

ARTICLE

Moral uncertainty

Krister Bykvist^{1,2}¹Dept. of Philosophy, Stockholm University²Institute for Futures Studies, Stockholm**Correspondence**

Krister Bykvist, Professor of Practical Philosophy, Dept. of Philosophy, Stockholm University, SE-10691 Stockholm, Sweden.
 Email: krister.bykvist@philosophy.su.se

Abstract

What should we do when we are not certain about what we morally should do? There is a long history of theorizing about decision-making under empirical uncertainty, but surprisingly little has been written about the moral uncertainty expressed by this question. Only very recently have philosophers started to systematically address the nature of such uncertainty and its impacts on decision-making. This paper addresses the main problems raised by moral uncertainty and critically examines some proposed solutions.

1 | INTRODUCTION

How important is the well-being of non-human animals compared with the well-being of humans?

How much should we spend on helping strangers in need?

How much should we care about future generations?

How should we weigh reasons of autonomy and respect against reasons of benevolence?

Few could honestly say that they are fully certain about the answers to these pressing moral questions. Part of the reason we feel less than fully certain about the answers has to do with uncertainty about empirical facts. We are uncertain about whether fish can feel pain, whether we can really help strangers far away, or what we could do for people in the far future. However, sometimes, the uncertainty is fundamentally moral. We can be uncertain about how to weigh reasons of autonomy against reasons of benevolence, not because we lack an understanding of what these reasons are (what autonomy and benevolence mean, whose autonomy and well-being are at stake, and how much autonomy and well-being are at stake, and so on) but because we are not sure which of the reasons is more morally important. This fundamental moral uncertainty can also be part of why we are uncertain about the answers to the other questions. Even if we were to come to know all the relevant non-normative facts, we could still waver about whether it is right to kill an animal for a very small benefit for a human, whether we have strong duties to help strangers in need, and whether future people matter as much as current ones. Fundamental moral uncertainty can also be more general as when we are uncertain about whether a certain moral theory is correct. Many first-year students express deep uncertainty about which moral theory is correct after having taken an introductory course on normative ethics, where all the standard theories are exposed with “warts and all.”

There is a long history of theorizing about decision making under empirical uncertainty, but surprisingly little has been written about fundamental moral uncertainty, except for a somewhat related discussion, starting in the late 16th century, among Catholic theologians about what to do under uncertainty about divine moral laws (*The New Catholic Encyclopedia*, 2002; Sepielli, 2010, ch. 1). Only very recently have secular philosophers started to systematically address the nature of such uncertainty and its impacts on decision-making. One of the first clear discussions of moral uncertainty is found in (Hudson, 1989). In (Oddie, 1995), moral uncertainty is applied to abortion. The first thorough

book-length project on moral uncertainty was (Lockhart, 2000). The discussion was revived in Andrew Sepielli's PhD thesis in 2010 and his papers (Sepielli, 2009, 2010, 2012a, 2012b, 2013, 2016). The latest thorough treatment of this topic is to be found in Will MacAskill's PhD thesis in 2014 and the subsequent papers (MacAskill, 2013, 2016, forthcoming). For an up to date and comprehensive discussion, see (MacAskill, Ord, & Bykvist, forthcoming).

2 | THE OUGHT OF MORAL UNCERTAINTY

One possible explanation of why so few philosophers have engaged with moral uncertainty might be serious doubt about whether it makes much sense to ask about what one ought do when one is uncertain about what one ought to do. The obvious answer to this question might be thought to be: "you ought to do what you ought to do, no matter whether or not you are certain about it" (Weatherson, 2002, 2014). However, this assumes the same sense of "ought" throughout.

