

Duty and Doubt

SETH LAZAR

Australian National University

ABSTRACT

Deontologists have been slow to address decision-making under risk and uncertainty, no doubt because the standard approaches to non-moral decision theory appear superficially similar to consequentialist moral reasoning. I identify some central tenets of simple decision theory and show that they should not put deontologists off, before showing where we should go next to develop a comprehensive deontological decision theory.



1. INTRODUCTION

As any bad gardener knows, the symptoms of overwatering and underwatering are often the same. The leaves start to yellow, the stems lose stiffness and vigour. If you do nothing, the plant dies. But if you base your response on the wrong diagnosis, you actively accelerate that process. Sometimes you can do more to discover what is wrong before acting, but the evidence is often inconclusive.

Almost all human decision-making takes place under uncertainty. When the stakes are purely prudential—whether my camellia flourishes or dies—this is challenging enough. But moral decision-making is no less dependent on facts. And the facts are just as commonly in doubt.

Doctors, of course, know this. The injunction to ‘first do no harm’ is oblique at best. It can only plausibly be interpreted as ‘first, do no harm, all things considered, in light of the information you ought to have’. In countries vulnerable to wildfire, firefighters know this too. Many of Australia’s fiercest bushfires have resulted from ‘hazard reduction’ burns, intended to reduce danger to life and property, but which

would never have been started had those in charge known how the weather conditions would change.

And soldiers at war know it. Indeed, ‘the fog of war’ is so endemic as to be a cliché. Life and death decisions must be made at all junctures—often without knowing who will live and die, whether they are combatants or noncombatants, and indeed precisely what one is fighting for.

Although this combination of high stakes and high uncertainty is normally the province of emergency practitioners, those with more pedestrian lives also face risky decisions, albeit usually with less uncertainty. For example, many of us routinely operate lethal projectiles at high speeds, knowingly imposing a low risk of substantial harm on those we drive past, often for the sake of trivial gains (a trip to the shops for chocolate, perhaps). Some of us eat meat, running the risk of grievously wronging animals that might in fact enjoy moral status. We try to raise our children to be fulfilled and secure—for many of us no other obligation weighs heavier. But we are forever making consequential decisions with incomplete information, which might, instead, lead to the result that Philip Larkin predicted.

Practically all real-life moral reasoning takes place under uncertainty. That’s not news. And yet a sizeable wing of moral philosophy has failed to adequately acknowledge this. What’s more, it’s the wing that otherwise seems closest to the truth.

There are many different ways to understand ‘deontological’ ethics. The simplest is by defining the opposition: ‘consequentialism’. Consequentialists think, very roughly, that when choosing what we are morally permitted to do, we must consider the outcome that each available act will realise. To each outcome we assign a value; we are morally permitted to perform an act if, and only if, no other act has a better outcome.

By contrast, deontologists think that sometimes it is permissible to perform an act, even though an alternative would have a better outcome. In other words, sometimes doing less than the best is not morally wrong. This is most obviously true when the act with the better outcome involves excessive personal sacrifice, in relation to the good done.

Suppose I have the last ticket for the last performance of a play you and I both want to see. And suppose that you would enjoy it somewhat more than I would. Hold everything else constant. The outcome where I give you the ticket is somewhat better than the one where I view the play myself: pleasure is good, and you’ll get more pleasure from the play than I will. By consequentialist lights, I am morally required to

give you the ticket. Deontologists will balk at that result. We could all make the world a much better place, if we had no regard for our own well-being. It might be heroic, or saintly, to sacrifice ourselves for the sake of marginally greater benefits for others. But doing so is not, in general, morally required.

Deontologists also think that sometimes it is wrong to bring about the best outcome (on a natural understanding of ‘best’). Suppose that you could ensure that more children were lovingly cared for by their parents, by neglecting your own children. Or that you can minimise promise-breaking by others by breaking your promise now. Or that if you kill one innocent person, you can use their organs to save five other innocent people’s lives. In each case, deontologists think that you are morally required not to bring about the outcome that is (apparently) best. It’s not okay to minimise breaches of a given duty overall; you must attend to your own wrongdoing first and foremost.

I’ll say more about the deontology/consequentialism divide in the next section. But here is the key difference. Consequentialism is in one sense more demanding than deontology—there is no scope to prioritise your own interests. But it is also much more permissive: it licenses imposing costs on some if you can thereby realise a slightly larger benefit for others.

There are many deontologists in the world (I am among them). But the other side has, as is often remarked, one apparent advantage. Consequentialism pairs very nicely with standard ‘decision theory’—the theory of how to rationally make decisions with imperfect information. Deontologists, on the other hand, have offered no systematic account of how to handle decision-making under doubt.¹ That’s my project: to develop a theory of duty under doubt or, in other words, a *deontological decision theory*.

2. WHAT IS DEONTOLOGY?

It will help, first, to flesh out what the ‘deontological’ in deontological decision theory means. Some philosophers have offered subtle structural diagnoses of this ‘fault line in ethical theory’ (e.g. Nair 2014). But for my purposes here, I’m less inter-

1. Colyvan et al. (2010); Olsen (2018); Spector (2016); Zimmerman (2008) all fail to accommodate supererogation, options to act suboptimally, and in general the justifying/requiring distinction. Nor do they offer an argument for *why* deontologists should adopt expected-utility theory. Aboodi et al. (2008); Zollman et al. (2017) are intentionally very limited in scope. Even its author does not endorse Isaacs (2014). My own approach in this paper draws on work including Lazar (2017a, 2017b, 2018a, 2018b); Lazar and Lee-Stronach (2018).

ested in structural features of ethical theories, and more concerned with their substantive commitments. Here, then, is a simple deontological credo.

- beings with sufficiently advanced rational and moral capacities enjoy moral status;
- beings with moral status have fundamental rights, corresponding to duties that others have not to harm them, and to aid them;
- our duties not to harm others are more demanding than our duties to aid others (or to advance better outcomes in general);
- we have special duties to care for and help our loved ones and friends;
- the manner in which one's action harms or fails to aid someone matters for its permissibility, and our reasons, intentions and beliefs can likewise matter;
- we should be sceptical about justifications for the few bearing significant costs, which appeal to individually smaller, but collectively greater, benefits to others;
- our actions can wrong others, even when we act permissibly, and this 'pro tanto' wronging demands a response (perhaps apology, perhaps reparation);

This grab-bag of commitments is obviously not shared by all deontologists. What's more, plenty of consequentialists would endorse at least some of them. Nonetheless, consequentialists are much more likely to deny these claims, and deontologists to affirm them. My project is to develop a decision theory for *this* kind of deontology.

3. CRITERIA OF SUBJECTIVE PERMISSIBILITY

Except in special cases, decision theory does not aim to provide an operationalisable decision procedure. It aims, instead, to develop a criterion of subjective permissibility. What does that mean, and why should we want one?

We are interested in *moral* permissibility, and *moral* wrongness. (I use 'wrong' and 'impermissible' interchangeably.) Objective permissibility is permissibility in light of all the non-moral facts that could bear on your action. Think back to the

hazard reduction burn. If in fact the weather will change this afternoon, blowing a hot wind towards town, then it is objectively wrong to begin the burn now.

