

A paradox for tiny probabilities and enormous values

Nick Beckstead and Teruji Thomas

Global Priorities Institute | September 2020

GPI Working Paper No. 10-2020



A Paradox for Tiny Probabilities and Enormous Values

NICK BECKSTEAD AND TERUJI THOMAS*

Abstract

We show that every theory of the value of uncertain prospects must have one of three unpalatable properties. *Reckless* theories recommend risking arbitrarily great gains at arbitrarily long odds for the sake of enormous potential; *timid* theories recommend passing up arbitrarily great gains to prevent a tiny increase in risk; *non-transitive* theories deny the principle that, if A is better than B and B is better than C , then A must be better than C . While non-transitivity has been much discussed, we draw out the costs and benefits of recklessness and timidity when it comes to axiology, decision theory, and moral uncertainty.

We will argue that every theory about the value of prospects must have one of three unpalatable properties. By a *prospect* we mean a situation in which different outcomes can arise with different probabilities. For convenience, we will only consider prospects in which the possible outcomes are adequately described in terms of a *payoff*: the quantifiable gain or loss of some type of good relative to a fixed baseline outcome. Different theories of value will, of course, care about different types of goods. We will be most interested in a moral context, in which a typical payoff might be some number of people benefited in a certain way, or else some number

*Working paper, September 2020. This paper was originally drafted by the first author based on his dissertation (Beckstead, 2013, chapter 6). The second author is responsible for most of the revisions and ideas newly appearing in this version. We are grateful to many people for feedback and assistance, including Amanda Askell, Andreas Mogenssen, Christian Tarsney, Elliott Thornley, Hayden Wilkinson, Hilary Greaves, Larry Temkin, Philipp Schoenegger, Stefan Riedener, Theron Pummer, Tomi Francis, and Will MacAskill.

of good lives brought into existence; in a prudential context, it might be some number of years of happy life. At any rate, the first two horns of the trilemma are to be understood with respect to a given type of good and a given baseline.

The first horn of the trilemma is

Recklessness: For any finite payoff x , no matter how good, and any positive probability p , no matter how tiny, there's a finite payoff y , such that getting y with probability p is better than getting x for sure.

Unless otherwise specified, 'getting y with probability p ' will mean getting nothing (i.e. the baseline outcome) with the remaining probability $1 - p$. So, a theory that is reckless with respect to future years of happy life will say that the prospect of living 100 more years for sure is worse than a mere one-in-a-trillion chance of living for some finite (but perhaps truly vast) amount of time—and thus, all but certainly, dying right away.¹

The second horn is

Timidity: Let s be any positive number, no matter how small. There's a finite payoff x and a positive probability p , such that getting x with probability p is no worse than getting any other finite payoff, no matter how good, with a slightly smaller probability, just s percent less than p .

A theory that is timid with respect to lives saved will say that, for some n and p , saving n lives with probability p is no worse than saving an arbitrarily large number of lives with probability 0.999 -times- p (just 0.1% smaller).

The third horn is

Non-Transitivity: There are prospects A , B , and C , such that A is better B , B is better than C , but A is not better than C .

We think that recklessness and timidity are both intuitively implausible for many types of goods. Not simply resting on intuition, we will draw out the troubling consequences of each. In contrast, we will not consider in depth

¹While, as here, we often speak of 'chance' for convenience, we are sympathetic to the view that, strictly speaking, the relevant probabilities are epistemic. The distinction may be especially important if one thinks that chance-propositions can themselves bear value (Stefánsson and Bradley, 2015).

the possibility of rejecting transitivity; for a comprehensive discussion of it, we defer to Temkin (2012) and the large associated literature. The arguments in this paper could, then, be seen as supporting non-transitivity, although we are not inclined toward that conclusion.

While our focus is officially on theories of value, we'll sometimes talk about reckless or timid *agents* as well as theories. These agents form their preferences and choose their options in line with the recommendations of a reckless or timid theory of value. We do this to make recklessness and timidity more vivid; we're not committed to the view that, in general, one ought to choose the best prospect, although we're inclined to think that this is often true.

In section 1, we will clarify the three problems and explain why we must choose between them, for any type of good. We'll also introduce 'negative' versions of timidity and recklessness, having to do with bad payoffs rather than good ones. The remainder of the paper discusses problems associated with timid and reckless views, and their broader implications for decision theory and axiology.

In section 2, we consider some well-known approaches to the evaluation of prospects, and how they can lead to timidity or recklessness. The discussion includes expected value theory, risk-weighted expected utility theory, and a cleaned-up version of the 'Nicolausian discounting' view advocated by Smith (2014) and Monton (2019). For example, we note that the best-supported versions of utilitarianism, which specifically involve expected value theory, will lead to recklessness with respect to the creation of good lives. In general, our analysis of timidity and recklessness will be relevant to identifying the costs and benefits of many axiological theories.

In section 3, we argue that timid approaches have implications that will be implausible for many kinds of goods. For example, timidity means that events that happened in remote regions of space and time could be relevant to which actions would be best in prospect, even though we cannot affect what happens in those parts of space and time. Moreover, timidity tends to require implausibly extreme risk aversion, and even timid views will tend to recommend certain kinds of long-shot bets, with counterintuitive results.

In sections 4 and 5, we describe some problems that arise from recklessness when one allows for unbounded or infinite payoffs. Versions of these problems will be familiar to many readers in the context of expected utility theory, where they are associated with the St Petersburg gamble

and Pascal's wager, respectively. What's new here is that similar problems arise very generally from recklessness, regardless of whether one accepts expected utility theory. So these problems are more general and—given that the alternative is timidity—more pressing than usually thought. We also want to emphasize the importance of these problems in moral contexts, whereas discussions of St Petersburg and Pascal almost always focus on prudential value or mere preference.

Section 4 shows, more specifically, that recklessness conflicts with some very plausible dominance principles, while in section 5 we argue that reckless theories are *infinity-obsessed*: they evaluate an arbitrarily tiny probability of an infinite payoff above any finite payoff for sure. Although, for reasons we will explain, it is unclear how much difference this makes in ordinary situations, infinity obsession would lead to very strange decisions in some possible circumstances, and would greatly alter the grounds that we would ordinarily invoke to justify claims about whether one prospect is better than another.

