

**Moral Uncertainty, Metaethics,  
and Natural Law**

Corin A. Katzke

Yale University

PHIL 423: Philosophy of Probability

**Moral Uncertainty, Metaethics,  
and Natural Law**

Suppose you are uncertain which theory of normative ethics is true, and the theories in which you have some credence disagree about how you should act. How, then, should you act? This is the problem of moral uncertainty.

An answer to the problem of moral uncertainty depends in part on our metaethics. After all, the extent to which we can compare options across ethical theories should plausibly depend on what exactly we think ethical theories *are*. In particular, I will argue that comparing what MacAskill and Ord (2018) call the *choice-worthiness* of options across theories isn't possible according to *moral realism*, but is — in principle — possible according to *moral constructivism*.

After introducing the problems of moral uncertainty and intertheoretic comparability, I will show that moral realism and moral constructivism involve opposite explanatory relationships between theories and particular judgements. Based on these explanatory relationships, I will make an analogy between moral constructivism and Humeanism about natural laws one hand, and moral realism and anti-Humeanism about natural laws on the other. This analogy will allow me to show why intertheoretic comparisons of choice-worthiness are in principle possible according to moral constructivism, but not according to moral realism.

### **Bayesian epistemology and subjective uncertainty**

I'll begin by motivating the problem of moral uncertainty by reference to both common epistemic humility and Bayesian epistemology.

Harman (2015) argues that the correct response to moral uncertainty is straightforward – what a person *should* do is act according to whichever theory is actually correct. If we take ‘should’ in the sense of an ethical imperative, then of course Harman is right. But we should not confuse this sense with the decision-theoretic sense of ‘should’, which makes a normative claim of its own. Suppose, for example, that someone is certain of the *wrong* ethical theory. There is a sense of ‘should’ in which they *should* act according to the wrong ethical theory; it would be irrational for them not to determine their actions according to their beliefs. As Lewis puts it, “[c]anons of reasonable belief need not be counsels of perfection. A moral code that forbids all robbery may also prescribe that if one nevertheless robs, one should rob only the rich.” (265) Therefore, Harman’s response only guides action if you are certain that a particular theory is true.

However, if you *are* certain that a particular theory is true, then you are likely not a rational believer. This follows from basic epistemic humility in the face of the complexity and lack of consensus in normative ethics. But it also follows — in a stronger form — from Bayesian epistemology. We can formalize ethical uncertainty applying the principles of Bayesian probability to belief formation. According to Bayesian epistemology, belief is characterized by credences between

full-blown disbelief and belief. Credences obey the laws of probability (Strevens 2017), so we can formalize them as subjective probabilities with values between 0 and 1.

Bayesian epistemology gives an account of how to update subjective probabilities in ethical theories in light of new evidence. (Evidence pertaining to ethical theories could be, for example, discovering an unwelcome consequence of a theory, or learning that a philosopher you respect advocates a theory.) For evidence  $e$  and theory  $t$ :

$$P(e|t) = \frac{P(t|e)}{P(t)}P(e)$$

$P(t)$  represents your prior belief in a theory. An implication of Bayesian epistemology is that, unless your prior in a particular theory is 0 or 1, no amount of evidence will ever lower your subjective probability of that theory to 0 or raise it to 1. It's implausible that a prior in any ethical theory should be 0 or 1 (unless, perhaps, it's logically inconsistent — but even then, can you be sure it's inconsistent?), so a Bayesian believer will always be to some degree uncertain about any particular ethical theory. Therefore, moral uncertainty is always a problem for a Bayesian believer. Harman's response should not counsel us.

## **Expected choice-worthiness and the problem of intertheoretic comparisons**

The first place we might look for an account of action under moral uncertainty is in the literature on empirical uncertainty. The most popular account of action under empirical uncertainty is **expected value theory** (EVT). According to EVT, the right action to take under empirical uncertainty is the action with the greatest *expected value*, where the expected value of an action is calculated by summing the values of each of the action's outcomes multiplied by their probabilities.

MacAskill and Ord (2018) propose an analogous account of action under ethical uncertainty: **expected choice-worthiness** (ECW). According to ECW, the right action to take under ethical uncertainty is the action with the greatest expected *choice-worthiness*, where the expected choice-worthiness of an action is calculated by summing the choice-worthiness's of each of the action's outcomes multiplied by their probabilities. ECW fits together nicely with Bayesian epistemology insofar as it admits precise subjective probabilities.