Subjective permissibility, by contrast, is permissibility in light of your imperfect information. This is an intentionally generic formulation, which covers a range of different ways in which your information could be imperfect. We might care, for instance, about whether an act is permissible in light of your beliefs about the world. Or we might care most about your evidence. There are other possibilities besides. In this paper, I will be neutral between them, focusing on questions that apply whichever subjective epistemic standard is at stake. Return to the firefighter. Suppose that, in the morning, there was no reason to believe that the wind would change in the afternoon. Then it may have been subjectively permissible to start the hazard reduction burn. Conversely, if a reliable forecast predicted the change, starting the burn would probably be subjectively impermissible.

A *criterion* of subjective permissibility is a set of necessary and sufficient conditions for an act being subjectively permissible. Ideally, these conditions would not only pick out subjectively permissible acts, it would explain why they are subjectively permissible. We're looking for conditions that state an act is subjectively permissible if and only if, *and because*, X, Y or Z.

This leads to the first substantive challenge to my project. Why should we want a criterion of subjective permissibility? One can understand this objection in different ways.

First, why bother with uncertainty at all? Why not let moral philosophers focus on decision-making under certainty, and leave decision-making under uncertainty to other branches of philosophy, indeed to other disciplines? Aren't these really technical questions about implementation, for which technical disciplines are necessary?

This objection might make sense if we were wholly clueless about subjective permissibility, lacking any intuitive insight into it. Assuming that we are more confident about our judgements of objective permissibility, we might then agree to outsource responsibility for decision-making under risk and uncertainty to a more technical discipline, anchored in our objective moral theory.

However, that is emphatically not where we stand. Our judgements of subjective permissibility are often as robust as our judgements of objective permissibility. For example, if the firefighter thinks it is more likely than not that the wind will change that afternoon, then without some compelling countervailing reason, it would be clearly wrong to start the burn in the morning. If we can draw conclusions like these,

then extending moral theory to decision-making under risk is not a purely technical task, and anyone engaged in it will be doing moral philosophy.

Other critics of criteria of subjective permissibility might argue that philosophers should focus on helping people *make* better decisions under uncertainty—developing actual decision procedures—rather than on criteria of right action (e.g. Smith 2018). My own background is in the ethics of war. I see the urgency of developing better decision procedures. But it is hubristic for philosophers to claim more than a bit-part in that process. Good decision-making involves much more than what philosophy alone can offer. Most notably, developing good decision procedures requires practical wisdom and experience—qualities that philosophers are not widely celebrated for possessing. Of course, we should contribute as best we can—and articulating criteria of subjective permissibility may be one of the most useful things we can do.

I agree, of course, that moral theory should be *action-guiding*. But this does not mean providing a formula that people can apply to actual decisions in realistic circumstances. Instead, it means *ranking* the options available to you, and saying which are permissible, in light of how they appeared to you at the time.

Others might reject ‘deontological decision theory’ just because it reduces the complex moral landscape to a few simple axioms. Deontologists are normally keen on drawing fine distinctions, and articulating complicated domain-specific principles. They may understandably feel uncomfortable with criteria of subjective permissibility, and instead prefer to accommodate risk in different domains of human activity with different principles (e.g. Bolinger 2017; Ferzan 2005; Quong 2015).

No doubt there is something to this idea. These domain-specific approaches, like deontological decision theory itself, are in their infancy. Both approaches should be developed, so that we can properly decide between them. At a first pass, though, I am sceptical about the domain-specific approach. However we divide up the relevant domains, we will find cases we cannot securely allocate to a specific domain. What principles do we apply then?

For example, suppose one set of principles governs beneficence under risk, while another set governs harming—that seems very likely. What do we do, then, if we don’t know whether our options will benefit or harm? Which domain-specific principles do we apply? I suspect that solving this problem in a principled way will lead inexorably to adopting a criterion of subjective permissibility.

A final observation: computer scientists and engineers are already building artificial agents to make consequential decisions. Artificial intelligence is fundamentally

grounded in probability theory and decision theory. Some of the reasons for scepticism that a human agent could ever actually deploy a criterion of subjective permissibility do not apply to artificial agents. They can gather and process much more information than humans can, at least within a narrow domain of expertise. They lack emotions, so cannot be blinded by fear or panic. A further reason to develop a deontological decision theory is that, if we don't, the AI agents of the future will all be consequentialists.

4. THE PROMISE AND PERILS OF ORTHODOX DECISION THEORY

Perhaps deontologists have ignored decision-making under uncertainty because a separate, interdisciplinary field of enquiry has already claimed that territory as its own. Economists, computer scientists, engineers, mathematicians, and others all contribute to, and make use of, various forms of decision theory. One might think that a group as smart as that would solve any problems that moral decision-making under risk might raise. And indeed deontologists have much to learn from decision theory. But there are some hurdles to overcome first.

Decision theory is, in essence, the attempt to devise a criterion of subjective permissibility for *rational* decision-making under risk. Vanilla decision theory says that an act is subjectively (rationally) permissible if and only if, and because, it maximises expected utility.² To understand this principle, we need some ingredients.

First, *acts*. An act is just something that it is in your control to do. It's easy to problematise this notion, but we'll work with a simple, intuitive sense of your available acts. For example, start the burn, or don't; fire at the person in your sights, or don't; and so on.

Second, *states*. A state is a complete description of the world (though to save time we can simply focus on the salient bits). Importantly, it is not just a description of how the world is at a time. It includes, at least in principle, the full history of the world too.

There are many ways the world might be. We do not know which one is actual. But some are more likely to be actual than others. *Probabilities* represent this idea.

2. Of course, this is as disputatious a field as any other in philosophy, so this characterisation of orthodox decision theory might itself be challenged. For two helpful overviews, see Briggs (2014); Buchak (2016a).

They raise many philosophical and mathematical questions, but we must ignore those here. They provide a measure from 0 to 1 of how likely a state is to be actual. On one appealing formulation: 1 is logically necessary (like a tautology); 0 is logically impossible (like a contradiction).

Economists often define acts as functions from states to outcomes. An outcome is what happens if you perform a specific act, given that a particular state of the world is the case. Again, outcomes are full descriptions of the world, including its history—in particular, including the fact that you performed that act.

To decide between acts, orthodox decision theory has us compare the outcomes that they might realise. We therefore need some way to represent which outcomes are better and worse than others. *Utilities* are the numbers that we assign to outcomes, to measure their level of normative support. The use of numbers here might make some nervous—numbers have properties that reasons might lack, such as precision and simple additivity. The numbers, however, are just a way of comparing the different possible outcomes. Those comparisons are what matters, not the specific unit or scale in which we are measuring. And the use of numbers to represent a comparison is consistent with denying that our reasons have other properties associated with numbers. Some might also be put off by the word ‘utility’, invoking as it does the paradigmatic consequentialist theory, utilitarianism. That would be premature. The concept of ‘utility’ is just a measure of normative support. As we will see below, it is flexible enough to accommodate non-consequentialist theories of what matters.