In section 6, we argue that recklessness poses a problem for creating an acceptable theory of normative or evaluative uncertainty. How should your preferences reflect your uncertainty about which theory of value is correct? Many philosophers have argued in favor of an approach to this and similar questions that relies on expected utility theory. We argue that, under such an approach, agents who have any credence at all in recklessness must themselves be reckless. (As we'll explain, this is related to, but in some ways goes beyond standard worries about extreme theories 'swamping' more moderate ones.) Although we will mention some approaches to normative uncertainty that avoid this problem, there is a conundrum here for philosophers broadly sympathetic to the use of expected utility theory in this context.

I Why We Have to Choose

In this section, we argue that one has to choose between timidity, recklessness, and non-transitivity when evaluating prospects. A simple example illustrates the point.

Devil at Your Deathbed: On your deathbed, God hands you a ticket that can be delivered to any of his angels, good for an additional year of happy life. As you celebrate, the devil appears and asks you, 'Won't you accept a small risk to get

something vastly better? Trade that ticket for this one: it's good for 10 years of happy life, with probability 0.999.' You accept and the devil hands you a new ticket. But the devil asks again, 'Won't you accept a small risk to get something vastly better? Trade that ticket for this one: it is good for 100 years of happy life, with probability 0.999²—just 0.1% lower.' After 50,000 trades, you find yourself with a ticket for 10^{50,000} years of happy life that only works with probability 0.999^{50,000}, less than one chance in 10²¹.

Predictably, you die shortly thereafter.

In tabular form, the options look like this:

Deal	0	1	2	3	...	n	...	50,000
Payoff	1	10	100	1,000	...	10 ^{n}	...	10 ^{100,000}
Probability	1	0.999	0.998	0.997	...	0.999 ^{n}	...	< 10 ⁻²¹

On the one hand, each deal seems better than the one before. Accepting each deal immensely increases the payoff that's on the table (increasing the number of happy years by a factor of 10) while decreasing its probability by a tiny percentage (specifically, by 0.1%). It seems unreasonably timid to reject such a trade. On the other hand, it seems unreasonably reckless to take all of the deals—that would mean trading the certainty of a really valuable payoff for all but certainly no payoff at all.

Why is it impossible to avoid all the timid choices and all the reckless ones? With some qualifications to be discussed below, a theory that is *not* timid with respect to years of happy life will say that deal 1 is better than deal 0, deal 2 is better than deal 1, deal 3 is better than deal 2, ..., deal 49,999 is better than deal 49,998, and deal 50,000 is better than deal 49,999. If transitivity holds, it follows that deal 50,000 is better than deal 0. But that amounts to recklessness: an extremely small probability of a sufficiently large payoff is better than getting an excellent payoff for sure.²

² Arguments with a similar structure are often called *spectrum* or *continuum* arguments; see Temkin (2012) for a survey. The best known arguments involve tradeoffs between the quantity and the quality of a payoff (e.g. the duration of a life and its typical wellbeing at a time) rather than, as here, the quantity and the probability. Our argument has the useful feature that probability has an extremely well-established quantitative structure, with no natural threshold between 'high' and 'low' probability, whereas the structure of well-being

That's why we must choose between timidity, recklessness, and non-transitivity, when it comes to years of happy life; the argument obviously generalises to other types of goods. Of course, timidity, as we formulated it in the introduction, does not specifically claim that a 0.1% decrease in probability cannot be outweighed by a 10-fold increase in the payoff. Similarly, recklessness does not specifically claim that one year of happy life is worse than a 0.999^{50,000} probability of 10^{50,000} years of happy life. But to avoid timidity altogether, there must *some* deal at each step that would improve your prospects: there must be some percentage decrease (0.01%, or 0.000001%, or...) such that, at each step, it would be better to reduce the probability of the payoff by that percentage in exchange for a sufficiently large increase in its size. By stringing together a sufficiently large number of such deals, we can deduce that, for any positive probability p , no matter how small, it would be better to trade the original payoff for the prospect of some vast payoff with probability *at most* p . But then it would also be better to trade the original payoff for the prospect of a vast payoff with probability *exactly* p . That's recklessness.

Someone might embrace the timid approach on the grounds that, after enough time, additional years of happy life have essentially no value (perhaps because of boredom).³ We are skeptical, but in any case the problem recurs for other types of goods, for which the analogous claim is less plausible. We'll consider the problems with timidity more systematically in sections 2 and 3, but for now, impressionistically: if you think that (i) it is better for human civilization to flourish for more time rather than less, and that (ii) if human civilization flourishes for much longer, that would be much better, then timidity with respect to years of flourishing will seem unattractive. Similarly for timidity with respect to saving lives, if you think that (i) it is better to save more lives rather than less, and that (ii) if we save many, many more lives, that would be much better.⁴

(say) is more mysterious. Temkin discusses a probability-based spectrum argument in his chapter 8, but it works quite differently from ours. At any rate, the point of this paper is not merely to identify the spectrum argument, but to explain how it fits into standard theories of value and to give a sustained discussion of the costs and benefits of different responses.

³For a sophisticated view along these lines, see Williams (1973).

⁴Won't we, though, eventually run out of lives to save? There are different ways to finesse this, but the simplest is to consider a baseline outcome with an infinite population. In general, while we have presented timidity and recklessness as posing a dilemma relative to *each* baseline, one can also think of it as a dilemma between recklessness for *some*

In general, timidity is *prima facie* unattractive with respect to any good for which (i) more of that good seems better than less of that good, and (ii) a lot more seems a lot better. On the other hand, recklessness leads to its own problems, which we will detail in sections 4 and 5.

It's worth noting that we can construct an analogous trilemma using negative value. In this version, the payoffs might instead be years of miserable life, rather than years of happy life. A person who is 'reckless,' in the negative variant, prefers even a very long period of misery for sure to an arbitrarily small probability of a sufficiently long period of misery (e.g. preferring 1,000 years of suffering for sure to a one-in-a-trillion chance of suffering for some much longer period). More generally and precisely:

Negative recklessness: For any finite payoff x , no matter how *bad*, and any positive probability p , no matter how tiny, there's a finite payoff y , such that getting y with probability p is *worse* than getting x for certain.