Respectively, EVT and ECW treat value and choice-worthiness equivalently. However, where value is defined *intratheoretically* (for example, utilitarianism defines value as wellbeing), choice-worthiness is defined *intertheoretically*. Choice-worthiness is intended as a measure with which to compare an action across

different theories. For example, suppose you have credence in two theories (0.51 in  $T_1$  and 0.49 in  $T_2$ ) and are deciding between two options ( $A$  and  $B$ ):

	Option $A$	Option $B$
Theory $T_1$ (0.51)	Slightly Worse	Slightly Better
Theory $T_2$ (0.49)	Much Better	Much Worse

Option  $B$  is only slightly better according to  $T_1$ . In contrast, option  $A$  is *much* better according to  $T_2$ . The intuition behind the possibility of intertheoretic comparability is that you should choose option  $A$  because this difference in relative evaluation of options ‘outweighs’ your slightly higher credence in  $T_1$ . Choice-worthiness fills this theoretical role.

The problem with choice-worthiness is that sometimes intertheoretic comparisons are senseless. For example, suppose your credence is split between total utilitarianism (TU) and average utilitarianism (AU). You have a 0.51 credence in AU, and a 0.49 credence in TU. Suppose that 100 people exist with wellbeing levels of 100, that you are deciding whether to A) bring 100 more lives with wellbeing levels of 98 into existence, or B) refrain from doing so.

	Option $A$	Option $B$
Theory AU (0.51)	Average wellbeing = 99	Average wellbeing = 100
Theory TU (0.49)	Total wellbeing = 19,800	Total wellbeing = 10,000

We might be tempted to say that option *B* is only *slightly* better according to *AU*, and option *A* is *much* better according to *TU*, so we should choose option *A*. But this intuition is senseless. It's true that option *A* is almost twice as good to *TU* than is option *B* — but, according to *TU*, the choice-worthiness of adding 100 lives worth living is the same regardless of how many other people exist. But how many other people exist is what determines the 'twice-as-good' ratio. If, instead of 100, 1,000,000 people existed with wellbeing levels of 100, then bringing 100 more lives with wellbeing levels of 98 into existence would be equally choice-worthy according to *TU*, but the resulting world would not be 'twice-as-good.' In contrast, according to *AU*, the choice-worthiness of adding 100 lives worth living is *not* the same regardless of how many other people exist. Therefore, the choice-worthiness of options are incomparable across *TU* and *AU*.

MacAskill, Bykist, and Ord (2020) agree that the choice-worthiness of options are incomparable across some theories, such as between *TU* and *AU*. In those cases, they argue against simply choosing according to your preferred theory or option, and instead propose using either the Borda rule (for merely ordinal theories) or variance voting (for interval-scale theories). In my view, both proposals are reasonable. However, the question with which this paper is concerned is whether ECW is *ever* appropriate — that is, whether choice-worthiness can ever be compared across theories.

## Cognitivism and Noncognitivism

MacAskill, Bykist, and Ord (2020) argue that it at least *seems* like choice-worthiness can sometimes be compared across theories, which gives *prima facie* plausibility to ECW. But they also recognize that the question of “[w]hat grounds intertheoretic comparisons of value” would “benefit from much greater study” (212)

I argue that this study should draw from the literature on metaethics. Our answer to the question of whether choice-worthiness can be compared intertheoretically in all likelihood depends on what an ethical theory *is* in the first place.

However, insofar as philosophers have discussed metaethics in the context of moral uncertainty, they have mainly distinguished between cognitivism and noncognitivism. According to *cognitivism*, moral statements are propositions, and therefore have truth-values. In contrast, *noncognitivism* describes moral statements in terms of non-propositional content, and therefore doesn’t assign them truth-values. For example, *expressivism* — a kind of noncognitivism — holds that moral statements merely *express* ‘pro’ and ‘con’ attitudes toward their objects.

Moral uncertainty presents a challenge for — and perhaps *to* — noncognitivism. It’s natural to understand moral uncertainty in terms of subjective probability and choice-worthiness, and cognitivism allows for both of these degrees of freedom. However, noncognitivism is not naturally rich enough to account for both. For example, expressivism allows for one degree of freedom: the strength of

the expressed attitude. This strength might correspond to subjective probability, or it might correspond choice-worthiness — but it can't correspond to both.