With these ingredients in place, we can define the *expected utility* of our available acts. Consider the decision whether to start the hazard reduction burn. To calculate the expected utility of doing so, we have to determine the possible outcomes of this choice, and how likely they are. To keep things simple, let’s imagine that there are only two relevant states: either the weather conditions will remain as they are, or else they will change unfavourably. If they remain as they are, the burn will proceed without a hitch. If they change unfavourably, then a catastrophic fire will ensue. To calculate the expected utility of undertaking the burn, we must assign utilities to these two outcomes, and then multiply each by the probability of the associated state being the case. We then sum the products to get the expected utility, which is a probability-weighted average of the utilities of the possible outcomes of your act. Orthodox decision theory says that it is rationally permissible to start the fire only if no alternative has greater expected utility.

This simplistic overview tells you how orthodox decision theory selects the per-

missible option from a risky choice. One question for deontologists, then, is whether they can adapt this approach to accommodate their view of what matters. However, there is a more pressing concern. Much—indeed most—work in decision theory focuses not on a criterion of right action, but on what justifies it. And the most commonly cited justifications seem ill-suited to a deontological approach to ethics. Before we ask whether deontologists *can* adopt and adapt orthodox decision theory, we must ask if they should want to.

There are three popular arguments for maximising expected utility. They are mathematical, and complex, so I will only sketch them here.

First, the most obvious and commonplace argument for maximising expected utility is that, if you do so, then in the long run and on average, you will realise more *actual* utility. Maximising expected utility in the short run is a great way to get what you want in the long run.

Second, many decision theorists think that one ought to maximise expected utility because any agent whose preferences over gambles (that is, acts with uncertain outcomes) obey some seemingly innocent ‘axioms of rationality’ can be mathematically *represented* as an expected-utility maximiser. The normative justification for their decision rule, then, is the plausibility of these underlying axioms. Indeed, many decision theorists would think it unbearably jejune even to suppose that they aim to defend a criterion of right action, and in particular to imagine that we can ever identify which act to choose under risk by assigning probabilities to states and utilities to outcomes. They argue that one’s preferences over gambles are the only real psychological phenomenon—probabilities and utilities are a construct generated by the ‘representation theorem’ (we can call these people ‘constructivists’). They would argue that orthodox decision theory gives no guidance on how to choose between risky options—only requiring that you obey the axioms of rationality.

Third, Graham Oddie and Peter Milne argue, roughly, that an additional 0.01 probability of some outcome occurring should have the same bearing on your decision regardless of whether the probability of that outcome has gone from 0.9 to 0.91, or from 0.1 to 0.11 (Oddie and Milne 1991). They then show that, if this premise is correct, we are committed to expected-utility maximisation (the details don’t matter here).

Notwithstanding the scepticism of constructivists, orthodox decision theory is more than just an elegant bit of maths. It underpins decision-making in many spheres of human—and non-human—endeavour, from governments to insurance compa-

nies, banks to casinos, self-driving vehicles to stock-trading algorithms. It is core to many different academic disciplines, from the behavioural sciences to some branches of physics. It has well earned its title as orthodoxy.

At first sight, this looks like bad news for deontologists. Let's start with the justifications. The first is the most obviously problematic. Justifying one's decisions now because of the long-run implications of deciding in that way is emphatically *not* the deontological approach! We are bound by constraints, which means that sometimes the ends *don't* justify the means. What's more, we've also seen that the ends don't necessarily *require* the means. If you are permitted to act suboptimally in individual choices, then there is no guarantee that you will realise the best outcomes in the long run. We must start by binning the justification grounded in the long run.³

Alas, the representation theorem justification is hardly more promising. First, constructivist decision theory clearly fails to be action-guiding. It cannot tell you which of your options maximises expected utility until you have made your choice. It can do no more than rule out options that would violate an axiom of rationality. But *moral* decision theory will need to be more robust than that. The axioms of rationality are merely formal constraints, demanding, in essence, nothing more than consistency. It is perfectly possible to be a consistent moral monster. Constructivist decision theorists harrumphing about probabilities and utilities not being real have abandoned normative theory altogether. Decision theory can—and must—do more than tell us how to be consistent.

Worse still, a deontological approach to ethics sits particularly ill with these supposedly innocent axioms. For example, *transitivity* means, roughly, that for any three options A, B and C, if A is preferred to B, and B to C, then A should be preferred to C. Deontological decision-making under risk is quite likely to violate transitivity. Agents typically have a range of permissible options, and the 'permissible to do X rather than Y' relation is emphatically not transitive, as Frances Kamm has shown. For example, it is (A) permissible to go golfing rather than (B) save a drowning child at high personal risk. It is permissible to (B) save a drowning child at high personal risk rather than (C) keep an appointment. But it is *not* permissible to (A) go golfing rather than (C) keep an appointment (Kamm 1985). Others have shown that the same is true of moral requirements, and indeed perhaps moral goodness (Temkin 2012; Voorhoeve 2014).

If deontologists reject transitivity, then they cannot be represented as expected

3. As many others have argued, e.g. in this context Tenenbaum (2017).

utility maximisers. They might also plausibly reject *continuity*, which roughly says that you should be prepared to trade off any sure thing against some probability of any other thing. And *completeness* is equally controversial, since it implies that we can, in principle, rank every option in every gamble against every other one—which is impossible if some options are incomparable to others.

Deontologists cannot reverse into orthodox decision theory via a representation theorem. They also can't justify adopting this approach on the grounds of its long-run success. So should they look elsewhere for inspiration? Is orthodox decision theory a dead end?

I think not. Oddie and Milne's argument, above, may point us in the right direction. It focuses attention on just what one is committed to, when one adopts the decision-theoretic framework. If we can pinpoint its fundamental elements, then perhaps we can either argue for those, or show that they cannot plausibly be controverted.

What is left of orthodox decision theory when we jettison long-run justifications and representation theorems? Suppose that we cannot *represent* you as an expected utility maximiser. You can still assign probabilities and utilities to states and outcomes, combine them to determine each act's expected utility, and then use those expected utilities to choose an action. Indeed, if you are going to make defensible decisions under risk, then you must find *some* way to combine a measure of your uncertainty with a measure of what matters. This is simply common sense: any risky choice involves an assessment of both what is at stake and what the odds are. The big question faced by deontologists is whether they should combine these in the way recommended by orthodox decision theory.

In this spirit, I think we can extract two basic ingredients from orthodox decision theory. Representation theorems and the maximising decision rule are not part of the core package. The core package is two simple theses, which I think deontologists can (with a little interpretation) adopt.