Similarly, the negatively timid person passes up a deal that would give them a much shorter period of suffering but with slightly higher probability. More generally and precisely:

Negative timidity: Let s be any positive number, no matter how small. There's a finite payoff x and a positive probability p , such that getting x with probability p is no *better* than getting any other finite payoff, no matter how *bad*, with a slightly smaller probability, just s percent less than p .

We think that, for many sources of value, negative recklessness and negative timidity are roughly as counterintuitive as the original positive versions, and they will lead to analogous difficulties. But, if anything, our arguments will weigh especially heavily against negative timidity.

2 Examples

We now give examples of different ways of evaluating prospects, and how they lead to timidity or recklessness. We organize these examples around three different ideas that one might invoke to explain timidity, or to avoid

baseline and timidity for *all*. Indeed, recklessness with respect to even one baseline leads to some quite general theoretical problems, as explained in section 4.

recklessness: the boundedness of value, risk aversion, and discounting small probabilities. To be clear, this discussion isn't meant to exhaust the logical space. But it will include the best-credentialed normative theories of evaluation under uncertainty.

Throughout this section, we'll focus on prospects that are *simple* in that each one has only a finite number of possible outcomes, all corresponding to finite payoffs. We thereby rule out, to begin with, any difficulties that might arise from infinite payoffs or St Petersburg gambles of the sort we'll describe in sections 4 and 5.

2.1 The boundedness of value

As we already suggested in section 1, the most natural explanation for timidity is that the value of finite payoffs is bounded above. Deal n would be worse than deal $n + 1$ in sequences like Devil at Your Deathbed, if the payoff of deal n were already so good that it was near the upper limit of how valuable an additional number of happy years could be. Though the number of years available in deal $n + 1$ might be much greater, their *value* would not, and therefore wouldn't justify taking even a tiny additional risk. It would be strange to say (on the other hand) that the possibility of a much better outcome could not justify a slight increase in risk.

As we'll now argue, the connection between timidity and the boundedness of value is especially clear if we use *expected value theory* (EVT), the most common normative model for evaluating simple prospects.⁵ According to EVT, one prospect is better than another just in case it has higher *expected value*. Here, the expected value of a prospect is computed by (i) identifying all of the possible outcomes of the prospect, (ii) multiplying the probability of each outcome by its value, and (iii) adding all those terms together. So the expected value (relative to the relevant baseline) of

⁵EVT is a special case of expected *utility* theory, which we'll discuss in the next subsection. There are two different ways to understand EVT. On one understanding, there are independently given cardinal facts about value, i.e. not only facts about which outcomes are better than which others, but also facts about *how much* better they are; EVT then makes sense relative to a suitable quantitative representation of those facts. But some authors (e.g. Broome, 2004, p. 90) think that the question 'how much better' is at best ambiguous, and regard EVT as providing a possible disambiguation. We note that the intuitive explanation of timidity in terms of the boundedness of value *itself* requires some cardinal facts: we want to be able to say that, at some point, increasing the payoff makes things better but *not by much*.

getting a payoff x with probability p just equals p times the value of getting x for sure.

It is easy to see that, given EVT, recklessness corresponds exactly to the value of payoffs being unbounded above, in the following precise sense:

For any finite payoff x , and any number $n > 0$, there exists a finite payoff y whose value is at least n times greater than that of x .

Similarly, timidity means that value has an upper bound, in the precise sense that there's some x and some n such that no finite payoff is more than n times better than x . Perfectly analogous arguments show that negative timidity requires a lower bound on the value of finite payoffs, and negative recklessness requires that there is no such lower bound.

Given EVT, then, the connection between timidity and the boundedness of value is especially tight. But even if EVT is not exactly right, the boundedness of value would be a natural way to avoid recklessness. It provides a clear and, at least, *prima facie* plausible explanation of why some of the devil's deals might not be worthwhile.

We'll discuss some of the general problems with this idea in section 3, but for now we'll just point out that many substantive theories of value do *not* put bounds on the value of various interesting goods. For example, the most natural understanding of utilitarianism as an evaluative theory says that improving n lives by a given amount improves the world by n times as much as improving one life; so the value of improving lives is unbounded. *Total* utilitarianism and its variants likewise put unbounded value on creating good lives.⁶ And of course utilitarianism is not special here: giving priority to the badly off, for example, or introducing further dimensions (like equality) along which an outcome can be good or bad will not automatically place bounds on value. So there is ample motivation to consider other possible justifications for timidity.

2.2 Risk aversion

The very names 'timidity' and 'recklessness' suggest different attitudes towards risk. If value is not bounded, one might guess that risk aversion is

⁶What's more, some of the best worked-out arguments for total (or critical level) utilitarianism, like that of Broome (2004), specifically rely on expected value theory. One can even argue *for* expected value theory from total utilitarianism, or for expected utility theory with an unbounded utility function from key utilitarian principles like the Pareto principle and anonymity; see McCarthy et al. (2020, especially Theorem 4.10).

the proper way to account for timidity. As we'll now argue, timidity can indeed be explained by *extreme* levels of risk aversion on *some*, but not all, ways of theorizing about risk attitudes. In the end, the risk aversion required may be too extreme to be plausible; that's an issue we'll return to in section 3.

It is not straightforward to give an account of 'risk aversion' in an ordinary intuitive sense. However, according to a standard story, the characteristic feature of risk aversion is that one systematically judges that a sure payoff is better than an uncertain prospect with the same expected size: for example, it's better to save 10 lives for sure than to have a $\frac{1}{2}$ chance of saving 9 lives and a $\frac{1}{2}$ chance of saving 11. (This is, specifically, risk aversion with respect to the *size* of payoffs; we'll also discuss risk aversion with respect to their *value* below.) Risk seeking requires the opposite judgment; risk neutrality means that these prospects are equally good.

We'll illustrate the connection between timidity and extreme risk aversion using what we take to be the two most popular normative theories of risk attitudes.

First, *expected utility theory* (EUT) assigns a numerical 'utility' to each outcome, and prospects are compared on the basis of their expected utility. The utility of an outcome should depend only on the outcome's value, and better outcomes should have higher utility. Given the similarity between EUT and EVT (we'll explain the contrast shortly) it should come as no surprise that timidity is equivalent to the claim that the utility of finite payoffs is bounded above, and negative timidity to the claim that it is bounded below. The argument is the same as before, with 'utility' instead of 'value'. What does this have to do with risk aversion?