Some expressivists have responded to this problem by building an extra degree of freedom into expressivism — for example, the strength of the attitude of an ideal advisor towards one's own attitude. MacAskill, Bykist, and Ord (2020) retort that such a maneuver is *ad hoc* and undermines the motivation for expressivism. I agree with them that the problem is at least much more straightforward for cognitivism. However, cognitivism is only a broad class of metaethical positions — I argue that some fare better than others.

### **Moral Realism and Moral Constructivism**

During the only passage in which they distinguish between any metaethical positions within cognitivism, MacAskill, Bykist, and Ord (2020) write that:

[...] we must assume that error theory is false (otherwise there would be no subject matter for us to investigate). That leaves us with some form of moral realism. (146-147)

But this conclusion implies a false dichotomy within cognitivism between error theory and moral realism.

**Error theory.** According to error theory, while moral statements are propositions, they are somehow confused and categorically false. MacAskill, Bykist, and Ord (2020) are right that error theory would make moral uncertainty a moot

point — according to error theory, our credence in every ethical theory should be 0. However, they aren't right to assume that the only cognitivist alternative to error theory is moral realism.

**Moral realism.** According to (Street 2006), the defining claim of moral realism is “that there are at least some [moral] truths that hold independently of all our evaluative attitudes.” (110) However, in order to reject error theory, we need only accept the first claim of moral realism: “that there are at least some moral truths that hold.” We need not specify the sense in which they hold. Moral antirealism — the claim that there are *no* moral truths that hold independently of all our evaluative attitudes — is compatible with cognitivism as long as we allow that there are some moral truths that hold *dependently* of our evaluative attitudes.

**Moral constructivism.** The alternative cognitivist position to moral realism is moral constructivism — a position which MacAskill, Bykist, and Ord (2020) entirely neglect. According to moral constructivism, a moral belief is true if it is implied by our evaluative attitudes and other moral beliefs. Or, more exactly: a moral belief is true if it is a member of the ideally consistent and economical set propositions that captures the greatest part of our evaluative attitudes.

As far as I can tell, no one has yet examined the implications of distinguishing between moral realism and moral constructivism on the problem of intertheoretic comparisons. But this examination is well overdue. For it seems that, while the question of *whether* there exists moral truth is relevant the problem of

intertheoretic comparisons, equally relevant is the *sense in which* there is moral truth.

In particular, moral realism and moral constructivism involve different explanatory relations between particular moral propositions (e.g., option A is better than option B) and general moral principles (e.g., wellbeing is the only intrinsic good). According to moral realism, general moral principles explain — or *found* — the truth of particular moral propositions. For example, particular utilitarian propositions might be true in virtue of the fact that wellbeing is the only intrinsic good. The picture is more complicated for moral constructivism, which does not prescribe any particular direction of explanation. For example, the principle “wellbeing is the only intrinsic good” might be true partially in virtue of the fact that someone has mostly utilitarian intuitions about particular cases.

The relevant features of moral realism and moral constructivism are not their independence or dependence on us, *per se*, but rather these contrasting explanatory directions. Technically, moral realism *could* take particular judgments as ‘real,’ and moral constructivism *could* consider evaluative attitudes about abstract principles. In that case, I would have to come up with different terms in order to make my argument (perhaps, to foreshadow, anti-Humeanism and Humeanism about value). But such relations, though technically consistent, would be strange for moral realism and constructivism, and certainly not how they have been conceived in the metaethical literature. So, with that caveat, I will use the

terms moral realism and moral constructivism in the context of their natural explanatory relations.

As I will argue over the next few sections, these contrasting explanatory relations determine the possibility of intertheoretic comparisons of choice-worthiness.

### **Realist intertheoretic comparisons**

What proponents of ECW need to show is that it's *in principle* possible to compare the value of an option according to one theory with the value of that option according to a different theory. If it is, then we can treat intratheoretic value as intertheoretic choice-worthiness. There are two main proposals for making such a comparison: the **content-based approach**, and the **absolute scale approach**. According to the *content-based* approach, if two theories only disagree about the extension of value, then we can treat where their evaluations overlap as equivalent. According to the *absolute value* approach, there exists an absolute scale of value according to which we can compare options between theories.

**A realist content-based approach.** MacAskill and Ord (2018) argue that there are at least some cases where the content-based approach works. They consider the comparison between total utilitarianism (TU) and anti-Nixon total utilitarianism (TU\*), where TU\* only differs from TU insofar as it doesn't count

Richard Nixon's wellbeing as valuable. According to the content-based approach, these two theories are comparable.