First, we can sensibly determine the level of normative support of the possible outcomes of our actions, and rank them accordingly, on one or more dimensions: I'll call this the *Ranking* hypothesis. Second, when deciding under uncertainty, we should discount an outcome's level of normative support (in some dimension) in linear proportion to its probability of coming about: I'll call this thesis *Linear Discounting*.⁴

For the project of developing a deontological decision theory to get off the

4. This is, I think, the same as Smith's 'Evaluative Compositionality': Smith (2014). See also Tenenbaum's description of the 'multiplicative model': Tenenbaum (2017).

ground, we need only to argue that *Ranking* and *Linear Discounting* are consistent with deontological ethics, and that we have positive reason to adopt them rather than some alternative.⁵

4.1 RANKING

Deontologists should not be deterred by *Ranking*. It requires only that, when figuring out whether to perform some act, the reasons for and against it, *within a particular dimension*, are roughly comparable to one another, so that we can somehow aggregate them. That is all. It does not presuppose that we can always *precisely* compare our reasons. Most importantly, it does not presuppose that our reasons all operate within the same normative dimension or domain. Vanilla decision theory's decision rule *does* presuppose that all reasons weigh on a single dimension. This explains why, according to the orthodoxy, it is irrational to choose a worse option when a better is available. But *Ranking* is consistent with endorsing a different decision rule, which registers different dimensions of normative strength (I'll explain this in more detail below).

Ranking does not assume that reasons interact with one another in a simple additive way. Some people think that if a morally relevant property tells for or against your action to a particular degree in one choice, it must do the same in any choice. On this view, a fact's reason-giving force is much like an object's mass in a controlled environment: if you add two properties together, you get the sum of their normative weights when separate. *Ranking* is consistent with this view, but also with its counterpoint—the idea that morally relevant properties can interact in unusual ways beyond mere addition, for example cancelling one another out or amplifying one another's force (e.g. Dancy 2004).

Ranking is conservative, but not vacuous. It insists that our reasons are comparable. For any two considerations, we can judge that one is weightier than the other, or they are equally weighty or roughly equal. If two considerations are incompara-

5. Notice that one of the main precursors to this paper, Colyvan *et al.* (2010), doesn't offer any argument to justify adopting orthodox decision theory and, in particular, does not defend either *Ranking* or *Linear Discounting*. Instead, it assumes that orthodox decision theory is the way to go, and takes the main challenge to be working out whether a utility function can adequately represent deontological moral theories. What's more, it does this without attending to the most challenging feature of deontological ethics—its commitment to supererogation and other options to act suboptimally. Oddie and Milne (1991) argue that all deontologists can be represented as expected-utility maximisers. They are wrong, again because of the failure to accommodate options to act suboptimally, which clearly lead to violations of transitivity. As far as I can make out, though Spector (2016) and Olsen (2018) endorse *Linear Discounting*, they offer no argument for it.

ble, none of these relations between them obtains. I think deontologists (and others) should enthusiastically give incomparability the boot. What's more, risky choices make this conclusion all the more compelling.

Here's the argument against incomparability. Suppose that two reasons are incomparable with one another: say, your reason not to kill an innocent person and your reason to preserve the natural beauty of the environment. Suppose that your parliament is about to vote in favour of allowing resource extraction in a national park. Your only way of preventing this from happening is to detonate an explosive, which will trigger a sequence of events leading (let's assume) to the parliament reversing its course and protecting the park. The downside is that if you detonate the explosive, you will definitely kill an innocent person. Simplifying a lot, suppose that the only considerations at stake here are the preservation of the environment's natural beauty and the innocent life that you will be taking. If these considerations are incomparable with one another, then there is simply no way to rationally weigh them. We cannot say that one is more important than the other, or that they are (roughly) equal.

Now this is hard to swallow, just on its face. But notice the further implications of endorsing incomparability. In any actual choice, there will be many other considerations at stake. Even if all of those are fully comparable, the presence on either side of incomparable considerations means that your options will be incomparable. That's already a problem (MacAskill 2013). Worse still, this problem arises even if there is a very low probability that incomparable considerations will be at stake. Suppose, for example, that the bomb is very unlikely to kill an innocent person, but sure to save the national park. One might naturally want to infer that the importance of saving the national park outweighs the risk of killing an innocent person, when that risk is low enough. But if these considerations are genuinely incomparable, then that inference is unwarranted.

Here's why. Let's use N as shorthand for natural beauty, and K as shorthand for the wrongness of killing an innocent person. If N and K are incomparable with one another, then any multiple of N and of K must also be incomparable with one another. If they were not, then we would be able to infer how N and K compare. Suppose, for example, that one wanted to say that, though N and K are incomparable, N is more important than $0.01(K)$. If that were true, then one could proceed by a series of comparisons, asking whether progressively larger probabilities of K are still less important than N , until we reach a point where they are roughly equal. If $p(K)$ is

roughly equal to N , then K must be at least roughly equal to N , if not more important than it. So if $p(K)$ is comparable with N , then K is comparable with N .

If you think that some reasons are incomparable under certainty, then you must also believe them incomparable under risk. And if the presence of any incomparable considerations makes it impossible to rationally choose between otherwise comparable options, then whenever there is any risk of a choice involving incomparable considerations, the options will themselves be incomparable. The result would be total paralysis. Incomparability, as Will MacAskill has nicely put it, is infectious (MacAskill 2013). That is not a tolerable or plausible result. Deontologists should reject incomparability, and endorse *Ranking*.

4.2 LINEAR DISCOUNTING

Nothing about discounting our reasons in proportion to their probability of being actual conflicts with core deontological commitments. If we should reject *Linear Discounting*, then, it is not because we are deontologists, but because there are independent reasons to reject it.

One might argue that probabilities are not the only tools for navigating an uncertain world. Most notably, we can distinguish between probabilities on the one hand, and ‘all-out beliefs’ on the other (Isaacs 2014; Tenenbaum 2017). We find something like this in the criminal law. Juries are instructed to convict if they *believe, beyond a reasonable doubt*, that the suspect is guilty.⁶ When things go wrong—like for the firefighter whose hazard reduction fire burns out of control—we often ask whether someone *believed* they were acting objectively permissibly, without regard to probabilities. Perhaps we could extend this kind of model to other cases, and so do without *Linear Discounting*?

I cannot debate the merits of these different approaches in this paper; nor should I have to. The prominent role of probabilities in philosophy, the natural and social sciences, and social life more generally, is indisputable. Since Cicero, people have agreed that ‘probability is the very guide to life’. Perhaps other guides are available. But it is at least defensible to base one’s decision theory on the guide that has the most currency.

There is a second line of argument against *Linear Discounting*, which we cannot so easily dismiss. One might agree that moral decision-making under doubt should

6. For a helpful overview, see Gardiner (2017).

use probabilities, but argue that probabilities should discount utilities in a different way. Perhaps we should totally disregard very low-probability states (Smith 2014). Or else perhaps we should give additional weight to very high-probability states, treating them as equivalent to certainty.⁷ Or we might give disproportionate weight to crossing the 50:50 threshold (Haque 2012). Alternatively, we might give some weight to ‘global’ properties of gambles, caring about the distance between the worst-case and best-case outcome, or else giving extra weight in our deliberations to one of those extremes. All of these and others could justify rejecting *Linear Discounting*. What can be said against them?

It is surprisingly hard to justify *Linear Discounting*. Oddie and Milne, as pointed out above, thought it just obviously true. They asked, incredulously, what could possibly make an additional 0.01 probability of some outcome matter more when added to 0.1 than when added to 0.9 probability of the same outcome. But this is really just table-thumping. I’m going to offer two arguments, which I hope will be more convincing.