Risk aversion, in this framework, means that payoffs have decreasing marginal utility: the utility function increases less rapidly with each additional unit of payoff. It is, in other words, *concave*. Now, concavity does not entail that the utility function is bounded above, so risk aversion does not entail timidity. But the utility function will be bounded above if it is *very* concave, when it comes to good outcomes—that is, if the contribution to utility of each additional unit of payoff falls off sufficiently quickly. So timidity can be explained by relatively extreme levels of risk aversion when it comes to large payoffs. Risk seeking, on the other hand, means that the utility function is *convex*: it increases more rapidly with each additional unit of value. Negative timidity can be explained by relatively extreme levels of risk seeking when it comes to very negative payoffs. This

connection between negative timidity and risk seeking, rather than risk aversion, shouldn't be too surprising. Negative timidity suggests that, in some cases, a gamble with a relatively high expected payoff and a relatively small spread of payoffs is no worse than one with a relatively low expected payoff and a relatively high spread. For risk attitudes to explain this pattern, they must tend to favour high-spread prospects over low-spread ones. That is risk seeking.

How is this story different from the one about expected *value* theory in §2.1? First of all, EUT doesn't presuppose that there are any cardinal facts about value (cf. footnote 5). But even if there are such facts, the utility function need not straightforwardly reflect them. In particular, even if the value of finite payoffs has no upper bound, their utility may. This can happen if the utility of positive payoffs is very concave as a function of their value, not just as a function of their size. In that case, getting a positive payoff for sure will be better than any uncertain prospect with the same expected value. This is 'pure' risk aversion, risk aversion with respect to value itself. And similarly, even if the value of finite payoffs has no lower bound, negative timidity can arise through pure risk seeking, when it comes to negative payoffs.

The EUT account of risk aversion has often been criticized, so we would not want to be its hostage. Let us also consider the recently popular *risk-weighted expected utility theory* (REUT), developed by Buchak (2013) as a 'subjective' version of the anticipated utility theory of Quiggin (1982). According to this view, each payoff again as an associated utility, but risk attitudes are captured by a separate 'risk function' r , which is used to transform the probabilities. In the case of risk aversion, this transformation is designed to put high weight on the relatively bad outcomes of each prospect, amounting to a sort of pessimism; for risk seeking it puts high weight on the relatively good outcomes.⁷

⁷Specifically, if, according to some prospect P , $P_{>x}$ is the probability of getting an outcome better than x , and $P_{\geq x}$ is the probability of getting an outcome at least as good as x , then the risk-weighted expected utility of the prospect is

$$\sum_x [r(P_{\geq x}) - r(P_{>x})] \cdot u(x)$$

with respect to the risk function r and the utility function u . In the special case where $r(x) = x$, the term in square brackets is just the probability of getting an outcome as good as x , and the whole sum is nothing but the expected utility of P . But, in general, r is only required to be an increasing, real-valued function with $r(0) = 0$ and $r(1) = 1$.

Even if this is an attractive theory of risk attitudes, risk-weighted expected utility theory will not avoid recklessness unless utility is bounded above: it is no different from expected utility theory in that respect. The reason is that the risk-weighted expected utility of getting x with probability p is still just some positive number times the utility of x . Because of the risk function, that number is no longer simply p , but that doesn't change the fact that, if the utility of x increases without bound, then so does the risk-weighted expected utility of getting x with probability p . So risk aversion in this sense is insufficient to avoid recklessness. Moreover, this time, the boundedness of the utility function cannot be explained by appeal to risk aversion: risk aversion is already supposed to be handled by the risk function r . A bound on the utility function is more naturally interpreted as a bound on value. So, according to the standard interpretation of this theory, not even extreme levels of risk aversion suffice for timidity.

There is, however, a relatively natural way to make risk-weighted expected utility theory *more* extreme, so that it does lead to timidity, even with an unbounded utility function. We'll discuss that tweak under the next heading.

2.3 Discounting small probabilities

A third strategy for avoiding recklessness is *Nicolausian discounting*. On common formulations of this view, one simply ignores outcomes whose probabilities are smaller than some threshold.⁸ So, when it comes to recklessness, one will simply end up ignoring the tiny probability of an enormous payoff. Indeed, a key motivation for Nicolausian discounting is precisely to avoid recklessness and similar phenomena.

Unfortunately, the common formulation of Nicolausian discounting is fatally flawed. But since *something* in this neighbourhood has often been considered *prima facie* plausible, we'll take the time to formulate a principle that avoids the most obvious pitfalls.

First, the problems. Consider the (not simple) prospect of getting x years of happy life, where the real number x will be chosen at random somewhere between 0 and 1. Then each specific payoff has probability

⁸See Smith (2014) and Monton (2019) for recent proponents of this very old view. Monton coins the name 'Nicolausian discounting' after Nicolaus Bernoulli, who, in a 1714 letter, suggested that 'cases that have a very small probability must be neglected and assumed to be zero' (in Monton's translation). The view is similar to the theory of '*de minimis* risks', for which see Peterson (2002).

zero. Ordinary and risk-weighted expected utility theory have ways of dealing with this phenomenon (roughly, one replaces sums with integrals), but applying Nicolausian discounting in this case is completely hopeless (we can't just ignore *every* outcome). Relatedly, whether or not any outcomes fall below the probability threshold will depend on how finely we individuate them. Suppose for the sake of concreteness that the probability threshold is one-in-a-billion. If the devil offers us one year of happy life, should we evaluate the prospect as if it has exactly one payoff with probability one, or as if it has a trillion slightly different payoffs, all below the probability threshold? Nicolausian discounting, insofar as it makes sense at all, must implausibly claim that this kind of redescription matters a lot.⁹ Finally, suppose that one prospect is very similar to another, except that it will turn out much better in certain cases, which individually have probabilities below the threshold. The first prospect is clearly better, but Nicolausian discounting denies it.