A better option is to assume that there are different kinds of moral ought. We are asking what we morally ought to do, in one sense of ought, when we are not certain about what we morally ought to do, in another sense of ought. One way to make this idea more precise is to think about the different senses as different *levels* of moral ought. When we face a moral problem, we are asking what we morally ought to do, at the first level. Standard moral theories, such as utilitarianism, Kantianism, and virtue ethics, provide answers to this question. In a case of moral uncertainty, we are moving up one level and asking about what we ought to do, at the second level, when we are not sure what we ought to do at the first level. At this second level, we take into account our credence in various hypotheses about what we ought to do at the first level and what these hypotheses say about the moral value of each action (MacAskill et al., forthcoming). This second level ought provides a way to cope with the moral uncertainty at the first level. It gives us a verdict of how to best manage the risk of doing first order moral wrongs. That there is such a second-level moral ought of coping with first-order moral risks seems to be supported by the fact that agents are morally criticizable when they, knowing all the relevant empirical facts, do what they think is very likely to be a first-order moral wrong when there is another option that is known not to pose any risk of such wrongdoing. (For more on this, see section on dominance reasoning below.)

One obvious objection to this understanding of ought is that it looks like it will lead to an infinite regress: when we are not sure what we are ought to do at the second level, we move up to yet another level and ask what we are ought to do at the third level, and so on ad infinitum. (Weatherson, 2014). It is not clear that there is a genuine threat of an actual infinite regress, however. Human agents have cognitive limitations and cannot consider an infinite number of levels. There is thus a natural limit on how many levels up we can go in our moral uncertainty.

On the other hand, this natural limit might not correspond to any intuitive notion of moral ought. It seems we can at least move up to the third or fourth level quite easily. But is there really an intuitive moral ought at these levels? (For more on the regress problem, see Harman, 2015 and Sepielli, forthcoming.)

An alternative way to understand the ought relevant to moral uncertainty is in terms of rationality (MacAskill et al., forthcoming; Sepielli, 2013). Rationality, in one important sense at least, has to do with what one should do or intend, given one's beliefs and preferences. This is the kind of rationality that decision theory often is seen as invoking. It can be spelled out in different ways. One is to see it as a matter of coherence: It is rational to do or intend what coheres with one's beliefs and preferences (Broome, 2013; for a critic, see Arpaly, 2000). Another way to spell it out is to understand it as matter of rational processes: it is rational to do or intend what would be the output of a rational process, which starts with one's beliefs and preferences (Kolodny, 2007). To apply the general idea to moral uncertainty, we do not need to take stand on which version is correct. We only need to assume that when a conscientious moral agent faces moral uncertainty, she cares about doing right and avoid doing wrong but is uncertain about the moral status of her actions. She prefers doing right to doing wrong and is indifferent between different right doings (at least when the right doings have the same moral value, that is, none is morally supererogatory). She also cares more about serious wrongdoings than minor wrongdoings. The idea is then to apply traditional decision-

theoretical principles, according to which rational choice is some function of the agent's preferences (utilities) and beliefs (credences). Of course, different decision-theories provide different principles (and require different kinds of utility information). But the plausible ones at least agree on cases where one option dominates another. Suppose that you are considering only two theories (which is to simplify considerably, but we only need a logically possible case): "business as usual," according to which it is permissible to eat factory-farmed meat and permissible to eat vegetables, and "vegetarianism," according to which it is impermissible to eat factory-farmed meat and permissible to eat vegetables. Suppose further that you have slightly more confidence in "business as usual." The option of eating vegetables will dominate the option of eating meat in terms of your own preferences: No matter which moral theory is true, by eating vegetables, you will ensure an outcome that you weakly prefer to the alternative outcome: if "vegetarianism" is true, you prefer the outcome; if "business as usual" is true, you are indifferent between the outcomes. The rational thing for you to do is thus to eat vegetables, given your beliefs and preferences.

Now, against the claim that a morally conscientious person cares about doing right and avoiding doing wrong one might object that I have depicted the conscientious agent as *moral fetishist*, someone who cares about rightness and wrongness as such rather than what makes actions right or wrong. One could argue that a conscientious agent should care about helping the needy, keeping promises, and not be concerned with doing the right thing as such (Smith, 1994, ch. 3; Weatherson, 2014).