The first raises a worry about attaching special significance to crossing particular probability thresholds—discounting low-probability outcomes, for example, or giving special weight to the 50:50 threshold. This approach makes the expected utility of an option depend acutely on how you define the relevant states. For example, suppose that the firefighter is permitted to ignore states with a lower than 0.01 probability. Before, I described the relevant states as being ‘weather remains the same’ or ‘weather changes unfavourably’. But we could equally well divide the latter possibility into arbitrarily many sub-possibilities (the wind increases by n km/h and shifts to the west, the temperature increases by n degrees, etc). What stops us dividing up the states until they fall below the threshold of indifference, and can safely be disregarded? If there is no obviously privileged way to partition the relevant states, how do we calculate the expected utility of an option?

My second argument gives a general reason to favour *Linear Discounting* over the alternatives. Our goal is to determine how to discount utilities by their probability of being actual. The product of that discounting is literally a product—the result of multiplying one factor by another. Any product that you can reach by tweaking one of the factors (for example, by altering the probability discount) you could instead reach by tweaking the other (by building it into the utility of the associated outcome). The question, then, is not whether any practical verdicts warrant a departure from

7. On the idea of a ‘moral certainty’, see Aboodi *et al.* (2008).

Linear Discounting. Instead, the question is whether those verdicts are best illuminated and explained by rejecting *Linear Discounting*, or else by tweaking the utility function.

Our deontic verdicts—what we are permitted to do in a given choice—are determined by the product of our probabilities and our utilities. We can reach precisely the same set of deontic verdicts with a decision theory that endorses *Linear Discounting* as we can with one that rejects it. The question is: which approach is better motivated and has more explanatory power?

Here is my simple thought: expected utilities combine epistemic and practical reasons. Probabilities are the epistemic component; utilities are the practical component. *Linear Discounting* is the default starting point. Probabilities share out the possibility of being actual in the way mandated by the evidence, or the agent's beliefs. *Linear Discounting* shares out the weight of possible utilities in just the same way. Your reasons in a state get to count in proportion to their possibility of being actual. I think that any departure from *Linear Discounting* should be justified on epistemic grounds. If the motivation for abandoning *Linear Discounting* instead comes down to one's sense of the relevant practical reasons, then it should be incorporated within one's utilities.

Here's an example. I think that when you breach some duties, your act is objectively worse (other things equal) the greater your degree of belief, when you acted, that your action would breach that duty. If the probability of your act breaching that duty was very high when you acted, then your act is objectively worse than if it was very unlikely to lead to that breach. So, when a combatant kills an innocent person in war, that killing is objectively more wrongful the likelier it was, when she acted, that the victim would be innocent (Lazar 2015). The reason: taking a bigger risk of breaching one's duty to another involves, other things equal, showing them a greater degree of disrespect. One *could* incorporate this into one's theory by proposing a departure from *Linear Discounting*. But this is fundamentally a practical reason—having to do with the respect that we owe to others—not an epistemic one. The proper place to figure it into our deliberations is among our practical reasons.

Here is another way to see the same point. Probabilities interact with our objective practical reasons in different ways. For example, our duties of rescue are fundamentally epistemically constrained, in this sense: the duties I have depend in part on the information that I have. There is someone in the world, right now, who will die in the next five minutes. I could save their life by phoning to warn them of the

impending threat. The cost to me would be negligible; the benefit to them would be their life. But I do not have a duty of rescue, because there is no way for me to know whom to call. Conversely, I think our duties not to harm others are not epistemically constrained in the same way. If I kill someone accidentally, and there was no way for me to know that my action would cause their death, I think I have still breached a duty to that person (albeit that I may be fully excused) (Thomson 1986).

If this is right, then any departure from *Linear Discounting* will have to be highly context-specific. It will depend, indeed, on the precise contours of the outcomes realised by one's options. This is in contrast, for example, with Lara Buchak's 'Risk-Weighted Expected Utility' Theory, which assumes that our risk attitudes can be represented by a constant function, which does not vary with context.⁸ If the probability discount is being determined by the content of the outcomes anyway, then Occam's razor suggests we should just fold this information into the utility assignment.

Deontologists have no special reason to reject *Ranking* or *Linear Discounting*. Nor are there general reasons to reject either. *Ranking* is a central tenet of practical reasoning generally; without it we would be quite lost; and any apparent challenges to *Linear Discounting* are better catered for by recognising that the objective weight of our moral reasons can vary depending on the information that the agent has when she acts. So deontologists have good reason to adopt vanilla decision theory—if they can.

5. CAN DEONTOLOGISTS USE ORTHODOX DECISION THEORY?

Headline: deontologists can adopt the concept of expected utility, but had better reject the vanilla decision rule. Some will balk at the first point. Doesn't the concept of expected utility require us to *evaluate outcomes*? Isn't the whole point of deontological ethics that it doesn't go in for that?

Some philosophers certainly take this view. They have been riled by others who have sought to 'consequentialise' moral theories—not only deriving a *ranking* on outcomes from every plausible moral theory, but pairing that *ranking* with a maximising decision rule, so that in some sense any moral theory can be represented by a conse-

8. Buchak (2013). Buchak's qualified departure from *Linear Discounting* is, I think, the most promising alternative to it. She argues that we have a distinct doxastic attitude—an attitude to risk—which is most felicitously captured by introducing a separate factor into the expected utility calculation. Although I disagree with Buchak's conclusion, I think this is the right way to approach this question.

quentialist counterpart.⁹ I think that consequentialisers are probably right that any moral theory can be represented as giving a ranking on outcomes. I think they're wrong that we can represent all plausible views with a maximising decision rule—this is true, at least, if we aim to represent those theories not only for full information cases, but also for decision-making under risk (Lazar 2017a).

However, the detour into consequentialising may be unnecessary here. The simple presentation of decision theory that I gave above said that we should evaluate outcomes, and, for consistency, I stated *Ranking* and *Linear Probability Discount* in those terms. But outcomes are not truly fundamental to decision theory. What really matters is the level of normative support an act has, along some dimension, given that some state of the world is the case. It helps to understand this in terms of assigning utilities to the outcome of the act, given that state. But we could just as readily ask about the strength of reason in favour of that act, given that state. Deontologists should be perfectly comfortable saying whether one act is more or less supported by reason than another, holding the state constant, or else comparing act–state pairs. Indeed, if we can't say this, then I don't know how we can explain our choices between options.

Deontologists might still be worried about this framing, however. Doesn't it imply that I should be prepared to commit one *prima facie* wrongful act-type, if by doing so I can prevent others from performing a greater number of such act-types? This risk can be catered for in a number of ways. My preferred approach is to recognise that the reasons for an act can be agent-relative as well as agent-neutral.

Agent-relative reasons either apply, or have particular force, for specific agents. Agent-neutral reasons apply universally, and have the same force for everyone. So, if my son is drowning, then everyone has a strong agent-neutral reason to save him. But I also have a more powerful agent-relative reason to do so. As is now well established, by recognising agent-relative reasons, we can accommodate the core deontological idea that there are some constraints that we must not breach, even if our doing so would prevent more breaches of the same duty by others (McNaughton and Rawling 1995). We can also do justice to special duties grounded in our valuable relationships. I must attend, first and foremost, to my own duties, rather than to the goal of maximising duty-compliance overall. Once we take agent-relative reasons into account, breaching a duty to minimise duty breaches overall is not, in fact, your best option.