Here's a solution. Instead of giving all small-probability outcomes zero weight, give *infinitesimal* weight to outcomes that are unusually extreme in value.¹⁰ We will call this type of view *tail discounting*. Here is a way to cash it out. There is some probability threshold ϵ . An outcome x is in the *left tail* of a given prospect if the probability of getting x -or-worse is at most ϵ . Similarly, x is in the *right tail* if the probability of getting x -or-better is at most ϵ . Say that an outcome is *extreme* with respect to the given prospect if it is in the left tail or the right tail. Otherwise, say that x is *normal*. On the simplest form of tail discounting, one prospect is better than another if the expected value of its normal outcomes is higher; but if the two prospects have the same expected value for normal outcomes, then use the expected value of extreme outcomes as a tie-breaker. Roughly, then, the view is that one should ignore the tails of a distribution, except for breaking ties. (If one is only interested in avoiding the positive form of recklessness, then one can ignore the right tail only.)

Tail discounting may be more defensible than Nicolausian discounting as commonly formulated. It does not suffer from the problems described

⁹See Gustafsson and Torpman (2014) for an attempt to find a canonical individuation of states, in a rather different context.

¹⁰Buchak (2013, pp. 73–74) makes (without endorsement) a similar suggestion in the context of St Petersburg gambles, which we'll consider in section 4; but her suggestion is to give extreme outcomes *zero* weight, which still leads to an analogue of the last problem mentioned above. At any rate, the problems we raise throughout this paper for tail discounting *also* apply to this simpler view.

above. And it can be seen as an extreme form of risk-weighted expected value theory, meaning that defences of the latter can, to some extent, be adapted to defend tail discounting as well.¹¹ Unfortunately, tail discounting leads to an especially extreme and implausible form of timidity (§3.3), along with some more general problems to which we'll now turn.

3 The Price of Timidity

We think that timidity is troubling on its face. For example, each deal offered by the devil really does seem better than the one before. But the problems don't stop there. In this section we explain some additional problems faced by timid theories. Some of these are completely general. Others relate to particular ways of implementing timidity that were sketched in section 2. Even if these latter problems may not be completely general, they illustrate the difficulty of constructing a plausible timid theory.

We'll give a brief summary at the end of the section.

3.1 Even small decreases in unlikely payoffs cannot be outweighed

In any instance of timidity, a small decrease in the probability of a payoff cannot be outweighed by any finite increase in the size of the payoff. One might think, 'This is not so implausible, because, after all, the payoff may have been enormous to begin with; a small decrease in the probability of an enormous payoff is still a weighty matter.' However, we can argue against timidity in a way that does not involve large decreases in the size of any possible payoff.

Consider, for the sake of concreteness, the prospect P of saving 1,000 lives with probability 0.1. A theory that is timid with respect to lives saved might say that P is no worse than the prospect of saving even a vast number of lives with probability 0.099, only one percent smaller. We can imagine that the outcome of P depends on the outcome of a raffle: if the golden

¹¹Specifically, this is a form of risk-weighted expected value theory in which the risk-function has infinitesimal values for probabilities below ϵ or above $1 - \epsilon$. For an infinitesimal number ι (e.g. a surreal or non-standard real number), one can choose the risk function r such that $r(x) = x\iota$ for $x \leq \epsilon$, $r(x) = 1 - (1 - x)\iota$ for $x \geq 1 - \epsilon$, and such that $r(x)$ increases linearly for x between ϵ and $1 - \epsilon$. Because of this connection, tail discounting satisfies the most important axioms of risk-weighted expected utility theory: dominance, comonotonic tradeoff consistency, and the related comonotonic sure thing principle, for which see Buchak (2013).

ticket is drawn, then a mechanism will be activated that will save 1,000 lives. Now consider a prospect P_1 that differs from P in two ways: first, the mechanism is designed to save many more lives—let’s say 10,000; but second, there is a one-percent failure rate, and in the fail-state, P_1 saves only 999 lives. So P_1 almost certainly leads to ten times as many lives saved as in P , and at worst, in the highly unlikely fail-state, it will save only one life fewer. Intuitively, P_1 is much better than P . But now consider a prospect P_2 . In P_2 , the mechanism usually saves 100,000, but only 998 in the fail-state. So, compared to P_1 , there are almost certainly 10 times as many lives saved, and at worst one fewer. P_2 is clearly better than P_1 , so also better than P . But if we continue this sequence along, we get to P_{1000} , in which no lives are saved in the fail-state. Thus P_{1000} leads to a vast number of lives saved (namely, 10^{1003}) with probability 0.099. Since P_{1000} is better than P , our theory of value must not be timid in this particular way. Nor is there anything special about this case; *mutatis mutandis*, we have an argument against timidity in general.

3.2 Strange dependence on distant space and time

Timidity requires our evaluation of prospects to depend in strange ways on what happens in distant places and times. At least, this is true insofar as the types of goods with respect to which we are timid can be realised far away.¹²

The problem is especially simple to understand if we suppose that timidity arises from the boundedness of value. How far away we are from the relevant baseline and how close we are to the upper bound can depend on how things are in remote regions of spacetime. In many situations, though, this would be deeply implausible. For example, if we had some way of greatly improving the lives of the roughly one billion people who now suffer from extreme global poverty, we could not plausibly claim that this would have only trivial value, on the grounds that a lot of great things had happened in the distant past or were now happening in other galaxies, so that we had already come close to the upper limit on how good improving lives could be.

¹²Perhaps, then, the point is less telling in the context of prudential value, but we leave this open. The argument to follow is closely related to well-known arguments from ‘separability’ for totalist views in population ethics (see Thomas, 2020b). However, the issue for us is not separability in general but the particularly strange violations of separability to which timidity leads.

What if value is not bounded, but timidity still holds? Then essentially the same problem appears in a more complicated guise. Suppose we would like to compare an option that would, relative to the status quo, significantly benefit a large number n of people with probability q , with another option that would benefit an even larger number N of people, to a similar extent, with a larger probability p . It seems obvious that the second option leads to a better prospect, especially if q is much smaller than p and N is much larger than n . However, if value is timid with respect to the number of people benefited, then whether this is true can depend on how the benefits happen to be correlated with what happens (and would have happened anyway) in far-away times and places.

To see this, remember that timidity is formulated relative to some baseline outcome. While it's usually a convenient assumption, this baseline need not be the practically relevant status quo. Consider, then the following prospects A, B, C . Relative to the baseline, A benefits $n+N$ people with probability p , while B benefits only n people with slightly larger probability $p+q$.