This objection assumes a false dichotomy, however. It is possible, and morally commendable, to care about both. An agent who cares only about moral wrongness seems deficient: She should also care about what makes actions wrong. But, equally, an agent who cares only about the well-being of individuals and not at all about whether her actions are wrong would also be deficient as a moral agent. After all, coming to see an action as wrong motivates a morally conscientious agent to change her intrinsic concerns so that she starts to care intrinsically about what makes actions wrong according to her newly acquired moral beliefs. (Dreier, 2000; Shafer-Landau, 2000) For example, a morally conscientious agent, who becomes fully certain that lying is wrong because it is disrespectful, is motivated to become intrinsically concerned with avoiding disrespectful actions. Furthermore, if she becomes fully certain that the severity of a wrongdoing is greater if it shows more disrespect, she is motivated to become more intrinsically concerned with avoiding more disrespectful actions.

When a morally conscientious agent is not fully certain that a feature is wrong-making, she should be sensitive to the probability of this feature being wrong-making and the extent to which it is wrong-making, that is, how strongly the feature speaks against an action. More exactly, if she is ideally conscientious, her intrinsic moral concern for a feature should be a function of both (a) the probabilities of her hypotheses about the extent to which the feature is wrong-making and (b) the extent to which the feature is wrong-making, according to these hypotheses. For example, in the case above with the agent facing the choice to eat vegetarian food or to eat meat, the morally conscientious agent is intrinsically more averse to harming animals by eating them, the more certain she is that harming animals is a wrong-making feature and the greater the extent is to which she thinks harming animals is such a feature. (Bykvist, 2014; Sepielli, 2016)

3 | MY FAVOURITE THEORY AND DOMINANCE REASONING

Another possible explanation of why moral uncertainty has not received much attention is that the challenge of moral uncertainty is thought to have an easy and obvious solution: If you are uncertain about which moral theory is correct (or which moral principle applicable to the situation at hand is correct), just follow the prescriptions of your "favourite theory," the theory you have most confidence in. This approach is very problematic, however, as (Lockhart, 2000) and (MacAskill, 2014) show.

First, the approach does not give us any advice when we have equally high credence in many different theories, because in this case, there is no theory you have most credence in. Second, it is not sufficiently sensitive to avoidable moral risks. Take our example above where you give slightly more credence to "business as usual" than to

“vegetarianism.” On the “my favourite theory”-approach, you are allowed to eat meat in a choice between eating meat and eating vegetables. But this seems wrong (especially if you find meat and vegetables equally tasty), because in this situation, not eating meat dominates eating meat: No matter which theory is true, you do something permissible if you eat vegetables; if you eat meat, however, you risk doing something impermissible. It seems wrong to risk a moral wrongdoing, when you can easily (and without any cost to yourself) avoid the risk. Finally, this theory is very sensitive to how you individuate different moral theories. Suppose you have 40% credence in Kantianism, according to which action A is wrong and action B is right, and 60% credence in utilitarianism, according to which A is right and B wrong. Then “my favourite theory” will tell us to choose A. But suppose the agent remembers that there are more than one version of utilitarianism and then splits the utilitarian hypothesis into two hypotheses, hedonistic utilitarianism and non-hedonistic utilitarianism, with 30% credence for each version of utilitarianism. Then “my favourite theory” will tell me to choose B. (Gustafsson & Torpman, 2014) defend a sophisticated version of “my favourite theory” -approach that avoids some of the problems mentioned here.

Dominance reasoning is no longer uncontroversial when, according to the theory you have most confidence in, some actions are supererogatory and very demanding. Suppose that you have very little credence in Peter Singer’s form of utilitarianism, according to which you are required to help others in need even at great costs to yourself, and a lot of credence in a common-sense theory, according to which you are permitted but not required to help others when it involves great self-sacrifice. Applying unqualified dominance reasoning gives us the result you are required to make the great sacrifice to help others even though you have very little confidence in the theory that tells you that you are morally required to do so. Any fully developed theory of moral uncertainty has to address this important problem. For a clear statement of this problem and new solution to it, see (Barry & Tomlin, 2016).