Agent-relative reasons allow us to account for *constraints*—the requirement not

9. For an overview of the debate, see Portmore (2009).

to bring about what seems on the face of it to be the best outcome. But deontologists are also committed to options to act suboptimally. How can vanilla decision theory accommodate those?

It cannot. I doubt whether we can make sense of options to act suboptimally without recognising, as alluded to above, that there are distinct dimensions of normative strength. Vanilla decision theory presupposes that there is only one dimension of normative strength, which in turn explains why it is committed to *maximising* expected utility. If there is only one dimension of normative strength, and one option is stronger than another on that dimension, then there is simply no rational explanation why one would choose the lesser option.¹⁰ By contrast, once we recognise that there are different dimensions of normative strength, this result is entirely predictable.

As with everything, there is some disagreement about precisely how to understand the different dimensions of normative strength.¹¹ My approach is quite simplistic. Sometimes a reason counts in favour of an act, without making it the case that not performing that act would be wrongful. That's the justifying dimension of normative strength. Sometimes a reason counts in favour of an act, in such a way that if you do not perform the act, you would be acting impermissibly. That's the requiring dimension of normative strength.

To illustrate this difference, consider a trio of trolley cases. In the first scenario, the trolley is headed towards ten people, whom I can save only by diverting it towards two people on another track. I am permitted to turn the trolley.¹²

In the second scenario, there are two tracks, with two levers. I can pull only one lever. On track A, trolley A is headed towards ten people, whom it will kill unless I pull lever A. On track B, trolley B is headed towards my son, whom it will kill unless I pull lever B. Pulling either lever diverts the relevant trolley down a side track where it stops harmlessly. I am permitted to turn trolley B.

In the third scenario, the trolley is again headed towards my son, and I can save him only by diverting it down a side track where it will kill two people. I am not permitted to turn the trolley.

10. One might endorse 'satisficing', but one of the main objections to that approach is that it is straightforwardly irrational. There are other potent objections too: see Lazar (2017a).

11. The distinction was introduced by Joshua Gert in the context of practical rationality, but also applies to moral reasons. I disagree with Gert's understanding of the distinction in various respects, but will not explore them here. See Gert (2007).

12. This example adapts one that I introduced in Lazar (2013).

Let's assume that you share my judgements on these cases. If there were only one dimension of normative strength, then that would make us both incoherent—unless we argue, implausibly, that each case is simply very finely balanced, so you can go either way (we could fix that by changing the numbers—I think the verdicts would be robust over quite a lot of variation).¹³ If saving ten is more important than killing two, and saving my son is more important than saving ten, then saving my son should also be more important than killing two.

Once we recognise that normative strength varies along two dimensions, however, this apparent incoherence goes away. On one dimension, saving ten is morally more important than not killing two, because saving ten can justify killing two. On another dimension, not killing two is morally more important than saving ten, because you can be required to let your child die, if that is necessary to avoid killing two people, whereas you cannot be required to let your child die in order to save ten.

Vanilla decision theory has no room for these subtleties. It recognises only one dimension of normative strength. The only remedy is to change the decision rule and, in particular, to reject its simple maximising framework. There are many different possibilities here, and adopting one rather than another may involve endorsing a particular account of what explains the requiring/justifying distinction. I don't pretend to have nailed that distinction, but as a proof of concept, here is an attempt that can capture a considerable range of the landscape.

COST: An act is subjectively permissible if and only if, and because, either (a) there is no all things considered expectedly better act or (b) every all things considered expectedly better act either (i) involves unreasonable marginal expected costs to the agent or (ii) is better only in virtue of expected benefits to the agent.

The concept of an 'expectedly better act' maps directly on to what vanilla decision theory would describe as an expected-utility *ranking*. One act is expectedly better than another insofar as it realises more expected utility. We calculate an act's expected utility by considering all the possible agent-relative and agent-neutral reasons that would tell for and against that act, given the different possible ways the world might be, and then discounting them for the probability of their associated state coming about. The 'expectedly best' act is always permissible.

13. You could appeal to vagueness to reach the same verdict, but that too is just a dodge.

This already gives us the machinery to accommodate most of the deontological credo described above. The constraint against violating fundamental rights is captured by that reason's agent-relative dimension. The same is true for our special duties. The role of intentions, beliefs, and the causal structure of harms and benefits can likewise be accounted for within our agent-relative and agent-neutral reasons.

Deontological constraints would be relatively easy to accommodate even within vanilla decision theory. But, as we have seen, its maximising rule elides the requiring/justifying distinction, so fails to account for options to act suboptimally. COST restores those options, and in so doing takes a particular approach to understanding justifying versus requiring. In COST, the moral betterness ordering captures the justifying dimension of normative strength. The expectedly best option is *always* justified. COST then captures the requiring dimension of normative strength by reflecting on the special authority we have over our own interests. In clause b(i) it states that, while an expectedly better option may be justified, it is not required unless the additional moral benefit is great enough to make the additional cost to the agent *reasonable*. And clause b(ii) states that if the better option is better only in virtue of how it serves the interests of the agent, then the agent cannot be required to pursue it.

This gives a coherent and explanatory account of the requiring/justifying distinction, allowing deontologists to adopt the central lessons of vanilla decision theory without abandoning options to act suboptimally. Some deontologists, however, might reject this attempt to explain the distinction by appealing to personal cost, seeing the latter as just one symptom of the requiring/justifying distinction, rather than the whole story. They might, then, favour a more schematic decision rule, along these lines:

RJ: An option is subjectively permissible if and only if, and because, either (a) there are no probability-weighted requiring reasons not to do it, or (b) the probability-weighted requiring reasons not to do it are outweighed by the combination of probability-weighted justifying and requiring reasons in its favour.¹⁴

RJ is consistent with COST, but not exhausted by it. It could also license the view, for example, that in standard trolley cases one is merely permitted, not required,

14. In formulating this principle, I was greatly helped by discussions with Tom Hurka about pro tanto duties and pro tanto permissions. And in an excellent paper, Daniel Muñoz has independently reached a very similar principle, for decision making under certainty: Muñoz (2018).

to turn the trolley (saving the five lives might justify killing the one, but not make it required).¹⁵ It still draws on vanilla decision theory, insofar as it weights one's reasons in linear proportion to probability, and of course relies on the thesis that reasons can be weighed. It can still incorporate agent-relative and agent-neutral reasons, in just the same way COST can. Of course, RJ is more schematic than COST, and therefore less informative. But it does offer a proof of concept that one can extend orthodox decision theory to accommodate the key positions in the deontological credo, without doing any violence to the requiring/justifying distinction.

There remains at least one relatively common (though not uncontroversial) deontological precept that this criterion of subjective permissibility might not adequately accommodate. Plenty of deontologists think that it is wrong to trade off lives against headaches: there is no number of headaches that one could avert that would make it permissible to avert those headaches rather than save a single life.¹⁶ Extending this kind of view to risky choices is extremely difficult.