		PROBABILITIES		
		p	q	$1-p-q$
A	benefits	$n+N$	0	0
B		n	n	0
C		n	0	0

According to timidity, there will be some such case in which A is no better than B , even though q is one trillionth the size of p and N is a trillion times the size of n . But now suppose that, in evaluating our options, the relevant status quo is not the baseline considered so far, but rather the prospect C , in which n people may already have been benefited in a galaxy far away—we have no influence on what happens to those people. Relative to *that* status quo, prospect B corresponds to benefiting only n people with tiny probability q , while A corresponds to benefiting N people with a much larger probability p . So, implausibly, the option of benefiting n people (relative to the status quo) with probability q is not better in prospect than the option of benefiting N people with probability p .

It seems strange to us that the importance of a given probability of benefiting a given number of people depends in any significant way on what might happen, entirely beyond our influence, in far-away times and places.

But, more importantly, the specific implications of timidity in cases like the one we just described seem rather implausible. And similar problems follow from negative timidity, but we won't discuss them here.

3.3 Extreme risk aversion in very positive outcomes

The next problem is that timidity leads to implausibly extreme forms of risk aversion, at least according to all the timid theories we surveyed in section 2.

Consider first a view according to which there is an upper limit on how good finite payoffs of the relevant sort could be. We will get deeply implausible results when thinking about prospects that mix high-value and middling-value outcomes. For example, suppose that 10^{10} extra years of happy life would take you extremely close to the upper limit to how good extra years of life could be. Now consider two prospects, A and B , in which the outcome depends on the toss of a fair coin:

	Heads	Tails
A	10^{10} years of happy life	1 hour of misery
B	10^{80} years of happy life	2 hours of misery

Since 10^{10} years of happy life would take you extremely close to the upper limit, the additional 70 orders of magnitude of years of happy life would represent a trivial improvement. However, against the baseline of a middling-value life, there is a non-trivial difference between two hours of misery and only one. Therefore, on this approach, we should conclude that A is better than B . This implausible conclusion can be interpreted as an extreme degree of risk aversion with respect to years of happy life. B has a much greater expected number of years of happy life than A , but it is still judged worse; this is the characteristic pattern of risk aversion.

What about versions of timidity on which value is not bounded? Views that generate timidity through risk aversion are likely to end up in the same boat. To avoid recklessness, the level of risk aversion must be very high; but then it will seem unacceptably high in other cases. Concretely, suppose that risk aversion is encoded by a bounded utility function, in the context of EUT. Then we can run the same example as before, supposing that 10^{10} extra years of happy life would bring you very close to the upper bound for utility. We will again conclude that A is better than B . In the case of

REUT, one could avoid this conclusion by combining the bounded utility function with a risk-function that is extremely risk-seeking. But then the theory will be implausibly risk-seeking in more ordinary cases.

Finally, what about tail discounting? Tail discounting has an advantage over the views we have just considered. Our example of extreme risk aversion seems particularly strange because the probabilities involved—those of a fair coin landing heads or tails—are middling in value, quite different from the exotically small probabilities that seem most relevant to the dilemma between timidity and recklessness. In such middling-probability cases, tail discounting will tend to agree with expected value theory, leading to generally plausible results. But it will still lead to weird results when we consider probabilities close to the threshold.

Specifically, tail discounting implies the following condition that is both stronger and stranger than ordinary timidity, and which isn't implied by the other timid theories we considered.

Threshold timidity: There is some positive probability threshold such that, for any finite, positive payoffs x and y , getting x with probability below the threshold is never better than getting y with probability above the threshold—no matter how much better x is than y and no matter how close together the two probabilities may be.

So, for example, the prospect of improving 10^{80} lives with some particular positive probability, one trillionth of a percent below the threshold, is no better than the prospect of improving only one life with a probability two trillionths of a percent greater. Roughly speaking, threshold timidity tells us that, close to the threshold, decreasing risk is infinitely more important than increasing expected value.

3.4 Extreme risk seeking in very negative outcomes

Just as timidity is associated with risk aversion for prospects with good outcomes, negative timidity is associated with risk *seeking* for prospects with very bad outcomes. Because of this, negatively timid theories face mirror-images of all the problems described in the previous subsection. In fact, these problems tend to look even worse, insofar as risk seeking generally seems more unreasonable than risk aversion does.

To vary the example, here's an interpersonal case. Suppose that a theory is negatively timid in that it puts bounds on the badness of many people suffering. This type of theory is pretty implausible right away. Some (e.g. prioritarrians) would argue that additional suffering within a single life has increasing moral urgency, and some utilitarians would argue that the importance of additional suffering neither diminishes nor increases, but there seems little to recommend the view that additional suffering could become decreasingly important as it piles up. Still, to see how risk seeking enters in, let us suppose that the suffering of 10 trillion people for a year would take us trivially close to the bound. Then this type of theory would tend to favour prospect B over A in the following case:

	Heads	Tails
A	10^{10} people suffer for a year	1 person has an enjoyable day
B	10^{80} people suffer for a year	1 person has two enjoyable days

After all, against the baseline of a middling-valued world, the happiness of one person for one additional day is a non-trivial consideration; it would outweigh the supposedly trivial consideration of all that extra suffering.

If negative timidity arises not from a bounded value function, but from a bounded utility function, then essentially the same problem arises. Although a view of this type gets to insist that additional suffering does not decrease in value, it would still lead to extreme risk seeking with respect to suffering, and this is almost as implausible.

Finally, tail discounting leads to a negative version of threshold timidity. It would say, for example, that the prospect of 10^{80} people suffering for a year, with some particular positive probability, very close to the threshold, is no worse than the probability of one person suffering for one hour, with a probability that is one part in a trillion greater. Roughly speaking, this negative variant of threshold timidity would say that, close to the threshold, *increasing* risk is infinitely more important than increasing expected value. This does not strike us as a tenable view.

3.5 Long-shots

Recklessness suggests that we should value some extreme long-shot bets in a way that often seems counterintuitive. However, all the timid views we considered in section 2 also recommend certain long-shot bets, so the purported advantages of timidity are undermined.