MacAskill and Ord (2018) argue that TU and TU\* agree about the value of every person's wellbeing excluding Nixon. In that case, TU and TU\* agree about the value of, for example, Obama's wellbeing. Therefore, since TU counts Obama's wellbeing as equally valuable to Nixon's, we can calculate the expected choice-worthiness (ECW) of actions affecting Nixon when our credence is split between the two theories.

However, according to moral realism, we aren't justified in assuming that TU and TU\* agree about the value of non-Nixonian wellbeing. If TU and TU\* are consistent theories, then there must be some difference in the general moral principles they endorse that explain their different evaluations of Nixon's wellbeing. For example, the general principle explaining TU's evaluation might be "wellbeing is the only intrinsic good." In contrast, we might gloss the general principle explaining TU\*'s evaluation as "non-Nixonian wellbeing is the only intrinsic good."

We can interpret this second principle as consequentialist: there is a state of affairs — non-Nixonian wellbeing — that constitutes the only intrinsic good. But if we do, then TU and TU\* don't only disagree about the extension of value; they also disagree about what value *is*. In that case, we aren't justified in assuming the TU and TU\* agree about the value of Obama's wellbeing. Strictly speaking, only TU values Obama's wellbeing. TU\* only values Obama's non-Nixonian wellbeing. These

are different states of affairs. *A fortiori*, if we can't compare options between even closely-related consequentialist theories like *TU* and *TU\**, then we should not expect to be able to compare option between more distantly related theories.

Similarly, in one of the first discussions of the problem of intertheoretic comparisons, Gracely (1996) argues that it is impossible to compare person-affecting utilitarianism and total utilitarianism. We might interpret the two theories as disagreeing only about the extension of value — in particular, whether to value the wellbeing of merely possible people. However, Gracely argues that, where total utilitarianism values wellbeing, person-affecting utilitarianism values *existing-person-wellbeing*. Though, like later writers, he didn't consider moral realism and moral constructivism, we should interpret Gracely as condemning a realist content-based approach to intertheoretic comparisons. We should do the same.

**A realist absolute scale approach.** Rather than try to show that two theories merely *agree* about the value of some option, we could try to assign the value each theory ascribes to that option to an intertheoretic absolute scale of value. The problem, of course, is coming up with such a scale.

I argue that an intertheoretic absolute scale of value is senseless in the context of moral realism. According to moral realism, the only existing scale of value the one delineated by the correct moral theory. For example, the scale of value according to total utilitarianism is delineated by wellbeing, and the scale of value according to person-affecting utilitarianism is delineated by existing-person-

wellbeing. In the context of realism, (at most) only one of these scales is correct — even if our credence is split between the two theories. There can't be an absolute scale of value that exists independently of the true moral theory. If there were, then *that* scale would constitute the true moral theory.

Therefore, neither the content-based approach nor the absolute scale approach found intertheoretic comparisons of value in the context of moral realism.

### **An analogy: interpretations of natural law**

Before examining intertheoretic comparisons in the context of moral constructivism, I'm going to take what is hopefully an enlightening detour through the philosophy of natural law. I am not the first writer to compare natural and moral properties. For example, MacAskill, Bykist, and Ord (2020) notice that:

Though work on the metaphysics of quantity has, so far, entirely focused on scientific quantities ('mass', 'size', 'temperature', etc), we can ask just the same questions about the metaphysics of quantities of value, or of choice-worthiness. We can ask: If it is true that the difference in choice-worthiness between *A* and *B* is twice as great as the difference in choice-worthiness between *B* and *C*, is that true in virtue of the fact that *A*, *B* and *C* each have an intrinsic property of a certain degree of choice-worthiness? Or is the metaphysical explanation the other way around? (142)

In particular, they have in mind the difference between absolutism and comparativism with respect to quantity. According to *absolutism* about mass, for example, the mass of an object is defined in terms of an absolute scale. In contrast, *comparativism* about mass defined the mass of an object only in relation to other objects.

This is a relevant analogy, and I will return to it. However, I claim there is another, more enlightening analogy to make: there is divide in the *interpretation of natural law* between *Humeanism* and *anti-Humeanism*. I claim this is relevantly analogous to the divide between moral constructivism and moral realism in the interpretation of moral law.

**Anti-Humeanism about natural laws.** According to Bhogal (2020), anti-Humean views “claim that laws are over and above the patterns of events and, in fact, that the laws govern those events, rather than merely describing them.” (Bhogal 1) Insofar as they ‘govern’ events, natural law *explains* why events take place. This mirrors the explanatory relationship between moral law and particular moral propositions according to moral realism. In the same way that realist moral law explains why a particular option has a particular value, anti-Humean natural law might explain, for example, why a particular object has a particular mass.

