. 2010. "Acting for the Right Reasons." *The Philosophical Review* 119: 201–42.
- Moller, Dan. 2011. "Abortion and Moral Risk." *Philosophy* 86: 425–43.
- Parfit, Derek. 2011. *On What Matters: Volume One*. Oxford: Oxford University Press.
- Pittard, John, and Alex Worsnip. 2017. "Metanormative Contextualism and Normative Uncertainty." *Mind* 126: 155–93.
- Podgorski, Abelard. 2020. "Normative Uncertainty and the Dependence Problem." *Mind* 129: 43–70.
- Pummer, Theron. 2016. "Whether and Where to Give." *Philosophy & Public Affairs* 44: 77–95.
- Robinson, Pamela. 2022. "Is Normative Uncertainty Irrelevant if Your Descriptive Uncertainty Depends on It?" *Pacific Philosophical Quarterly* 103: 874–99.
- Rosenthal, Chelsea. 2021. "What Decision Theory Can't Tell Us About Moral Uncertainty." *Philosophical Studies* 178: 3085–3105.
- Ross, Jacob. 2006. "Rejecting Ethical Deflationism." *Ethics* 116: 742–68.
- Russell, Jeffrey Sanford. n.d. "The Value of Normative Information." *Australasian Journal of Philosophy*. Forthcoming.
- Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Sepielli, Andrew. 2009. "What to Do When You Don't Know What to Do." *Oxford Studies in Metaethics* 4: 5–28.
- Sepielli, Andrew. 2013. "Moral Uncertainty and the Principle of Equity among Moral Theories." *Philosophy and Phenomenological Research* 86: 580–89.
- Sepielli, Andrew. 2019. "Decision Making Under Moral Uncertainty." In *The Routledge Handbook of Moral Epistemology*, ed. Aaron Zimmerman, Karen Jones, and Mark Timmons, 508–21. New York: Routledge.

- Sinclair, Thomas. 2018. "Are We Conditionally Obligated to be Effective Altruists?" *Philosophy & Public Affairs* 46: 36–59.
- Sliwa, Paulina. 2016. "Moral Worth and Moral Knowledge." *Philosophy and Phenomenological Research* 93: 393–418.
- Slote, Michael. 1985. *Common-Sense Morality and Consequentialism*. London: Routledge.
- Slote, Michael, and Philip Pettit. 1984. "Satisficing Consequentialism." *Aristotelian Society Supplementary Volume* 58: 139–76.
- Staffel, Julia. 2021. "Normative Uncertainty and Probabilistic Moral Knowledge." *Synthese* 198: 6739–6765.
- Sung, Leora. n.d. "Supererogation, Suberogation, and Maximizing Expected Choiceworthiness." *Canadian Journal of Philosophy*. Forthcoming.
- Tarsney, Christian. 2019. "Normative Uncertainty and Social Choice." *Mind* 128: 1285–1308.
- Tarsney, Christian. 2021. "Vive la Différence? Structural Diversity as a Challenge for Metanormative Theories." *Ethics* 131: 151–82.
- Way, Jonathan, and Daniel Whiting. 2016. "If You Justifiably Believe That You Ought to Φ , You Ought to Φ ." *Philosophical Studies* 173: 1873–1895.
- Weatherson, Brian. 2002. "Review: Moral Uncertainty and Its Consequences." *Mind* 111: 693–96.
- Weatherson, Brian. 2014. "Running Risks Morally." *Philosophical Studies* 167: 141–63.
- Wedgwood, Ralph. n.d. "Decision-Theoretic Virtue Ethics." Unpublished manuscript.
- Zimmerman, Michael J. 2008. *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge: Cambridge University Press.