To take the simplest case first, consider tail discounting. Suppose for concreteness that the probability threshold is 0.00001. Then we avoid recklessness, but, if the value of finite payoffs is unbounded, it will still be true that any finite payoff for sure, no matter how good, is worse than a 0.000011 probability of some other finite payoff. But a longshot bet with a 0.000011 probability of a payoff is not really less objectionable than one with a 0.00001 probability. This illustrates the more general problem: there seems to be no way to set the threshold so that tail discounting rules out all and only the objectionable longshots. Perhaps, however, this problem can be mitigated by allowing the probability threshold to be vague, to match the way in which it is vague which longshots are objectionable.¹³ We won't press the issue here.

Now here's an example of the type of longshots that are encouraged when value or utility is bounded above.

Lingering doubt: In an alternate possible world, people live in a utopia. Life is extremely good for everyone, society is extremely just, and so on. Their historians offer a reasonably well-documented history where life was similarly good. However, the historians cannot definitively rule out various unlikely conspiracy theories. Perhaps, for instance, some past generation 'cooked the books' in order to shield future generations from knowing the horrors of a past more like the world we live in today.

Against this background, let us evaluate two options: one would modestly benefit everyone alive *if* (as is all but certain) the past was a good one; the other would similarly benefit only a few people, and only if the conspiracy theories happened to be true.

In this case, the second option might yield a better prospect. Why? If the conspiracy theories are not true, then the status quo is already about as good as it could be, and the benefits to everyone alive would make almost no difference to the value (or utility) of the world. But if the conspiracy theories are true, and the world overall is middling in value or utility, then benefiting relatively few people would make a relatively large difference.

¹³See Peterson (2002, p. 53) for this idea in the context of *de minimis* risk, and Thomas (2020a) for vagueness as a general response to spectrum arguments.

Of course, what's strange about this case is not *only* that it involves seemingly undue attention to a highly unlikely scenario. Contrary to the timid analysis of the case, we feel that modestly benefitting everyone alive would in fact be a weighty consideration, even in the presence of some risk.

3.6 Itemized billing for the timid

To summarize, timid theories have the following features that seem problematic for many types of goods. The first three are general.

1. They pass up at least one seemingly great low-risk deal in Devil at Your Deathbed, or in a similar case for other types of goods.
2. They must claim that, in some cases, even a small decrease in one unlikely payoff cannot be outweighed by any increase in a much more likely payoff (§3.1).
3. Their ranking of prospects is sensitive in implausible ways to how things are in the far removes of space and time (§3.2).

The others features are had by all the timid theories we have identified, and it is hard to see how we could avoid them; but, for all that, they may not be completely general.

4. They rank prospects in an extremely risk-averse way (§3.3).
5. In the case of negative timidity, they rank other prospects in an extremely risk-seeking way (§3.4).
6. Like reckless theories, they still recommend betting on certain kinds of counterintuitive longshots (§3.5).

All together these problems show that justifying timidity in a plausible way is a serious challenge, especially for morally relevant goods. But recklessness will also have its costs.

4 Recklessness and Dominance

There is some hope that the basic implausibility of recklessness might be open to debunking. It is hard to comprehend payoffs that are getting arbitrarily large, so perhaps intuition is not to be trusted in these cases.¹⁴ And our intuitions might also be confused by all kinds of confounding factors that we would expect to see in practical situations where the stakes are very high and one's evaluation of prospects is sensitive to miniscule absolute changes in probabilities: it's natural to start worrying about the reliability of one's evidence and one's cognitive facilities, one's powers of introspection, whether one is being tricked or one is simply hallucinating, and so on. Perhaps in the unusually simple and clearly specified cases which recklessness strictly speaking involves, 'reckless' behavior really would be reasonable.

Still, the case against recklessness can be pressed, and that's what we'll do in this section and the next. Here, we'll argue that recklessness leads to violations of a very compelling dominance principle. Next, we'll argue that (under some further assumptions) recklessness leads to the single-minded pursuit of infinite payoffs.

In the context of expected utility theory, we know that recklessness corresponds to the use of an unbounded utility function. And there is a large literature that focuses on the problems that arise from an unbounded utility function, centered around the *St Petersburg gamble*. It turns out that at least some of these problems arise directly from recklessness, given background assumptions that are much weaker than the assumption of expected utility theory. Given that the alternative is timidity, the problems are deeper and more general than usually presented. That's what we'll now explain.

Recall that the St Petersburg gamble, in its modern guise, offers a $\frac{1}{2}$ chance of 2 units of utility, a $\frac{1}{4}$ chance of 4 units, a $\frac{1}{8}$ chance of 8 units, and in general a $\frac{1}{2^n}$ chance of 2^n units. The expected utility of this prospect is infinite.¹⁵ Some of the problems associated with the St Petersburg gam-

¹⁴See Baron and Greene (1996) for enlightening examples about how pre-theoretic intuition fares very badly with respect to evaluating alternatives involving large numbers of people.

¹⁵The original St Petersburg gamble had monetary payoffs, and so infinite expected *monetary* value. The observation that one can create St Petersburg gambles with infinite expected utility, as long as the utility function is unbounded, goes back at least to Menger (1934). To emphasize: our own argument below in no way assumes expected utility

ble involve handling this infinity in a sensible way. As but one example, the ‘Petrograd’ gamble (Colyvan, 2008) is like St Petersburg except that it invariably pays off one extra unit of utility. It is presumably better than St Petersburg (indeed, one unit better!), but it apparently has ‘the same’ infinite expected utility. The more fundamental problem, however, is that the St Petersburg gamble leads to violations of some extremely attractive normative principles which are often themselves assumed in axiomatic approaches to expected utility theory. Specifically, consider the following two principles:

Outcome-outcome dominance: If, no matter what, the outcome of A would be at least as good as the outcome of B , then A is at least as good as B .¹⁶

Prospect-outcome dominance: If prospect A is strictly better than each possible outcome of prospect B , then A is strictly better than B .

By outcome-outcome dominance, the St Petersburg gamble is at least as good as a version that caps the payoff at 2^n units of utility.¹⁷ But (one can check) this capped version has expected utility greater than n ; it is better than getting n units of utility for sure. So, by transitivity, the St Petersburg gamble is strictly better than getting n units for sure, for every natural number n . Therefore, the St Petersburg gamble is strictly better than every one of *its own* possible outcomes. Since the St Petersburg gamble is not (however) strictly better than itself, prospect-outcome dominance fails.¹⁸

theory.