The main problem can be articulated quite simply: deontological hostility to aggregation has to do with *how we weigh our reasons*. If the more important class of reason is at stake, the less important class of reason cannot even be counted. But the decision-theoretic approach has us total the expected reasons in favour of each act, and those reasons must account for all the relevant ways the world might be. Some of the ways the world might be include scenarios in which the higher-weighted reasons are not, in fact, actual. In those scenarios, the lower-weighted reasons can count. We then have to discount those reasons for the probability of that scenario arising. But once we have done so, we have an expected utility like any other, and it goes into the mix, to weigh against the expected utility associated with the higher-weighted consideration.

Suppose, for example, that you have to choose between treating someone who is at risk of dying and treating some number of people at risk of suffering a headache. If the person will certainly die if you do nothing, then you have no reason to save those who may suffer a headache. But he might not die. So there are possible states of the world where he won't die if you treat those with headaches. But if that's right, then we can weigh the expected utility of treating the headaches against the expected

15. Thanks to Peter A. Graham for raising such cases with me. For further discussion of them, see our paper, Lazar and Graham (2019). For the contrary view about 'lesser evil options', see Frowe (2018).

16. For the state of the art, see Tomlin (2017); Voorhoeve (2014).

utility of saving the life, even though we are not allowed to weigh actual headaches against actual lives.

There are ways to fix this problem, with an exotic theory of the weight of our reasons, or by articulating a complicated principle that accommodates hostility to aggregation. I develop those in detail elsewhere (Lazar 2018a). An alternative option of course is to reject the deontological hostility to aggregation.

6. WHERE NEXT?

I think the prospects for deontological decision theory are good. We can accommodate all of the central deontological tenets within a relatively simple decision rule—indeed, one that is only a little more complex than vanilla decision theory. Anti-aggregationism may be a challenge to accommodate, but not impossible, and anti-aggregationism itself faces serious questions (Tomlin 2017). And we have positive reason to endorse this decision rule, since *Ranking* and *Linear Discounting* seem to be very well supported.

Of course, RJ and COST are not the only alternatives. RJ is more general, COST is more partisan. We could propose other alternatives that account for the requiring/justifying distinction differently. Much of the work would go into figuring out how to enumerate and weigh the many different reasons that might apply to our actions. But that, of course, is simply the inescapable task of moral philosophy in general. More pressing, perhaps, are four additional kinds of question that deontological decision theory—and indeed any moral decision theory—must address.

First, which probabilities are relevant for subjective permissibility? Perhaps it depends on what we're trying to do. Quite plausibly, we can care about all of the following: the agent's actual probabilities; what her probabilities would be if her beliefs were consistent; the probabilities justified by the evidence she actually had; and the probabilities justified by the evidence she could have had, had she done more research. I think that we hold people to high epistemic standards, and that we do so on moral grounds: the relevant probabilities are those that you would have if you did the morally appropriate research. But that might involve some problematic circularity (how do we work out what amount of research counts as morally appropriate?) (Zimmerman 2008), and it might also contravene my earlier argument that practical matters should be addressed through our practical reasons, rather than through the probabilities.

Second, how should moral philosophers approach attitudes to risk? For example, suppose a villain presents me with the following choice. I can either flip a coin or roll a die. If I flip the coin and it lands heads, he will kill Alf and Betty; if tails, he kills nobody. If I roll a die and it comes up even, then he will kill Alf; if odd, he will kill Betty. If I refuse, he will kill us all. If I am risk-averse, then I should give additional weight to the worst-case scenario. This means I should roll the die, because its worst outcome is better than the worst outcome from tossing the coin. If I am risk-seeking, then I should give more weight to the best-case scenario, and so flip the coin, since it's the only option with the possibility of saving everyone. If I am risk-neutral, then I should plausibly be indifferent between coin and die.¹⁷

Recent work in decision theory has seen the resurgence of non-expected-utility theory, including ambitious attempts to build non-neutral attitudes to risk into decision theory from the ground up.¹⁸ These philosophers tend to argue that it is sometimes permissible to adopt a non-neutral attitude to risk. That seems fine for rational decision theory. But if we are moral realists, then we must think that moral decision theory is more prescriptive. So, for any given choice, is there a morally appropriate attitude to risk that we should take?

One issue, which I have already addressed, is whether risk attitudes should be reflected in the utility, the probability-discount, or some third factor. But that is basically just a modelling choice. The more substantive question is just what attitude to risk one ought to take.

Scant work exists on this topic.¹⁹ One natural view is that when one's action affects only oneself, one may adopt whatever reasonable attitude to risk one chooses, but when it affects others, one must instead act on their attitude to risk (Altham 1983). Lara Buchak argues that, if they have different risk attitudes, we should defer to the most risk-averse (within the reasonable range) (Buchak 2016b). I'm not sure about this. I think the risk attitude we adopt should also depend on what the stakes are—we should be more risk-averse when the worst-case scenario is really bad.

Third, just as people often have non-neutral attitudes to risk, so they also tend to prefer to decide on what they know, rather than act on ambiguous or impoverished information (Bradley 2016). We often face situations where one option has relatively

17. This is actually not quite right; one could be risk-neutral but still care about the distribution of harms, so view the two options as different on those grounds.

18. See especially Buchak (2013).

19. Altham (1983); Buchak (2016b). It's also often easily confused with other distinct, but related topics—that's especially true in Spector (2016).

sharply defined probabilities, while another is much less precise; we tend to go for the more clear-cut alternative, even if the stakes of the other are higher.

For example, suppose you're a soldier deciding whether to go to war. You know that your friends are going, and that their lives will be at serious risk. You are a very talented soldier, so it is very likely that, if you go to war, you will at some point save some of your friends from suffering serious harm. But you also know that there's a chance the war is unjust, and that if it is unjust, then fighting is very seriously wrong, since every act of killing that you commit is the equivalent of murder (and can't be justified by the importance of saving your friends). You don't know enough, however, to assign any sharp probability to whether the war is just or unjust. It depends on information that neither you, nor anybody in the public, can find out. Many people would argue that you should give somewhat more weight to the outcomes about which you are more precisely confident (e.g. Betz 2016). You know that if you don't go, you will fail to save your friends from serious risks. That clear knowledge is worth something (Al-Najjar and Weinstein 2009; Voorhoeve et al. 2016).

The fourth topic is not, strictly speaking, one for moral decision theorists alone, but it does have special importance for us, and especially for deontological decision theorists. In many situations, if we assess our options in isolation from one another, we reach one verdict on their permissibility, but if we consider them as a set, we reach a different verdict (McClennen 1990). There are also interaction effects between choices within a series: sometimes whether it is permissible to perform an act depends either on what one did earlier, or on what one will do later. I have touched on this topic in a number of papers, arguing, for example, that sometimes moral 'sunk costs' should have a bearing on what it is permissible to do now (Lazar 2018C). I have also argued that whether we should assess individual options in isolation or as a campaign can depend on whether they are a necessary part of a series that is justified in the aggregate (Lazar and Lee-Stronach 2018). And I've considered cases in which one can seemingly be permitted to perform some suboptimal beneficent act when it is parcelled out into individual choices, but not when those choices are taken as a whole (Barry and Lazar 2018).