4 | MAXIMIZING EXPECTED MORAL VALUE

Obviously, an account of moral uncertainty would be too restricted if it could only give verdicts in cases where one action dominates another. We also need guidance in cases where the relevant moral theories (principles, considerations, or reasons) favour different options. For example, suppose you are morally uncertain about the relative importance of non-human versus animal well-being. Suppose that you are certain that giving a certain pill to John will cure him partially and you are certain that giving this pill to Jane will cure her completely. You know that John is a human being and Jane is an animal, but you are not sure how much animal welfare matters in comparison to human welfare. Your probabilities divide equally between *impartialism*, according to which partially curing John is impermissible and fully curing of Jane is permissible, and *speciesism*, according to which partially curing John is permissible and fully cure of Jane is impermissible. Unfortunately, there is only one pill to go. What should you do? (the example is borrowed from Zimmerman, 2008).

One option is to try to identify, for each theory, the moral value it assigns to the options, for example, the degree of severity of the wrongdoings. How much of a difference in moral value is there between the two options, according to speciesism, and how much of a difference is there, according to impartialism? If the value difference is greater according to speciesism, then we should partially cure the human rather than fully cure the animal. Of course, this only works if the probabilities are split even between the theories. What should we do when we have more confidence in impartialism, according to which here is a smaller difference in value between the options? If we can compare the ratios of these differences (so that moral value is measured on an interval scale), then we can weigh these value differences against the probabilities of the two theories. We can even apply the standard expected utility approach to moral uncertainty and go for the maximization of expected moral value.

5 | INTER-THEORETICAL COMPARISONS OF MORAL VALUE

Of course, this is a big “if.” One pressing question is whether we can at all meaningfully say that the value difference between two options, according to one moral theory, is smaller or greater than the value difference between the

options, according to another (for an early statement of this problem, see Gracely, 1996). It should be noted, however, that we do seem to assume that such comparisons are possible at least in some cases. It seems sensible to say things like “the importance of saving animals rather than humans is greater, if impartialism is true than if speciecism is true.” As in the case of comparisons of well-being, across different people, there are many possible views on how to compare value differences across theories. One idea is that the end points of value for one theory is treated equally as the end points for another. This idea is defended in (Lockhart, 2000) and criticized in (Sepielli, 2012a). More exactly, the best option for one theory is assumed to be as good as the best option for another, and the worst option for one theory is assumed to be as good as the worst option for another. But this account has serious drawbacks. First, some theories might not have a bounded moral value, so there is no best (or no worst) option. Furthermore, could not one theory’s best option have greater moral value than another theory’s best option? (For more problems with this approach, see Sepielli, 2009 and MacAskill, 2014). In the recent literature on this problem, there are approaches that avoids these problems, but, the jury is still out on whether any of these approaches are overall defensible (Hedden, 2016; MacAskill, 2014; MacAskill et al., forthcoming; Ross, 2006; Sepielli, 2009).

Now, even if it is possible in principle to make these inter-theoretical comparisons, it is clear that this will not always help, for some of the theories we have credence in might explicitly deny that there are any clearly defined value differences between the options. Perhaps, certain absolutist deontological theories, which only talk in terms of absolute duties, are good examples, because according to these theories, you either act in accordance with duty or you act wrongly—no gradable moral value is invoked. For these cases of uncertainty, the expected value approach will not work. The right approach might thus be to go pluralistic and apply different accounts to different cases. For those cases where we can make inter-theoretical comparisons of ratios of value differences, we should maximize expected value. For other cases with less rich information about measurability, we need other accounts. One recent approach is to borrow a page from voting theory and see the degree of credence in a theory’s ranking of the options as analogous to the number of voters who rank the options in a certain way (MacAskill, 2014, forthcoming). Various voting principles can then be applied, without making any inter-theoretical comparisons of value.