Now consider the problem of *natural uncertainty* (bear with me). For example, suppose you have credence in two anti-Humean fundamental natural

theories (0.51 in  $T_1$  and 0.49 in  $T_2$ ) and are choosing between two objects ( $A$  and  $B$ ).

Your ‘goal’ is to choose the object with the greatest expected mass.

	Object $A$	Option $B$
Theory $T_1$ (0.51)	Slightly Less Massive	Slightly More Massive
Theory $T_2$ (0.49)	Much More Massive	Much Less Massive

Object  $B$  is slightly more massive according to your preferred theory,  $T_1$ . However, object  $A$  is *much* more massive according to  $T_2$ . Should you choose object  $A$ ?

No, you shouldn’t — at least not for reason of intertheoretic comparability. Comparisons of mass between fundamental natural theories are senseless in the context of anti-Humeanism because *mass is defined in terms of those theories*. Consider Newtonian mechanics and special relativity as rival natural theories. According to Newtonian mechanics, mass is a measure of resistance to force. While Newton didn’t specify exactly what constitutes mass, we can think of it as just a property of physical stuff: the more stuff, the more mass. In contrast, relativity defines mass as the energy of a system of particles. What’s more, according to relativity, the mass of an object depends on its relative velocity to an observer. The question, “how massive is this object” is not a well-specified question outside of context of a particular theory. This is why physicists use the terms *inertial* mass and *relativistic* mass.

Therefore, anti-Humeanism doesn't allow for intertheoretic comparisons of mass. This mirrors our conclusion about intertheoretic comparisons of value in the context of moral realism.

**Humeanism about natural laws.** In contrast with anti-Humeanism, Humeanism holds that the explanatory relation between events and laws runs the other way — events *explain* laws insofar as laws reduce to descriptions of events:

“[...] laws of nature reduce to the patterns of occurrent, nonmodal, events that occur in the world. The laws of nature are just patterns, or ways of describing patterns, in the mosaic of events. [...] To be more precise,

Humeanism about laws of nature is the view that the laws of nature reduce to the Humean Mosaic—that is, the intrinsic physical state of each space-time point [...].” (Bhogal 1)

This mirrors the explanatory relationship between moral law and particular moral propositions according to moral constructivism. In the same way that intuitions about the value of particular propositions might explain why a constructivist moral law is true, the properties of objects might explain why a particular a Humean law is true.

Consider (again) the problem of natural uncertainty: can we compare the mass of two objects across different Humean theories? We rejected the possibility of such a comparison in the context of anti-Humeanism because, according to anti-Humeanism, mass is defined in terms of a particular theory. But that isn't true in

the context of Humeanism. If mass is a feature of the “the intrinsic physical state of each space-time point,” then mass is defined independently of any particular natural theory, which are only descriptions of those physical states.

Therefore, Humeanism *in principle* allows for intertheoretic comparisons of mass. This is not to say that such comparisons would be straightforward, or even epistemically tractable. It is only to say that the question isn’t senseless.

### **Constructivist intertheoretic comparisons**

I suspect that even the most ornery, stubborn, slow-witted, and uncharitable reader already expects my conclusion: since Humeanism *in principle* allows for intertheoretic comparisons of mass, and Humeanism is relevantly analogous to constructivism, constructivism *in principle* allows for intertheoretic comparisons of value.

Nonetheless, it admits some spelling out. First, by “relevantly analogous,” I mean that, in both Humeanism and constructivism, the direction of explanation runs *from* particular facts (e.g., about mass and value) and *to* theories (of natural law and moral law). This does not imply that Humeanism and constructivism are perfectly analogous. For example, Barry Loewer once became quite bothered during a talk of his at my suggestion of an association between anti-realism and Humeanism. He is right, of course, that Humeanism is ‘realist’ in the sense that it holds that the physical world exists independently of us. In contrast, constructivism

is ‘anti-realist’ in the sense that it holds that no moral propositions are true independently of us. But this distinction is not important to the success of my argument.