All of these cases suggest that moral philosophy needs a theory of dynamic choice—how to do the right thing over time, over a sequence of choices, not just one choice at a time. This is especially pressing for theories of moral decision making under risk, since choices with one risk profile when considered in isolation will have quite different risk profiles when taken as a sequence.

REFERENCES

- Aboodi, R., A. Borer, and D. Enoch (2008) 'Deontology, Individualism, and Uncertainty', *Journal of Philosophy*, 105/5: 259-72.
- Al-Najjar, N. I. and J. Weinstein (2009) 'The Ambiguity Aversion Literature: A Critical Assessment', *Economics and Philosophy*, 25/3: 249-84.
- Altham, J. E. J. (1983) 'Ethics of Risk', *Proceedings of the Aristotelian Society*, 84: 15 - 29.
- Barry, C. and S. Lazar (2018) 'Moral Options and Optimisation', Unpublished MS.
- Betz, A. (2016), *The Ethics of War and Friendship: The Moral Significance of Fellowship of Arms*, (University of Illinois, Chicago).
- Bolinger, R. J. (2017) 'Reasonable Mistakes and Regulative Norms: Racial Bias in Defensive Harm', *Journal of Political Philosophy*, 25/2: 196-217.
- Bradley, R. (2016), 'Ellsberg's Paradox and the Value of Chances', *Economics and Philosophy*, 32/2: 231-48.
- Briggs, R. (2014) 'Normative Theories of Rational Choice: Expected Utility', in *Stanford Encyclopaedia of Philosophy*, Edward Zalta (ed.).
- Buchak, L. (2013) *Risk and Rationality*, Oxford: Oxford University Press.
- (2016a) 'Decision Theory', in *Oxford Handbook of the Philosophy of Probability*, Alan Hájek and Christopher Hitchcock (ed.), Oxford: Oxford University Press.
- (2016b), 'Taking Risks Behind the Veil of Ignorance', *Ethics*, 127/3: 610-44.
- Colyvan, M., D. Cox, and K. Steele (2010) 'Modelling the Moral Dimension of Decisions', *Noûs*, 44/3: 503-29.
- Dancy, J. (2004) *Ethics without Principles*, Oxford: Clarendon Press.
- Ferzan, K. K. (2005) 'Justifying Self-Defense', *Law and Philosophy*, 24/6: 711-49.
- Frowe, H. (2018) 'Lesser-Evil Justifications for Harming: Why We're Required to Turn the Trolley', *The Philosophical Quarterly*, 68/272: 460-480.
- Gardiner, G. (2017) 'In Defence of Reasonable Doubt', *Journal of Applied Philosophy*, 34/2: 221-41.
- Gert, J. (2007) 'Normative Strength and the Balance of Reasons', *The Philosophical Review*, 116/4: 533-62.
- Haque, A. A. (2012) 'Killing in the Fog of War', *Southern California Law Review*, 86/1: 63-116.
- Isaacs, Y. (2014) 'Duty and Knowledge', *Philosophical Perspectives*, 28/1: 95-110.
- Kamm, F. M. (1985) 'Supererogation and Obligation', *Journal of Philosophy*, 82/3: 118-38.

- Lazar, S. (2013) 'Associative Duties and the Ethics of Killing in War', *Journal of Practical Ethics*, 1/1: 3-48.
- (2015) 'Risky Killing and the Ethics of War', *Ethics*, 126/1: 91-117.
- (2017a) 'Deontological Decision Theory and Agent-Centred Options', *Ethics*, 127/3: 579-609.
- (2017b) 'Anton's Game: Deontological Decision Theory for an Iterated Decision Problem', *Utilitas*, 29/1: 88-109.
- (2018a) 'Limited Aggregation and Risk', *Philosophy & Public Affairs*, 46/2: 117-59.
- (2018b) 'In Dubious Battle: Uncertainty and the Ethics of Killing', *Philosophical Studies*, 175/4: 859-83.
- (2018c) 'Moral Sunk Costs', *The Philosophical Quarterly*, 68/273: 841-61.
- Lazar, S. and C. Lee-Stronach (2018) 'Axiological Absolutism and Risk', *Noûs*, 53/1: 97-113.
- Lazar, S. and P. A. Graham (2019) 'Deontological Decision Theory and Lesser-Evil Options', *Synthese*.
- MacAskill, W. (2013) 'The Infectiousness of Nihilism', *Ethics*, 123/3: 508-20.
- McClellan, E. F. (1990) *Rationality and Dynamic Choice: Foundational Explorations*, Cambridge: Cambridge University Press.
- McNaughton, D. and P. Rawling (1995) 'Value and Agent-Relative Reasons', *Utilitas*, 7/1: 31-47.
- Muñoz, D. (2018) 'Better to Do Wrong', Unpublished MS.
- Nair, S. (2014) 'A Fault Line in Ethical Theory', *Philosophical Perspectives*, 28/1: 173-200.
- Oddie, G. and P. Milne (1991) 'Act and Value: Expectation and the Representability of Moral Theories', *Theoria*, 57/1-2: 42-76.
- Olsen, K. (2018) 'Subjective Rightness and Minimizing Expected Objective Wrongness', *Pacific Philosophical Quarterly*, 99/3: 417-41.
- Portmore, D. W. (2009) 'Consequentializing', *Philosophy Compass*, 4/2: 329-47.
- Quong, J. (2015) 'Rights against Harm', *Aristotelian Society Supplementary Volume*, 89/1: 249-66.
- Smith, H. (2018) *Making Morality Work*, Oxford: Oxford University Press.
- Smith, N. J. J. (2014) 'Is Evaluative Compositionality a Requirement of Rationality?', *Mind*, 123/490: 457-502.
- Spector, H. (2016) 'Decisional Nonconsequentialism and the Risk Sensitivity of Obligation', *Social Philosophy and Policy*, 32/2: 91-128.
- Temkin, L. S. (2012) *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*, Oxford: Oxford University Press.
- Tenenbaum, S. (2017) 'Action, Deontology, and Risk: Against the Multiplicative Model', *Ethics*, 127/3: 674-707.

Thomson, J. J. (1986) *Rights, Restitution, and Risk: Essays in Moral Theory*, Cambridge, Mass.: Harvard University Press.

Tomlin, P. (2017) 'On Limited Aggregation', *Philosophy & Public Affairs*, 45/3: 232-60.

Voorhoeve, A. (2014) 'How Should We Aggregate Competing Claims?', *Ethics*, 125/1: 64-87.

Voorhoeve, A., et al. (2016), 'Ambiguity Attitudes, Framing, and Consistency', *Theory and Decision*, 81/3: 313-37.

Zimmerman, M. J. (2008) *Living with Uncertainty: The Moral Significance of Ignorance*, Cambridge: Cambridge University Press.

Zollman, K. J. S., A. Bjorndahl, and A. J. London (2017) 'Kantian Decision Making under Uncertainty: Dignity, Price, and Consistency', *Philosophers' Imprint*, 17/7: 1-22.