For example, consider the Borda rule. In the voting context, this rule tells us first to look at the voters’ preference rankings of the alternatives and assign number n , which is the number of alternatives, to the first ranked alternative in each ranking, number $n-1$, to each second ranked alternative in each ranking, and so on. Then we simply add up these individual scores for each alternative and choose the one with the highest total score. So, if we have the following preference rankings,

Chris: A, B, C

Lisa: B, C, A

Eva: A, C, B

we calculate the total scores in this way. Total score of A = $(2 \times 3) + (1 \times 1) = 7$; total score of B = $(1 \times 3) + (1 \times 2) + (1 \times 1) = 6$; total score of C = $(2 \times 2) + (1 \times 1) = 5$. Here, the winner is A. Applied to the context of moral uncertainty, we replace people with moral theories, individual preference rankings with moral value rankings, and alternatives with options. Finally, we weight the score a moral theory assigns to an option by the credence the agent has in this theory. For example, suppose Chris is a moral theory we have 60% credence in. Then the credence-weighted score of A, according to this theory, will be 0.6×3 . The alternative action with the highest total credence-weighted score is to be chosen.

6 | HIGH-STAKE THEORIES SWAMPING LOW-STAKE THEORIES

Even for those cases where we can make inter-theoretical comparisons, it is not clear that it is always best to maximize expected moral value. A theory, according to which a lot is at stake but in which we only have small credence, can easily swamp other theories, according to which little is at stake, even if we have much more confidence in them. For instance, suppose we have little confidence in a strict deontological theory, according to which we have always

very strong reason not to lie. Suppose further that we have much more confidence in a less strict theory, according to which there is some, but not very strong, reason to lie, especially when you can save people from harm. If we should maximize expected moral value, the strict theory will trump less strict theory (if the strict theory's reasons are sufficiently strong), even though we have very little confidence in it (MacAskill, 2014; MacAskill et al., forthcoming; Barry & Tomlin, 2016).

That a low high-stake theory often wins over a low-stake one, even when the latter is much more probable, cannot in general be troublesome, because it seems sensible to take seriously even small risks of doing something that turns out to be very bad. But some would argue that the case above shows that maximizing expected moral value gives too much weight to high-stake theories.

One option here is simply to modify the expected moral value approach so that it no longer has these swamping effects. Perhaps it is better to think we should satisfice rather than maximize expected moral value, so that it is permissible but not required to maximize expected moral value. Or perhaps, we should adopt a more radical departure from the expected value theory, as is argued in (Barry & Tomlin, 2016), which proposes that we should value and choose option sets rather individual actions in cases of moral uncertainty.

7 | META-ETHICAL IMPLICATIONS

So far I have focused on the normative implications of moral uncertainty, but there are also important meta-ethical implications. One is that certain meta-ethical theories have grave difficulties accommodating the very existence of moral uncertainty. A case in point is non-cognitivism, according to which moral judgements are seen as desires rather than beliefs. The problem for non-cognitivism is that desires seem to have too little structure to account for both degree of moral uncertainty and importance, the degree of value or normative importance you ascribe to an action or a state of affairs. If degrees of moral uncertainty is identified with degrees of desire strength—so that the stronger the expressed desire is, the more certain you are about the moral judgement—then there is nothing left to explain importance. On the other hand, if importance is identified with degrees of desire strength—so that the stronger the expressed desire is, the more value or importance is ascribed to the act or state of affairs—then there is nothing left to explain uncertainty. Of course, this would not be a problem if uncertainty and importance always co-varied, but that is not true. One might, for example, be very uncertain that leading an autonomous life is of great intrinsic value and very certain that experiencing bodily pleasure is of moderate intrinsic value. Similarly, one might be very uncertain that one has a strong reason to save a stranger's two children at the cost of the life of one's own child and be very certain that one has a weak reason to satisfy one's whims. For another meta-ethical implication of moral uncertainty, see (MacAskill, 2013), where it is argued that credence in nihilism, the idea that moral judgements are false beliefs, causes havoc for the idea that we should maximize expected moral value, because value will simply be undefined according to the nihilistic hypothesis.