Second, I say ‘in principle’ because moral constructivism doesn’t *categorically* admit intertheoretic comparability. To see why, consider that, in order for Humeanism to admit intertheoretic mass comparability, mass must be an intrinsic feature of the Humean mosaic. Additionally, mass must be absolutist rather than comparativist — that is, defined according to an absolute rather than relative scale. Similarly, in order for constructivism to admit intertheoretic value comparability, our evaluative attitudes must define value according to an absolute rather than relative scale. In other words, our intuitions about value have to be such that, not only do we *order* options, but we also *quantify* them according to an absolute scale of value.

So, depending on the nature of our evaluative attitudes, constructivism may or may not admit intertheoretic value comparability. Officially, I’m not going to take a position on this question. But I will suggest that it’s at least plausible our evaluative attitudes do reference an absolute scale of value. MacAskill, Bykist, and Ord (2020) make this point in the context of realism with the following thought experiment:

Consider Thomas, who initially believes that human welfare is ten times as valuable as animal welfare, because humans have rationality and sentience,

whereas animals merely have sentience. He revises this view, and comes to believe that human welfare is as valuable as animal welfare. He might now think that human welfare is less valuable than he previously thought because he has rejected the idea that rationality confers additional value on welfare. Or he might now think that animal welfare is more valuable than he previously thought, because he has extended his concept of rationality, and thinks that animals are rational in the morally relevant sense. (129)

They conclude that, in revising his judgement that human welfare is ten times as valuable as animal welfare to the judgement that human welfare is equally as valuable as animal welfare, Thomas could have either 1) raised his evaluation of animal wellbeing, or 2) lowered his evaluation of human wellbeing. Therefore, his relative evaluation of human and animal wellbeing is not sufficient to fully describe his view. This implies that his judgments reference an absolute scale of value.

This thought experiment is realist in the sense that Thomas *first* revises his theoretical commitments (e.g., that rejecting the idea that rationality confers additional value on wellbeing), which *in turn* explain his revised judgement (that human welfare is as valuable as animal welfare). But, as I argued above, realist absolute scale of value only makes sense in the context of a particular theory, because, according to realism, value is defined by theoretical commitments. So, an intertheoretic absolute scale of value is senseless in the context of realism. But we can reach the same conclusion in the context of constructivism: just suppose that

Thomas *first* revises his judgement which *in turn* explain his revised theoretical commitments. I conclude that it's plausible our evaluative judgments reference an absolute scale of value in the context of constructivism.

### **Conclusion**

I have argued that our answer to the problem of moral uncertainty depends in part on our metaethics. In particular, I have argued that — at least within the context of cognitivism — intertheoretic comparisons are only *in principle* possible according to moral constructivism. This is good news for expected choice-worthiness as a decision theory under moral uncertainty. That being said, more work needs to be done to show exactly *how* to make constructivist intertheoretic comparisons.

On the other hand, this is bad news for moral realism. Of course, a moral realist might interpret my argument as reason to reject expected choice-worthiness. But one person's *modus ponens* is another's *modus tolens*. Therefore, if expected choice-worthiness is *prima facie* plausible, then it presents a challenge to moral realism.

### References

- Bhogal, H. (2020). Humeanism about laws of nature. *Philosophy Compass*, 15(8). <https://doi.org/10.1111/phc3.12696>
- Gracely, E. J. (1996). On the noncomparability of judgments made by different ethical theories. *Metaphilosophy*, 27(3), 327–332. <https://doi.org/10.1111/j.1467-9973.1996.tb00212.x>
- Harman, E. (2015). The Irrelevance of Moral Uncertainty. *Oxford Studies in Metaethics*, 10. <https://doi.org/10.1093/acprof:oso/9780198738695.003.0003>
- Lewis, D. (1987). A subjectivist's guide to objective chance. In D. Lewis, *Philosophical Papers Volume II* (1st ed., pp. 83–132). Oxford University Press New York. <https://doi.org/10.1093/0195036468.003.0004>
- MacAskill, W., Bykvist, K., & Ord, T. (2020). *Moral Uncertainty*. Oxford University Press. <https://doi.org/10.1093/oso/9780198722274.001.0001>
- MacAskill, W., & Ord, T. (2020). Why Maximize Expected Choice-Worthiness? . *Noûs*, 54(2), 327–353. <https://doi.org/10.1111/nous.12264>
- Street, S. (2006). A darwinian dilemma for realist theories of value. *Philosophical Studies*, 127(1), 109–166. <https://doi.org/10.1007/s11098-005-1726-6>
- Strevens, M. (2017). *Notes on Bayesian Confirmation Theory*. <http://www.strevens.org/bct/BCT.pdf>