Obviously, the cognitivist, who simply identifies moral judgements with beliefs, has no problem capturing these cases. Degrees of moral judgements are simply degrees of beliefs and the degree of belief in a moral proposition can vary independently from the degree of moral importance ascribed to an action or a state of affairs. But for the non-cognitivist, it is a real challenge, and so far, no fully satisfactory solution has been found (for more on this problem, see Bykvist & Olson, 2009; Bykvist & Olson, 2012; Lenman, 2003; Ridge, 2007; Sepielli, 2012b; Smith, 2002).

8 | EXTENSIONS AND CONCLUDING REMARKS

As seen, moral uncertainty has its own stock of tricky questions, some of which are still relatively unexplored. It offers therefore a fertile ground for new research. This research need not be confined to the moral domain, however, for the problems can easily be generalized to other forms of normative uncertainty: uncertainty about rationality, about epistemic principles and reasons, and about prudential reasons (for a discussion of uncertainty

about what to do under decision-theoretic uncertainty, see Nozick, 1993, Sepielli, 2013 and MacAskill, 2016; for a discussion of second-order credence, see Christensen, 2010 and Lasonen-Aarnio, 2014). There is also a host of uncertainty questions about value. For example, one can be uncertain about the final or intrinsic value of things, which well-being theory to adopt and which aesthetic values are relevant (for a discussion of uncertainty about final value, see Zimmerman, 2008, 2014).

Today's most important moral questions involve deep uncertainty in empirical, epistemic, evaluative, and normative matters. Just think about the deep empirical, epistemic, and moral questions raised by global warming and climate change: "What are the probabilities of the possible future scenarios?", "How do we know these probabilities?" and "How should we evaluate the scenarios?" What is urgently needed is an approach that takes all these uncertainties into account.

WORKS CITED

- Arpaly, N. (2000). On acting rationally against one's best judgement. *Ethics*, 110(3), 488–513.
- Barry, C., & Tomlin, P. (2016). Moral uncertainty and permissibility: evaluating option sets. *Canadian Journal of Philosophy*, 46(6), 1–26.
- Broome, J. (2013). *Rationality through reasoning*. Wiley-Blackwell.
- Bykvist, K. (2014). Evaluative uncertainty and consequentialist environmental ethics. In L. Kahn, & A. Hiller (Eds.), *Environmental ethics and consequentialism*. Routledge.
- Bykvist, K., & Olson, J. (2009). Expressivism and moral certitude. *The Philosophical Quarterly*, 59(235), 202–215.
- Bykvist, K., & Olson, J. (2012). Against the being for account of normative certitude. *Journal of Ethics and Social Philosophy*.
- Christensen, D. (2010). Higher-order evidence. *Philosophy and Phenomenological Research*, 81(1), 185–215.
- Dreier, J. (2000). Dispositions and fetishes: Externalist models of moral motivation. *Philosophy and Phenomenological Research*, 61(3), 619–638.
- Gracely, E. J. (1996). On the noncomparability of judgments made by different ethical theories. *Metaphilosophy*, 27, 327–332.
- Gustafsson, J., & Torpman, T. (2014). In defence of my favourite theory. *Pacific Philosophical Quarterly*, 95(2), 159–174.
- Harman, E. (2015). The irrelevance of moral uncertainty. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics*. Oxford University Press.
- Hedden, B. (2016). Does MITE make right? On decision-making under normative uncertainty. In R. S. Landau (Ed.), *Oxford studies in metaethics 11*. Oxford University Press.
- Hudson, J. L. (1989). Subjectivization in ethics. *American Philosophical Quarterly*, 26, 221–229.
- Kolodny, N. (2007). State of process requirements? *Mind*, 116(462), 371–385.
- Lasonen-Aarnio, M. (2014). Higher-order evidence and the limit of defeat. *Philosophy and Phenomenological Research*, 88(2), 314–345.
- Lenman, Jimmy. (2003). "Non-cognitivism and the dimensions of evaluative judgement." Brown Electronic Article Review Service, <http://www.brown.edu/Departments/Philosophy/bears/homepage.html>, posted 15 March 2003.
- Lockhart, T. (2000). *Moral uncertainty and its consequences*. Oxford University Press.
- MacAskill, W. (2013). The infectiousness of nihilism. *Ethics*, 123(3), 508–520.
- MacAskill, W. (2014). Normative uncertainty, Ph.D. thesis in Philosophy, Department of Philosophy, Oxford University.
- MacAskill, W. (2016). Smokers, Psychos, and Decision-Theoretic Uncertainty. *Journal of Philosophy*.
- MacAskill, W. (forthcoming). Normative uncertainty as a voting problem. *Mind*.
- MacAskill, William, forthcoming, with Ord, Toby, and Bykvist, Krister. *An introduction to moral uncertainty*, Oxford University Press, forthcoming.
- Nozick, R. (1993). *The nature of rationality*. Princeton University Press.
- Oddie, G. (1995). Moral uncertainty and human embryo experimentation. In K. W. M. Fulford, G. Gillett, & J. M. Soskice (Eds.), *Medicine and moral reasoning*. Oxford University Press.
- Ridge, M. (2007). Ecumenical expressivism: The best of both worlds? In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 2). (pp. 51–77. Oxford University Press.
- Ross, J. (2006). Rejecting ethical deflationism. *Ethics*, 116, 742–768.

- Sepielli, A. (2009). What to do when you don't know what to do. In R. Shafer-LandOau (Ed.), *Oxford studies in metaethics: 4*. Oxford University Press.
- Sepielli, Andrew. (2010). 'Along an imperfectly lighted path': Practical rationality and normative uncertainty. PhD Thesis, Rutgers University.
- Sepielli, A. (2012a). Moral uncertainty and the principle of equity among moral theories. *Philosophy and Phenomenological Research*, 86(3), 580–589.
- Sepielli, A. (2012b). Normative uncertainty for non-cognitivists. *Philosophical Studies*, 1–17.
- Sepielli, A. (2013). What to do when you don't know what to do when you don't Know what to do.... *Noûs*, 48(3), 521–544.
- Sepielli, A. (2016). Moral uncertainty and fetishistic motivation. *Philosophical Studies*, 173(11), 2951–2968.
- Sepielli, A. (forthcoming). How moral uncertainty can be both true and interesting. In M. Timmons (Ed.), *Oxford studies in normative ethics*. Oxford University Press.
- Shafer-Landau, R. (2000). A defence of motivational externalism. *Philosophical Studies*, 97, 267–291.
- Smith, M. (1994). *The moral problem*. Wiley-Blackwell.
- Smith, M. (2002). Evaluation, uncertainty and motivation. *Ethical Theory and Moral Practice*, 5(3), 305–320.
- The New Catholic Encyclopedia, 2nd ed. 2002. The Catholic University Press.
- Weatherson, B. (2002). Review of Ted Lockhart's 'moral uncertainty and its consequences'. *Mind*, 111, 693–696.
- Weatherson, B. (2014). Running risks morally. *Philosophical Studies*, 167(1), 141–163.
- Zimmerman, M. J. (2008). *Living with uncertainty: The moral significance of ignorance*. Cambridge University Press.
- Zimmerman, M. J. (2014). *Ignorance and moral obligation*. Oxford University Press.

How to cite this article: Bykvist K. Moral uncertainty. *Philosophy Compass*. 2017;e12408. doi: 10.1111/phc3.12408

AUTHOR BIOGRAPHY

Krister Bykvist is Professor of Practical Philosophy at the Department of Philosophy, Stockholm University and a Research Fellow at the Institute for Futures Studies. Before moving to Stockholm, he was a Tutorial Fellow at Jesus College, Oxford, and a CUF Lecturer in Philosophy at Oxford University, (2001–2013).