¹⁶This principle is often called ‘statewise dominance’. The outcome of a prospect depends on what state the world is in initially; the antecedent in the principle is that, whichever state the world is in, A results in an outcome that is at least as good as that of B . Some decision-theoretic frameworks do not formally represent prospects using states, but it should still be possible informally to specify prospects by their effects in different states, and to assert outcome-outcome dominance for prospects so specified. Alternatively, one can replace outcome-outcome dominance in the following discussion by the notionally stronger but also very popular condition known as ‘stochastic dominance’ that does not refer to states.

¹⁷This ‘capped’ version, in other words, offers a $\frac{1}{2}$ chance of 2 units of utility, a $\frac{1}{4}$ chance of 4 units, a $\frac{1}{8}$ chance of 8 units, and so on, up to a $\frac{1}{2^{n-1}}$ chance of 2^{n-1} units, and offers 2^n units with the remaining probability.

¹⁸See Chalmers (2002) for a version of this problem, and Hammond (1998) for the use of similar dominance principles to rule out unbounded utility functions.

The problem as we have described it so far depends on expected utility theory. Unfortunately, it generalises to all reckless views. For, assuming recklessness, we can construct a ‘generalized St Petersburg gamble’ in the following way. Take any payoff x_0 . By recklessness, there is a payoff x_1 such that a $1/2$ chance of x_1 is better than x_0 for sure. By recklessness, there is an payoff x_2 such that a $1/4$ chance of x_2 is better than x_1 for sure. And so on, obtaining a sequence of payoffs x_1, x_2, x_3, \dots , such that a $1/2^n$ chance of x_n is better than x_{n-1} for sure. Now consider a ‘generalized St Petersburg gamble’ A that offers a $1/2^n$ chance of x_n , for every $n = 1, 2, 3, \dots$. By outcome-outcome dominance, A is at least as good as simply getting the $1/2^n$ chance of x_n and nothing otherwise. And by construction this is better than getting x_{n-1} for sure. So, again, A is better than every one of its possible outcomes, violating prospect-outcome dominance.

We think that prospect-outcome dominance and (especially) outcome-outcome dominance are both compelling principles; it’s hard to fathom why they should be violated in the clean kinds of cases we have in mind, where the value of outcomes is apparently the only thing at stake and there are no violations of rights or objectionable unfairness or other complicating factors. Moreover, the timid theories with bounded value or utility functions that we discussed in section 2 are immune to these problems. This is a serious strike against recklessness.¹⁹

5 Recklessness and Infinity Obsession

Now we introduce a problem that arises for reckless theories when there is a possibility of an infinite payoff. While theorizing about infinite cases in ethics is notoriously difficult (see Bostrom (2011) for an entry to the literature), the timid approaches we surveyed in section 2 don’t lead to the particular problem we identify here. And we aren’t aware of any infinitarian problems that are specific to timid views. That speaks in favour of

¹⁹What about views that discount small probabilities (§2.3)? While such views are immune to the the recklessness-based argument we gave in this section, the problem still arises. To see this in the context of tail discounting, consider a prospect in which one will either, with probability below the threshold, face a St Petersburg gamble, or, with the remaining probability, get a zero payoff. This will lead to a violation of prospect-outcome dominance. On the other hand, simply ignoring the tails of the distribution (as in fn. 10) would violate outcome-outcome dominance. We do not know whether there is a reasonable way for probability discounters to thread the needle here—it’s yet another challenge for advocates of these views.

timidity, and against recklessness.

While in philosophy and in life we usually don't think about the possibility of achieving infinite payoffs, some philosophers, such as Pascal, claim that such infinite considerations should be decisive, at least in theory.²⁰ More precisely, these people claim:

Infinity obsession: Any non-zero probability, no matter how small, of an infinite payoff, is better than any finite payoff for sure.²¹

Throughout this discussion, by 'an infinite payoff' we will mean specifically a *positively* infinite payoff, like an infinite number of years of happy life, or an infinite number of lives saved, or something equally good. For example, in Pascal's case, infinity obsession arises from the claim that an arbitrarily small chance of eternity in heaven is better than any finite reward. But of course the negatively reckless will face similar problems with respect to negatively infinite payoffs.

In this section, we argue that, given minimal assumptions, reckless agents *must* be infinity-obsessed (§5.1), and that, given slightly stronger assumptions, recklessness leads to other, potentially even stranger forms of obsession (§5.2). Finally, we argue that it's unclear to what extent these forms of obsession would affect our judgments about cases of immediate practical relevance, but they remain troubling in other ways (§5.3).

First, though, let us make sure the differences between recklessness and infinity obsession are quite clear. Of course, only the latter involves infinite payoffs, but there is another difference too. Though the reckless are willing to take some extreme risks for finite rewards, their willingness to take a risk for any given reward depends on the probability involved. When it comes to infinite rewards, the infinity-obsessed have no such concern. *Any* non-zero probability of an infinite reward seems better to them than any finite reward for sure. To emphasize how strange this is: in any instance of recklessness, there's some conceivable evidence you could show

²⁰See Hájek (2003) for an insightful discussion of Pascal's wager from the point of view of modern decision theory.

²¹We continue with the standard assumption that probabilities are real numbers, and so exclude infinitesimal probabilities. We thus set aside the natural question of what to do about infinitesimal probabilities of infinite payoffs. Infinity obsession involves a violation of the standard 'continuity' or 'Archimedean' axiom used by expected utility theory and many other decision theories.

the reckless agent, short of definitively proving, with probability $\frac{1}{2}$, that his longshot won't pay off, that will make him give up his interest in it. Not so for the infinity-obsessed; unless you definitively prove, with probability $\frac{1}{2}$, that his infinite longshot won't pay off, he'll prefer to take his chances.

5.1 How recklessness leads to infinity obsession

Despite their differences, recklessness leads to infinity obsession, given only very weak assumptions. For example, suppose we have an agent who is reckless about the number of happy years of life he has. Compared to getting any finite number of years of happy life for sure, he'd prefer even a tiny chance of living for a sufficiently long but finite time. *A fortiori*, he'd prefer the same tiny chance of living for infinitely long. So he'd prefer any tiny chance of living forever to any finite number of years of happy life for sure. In other words, he'll be infinity-obsessed.

More formally, the premises we need to derive infinity obsession are recklessness, transitivity, outcome-outcome dominance, and the claim that an infinite payoff is better than a finite one.²² The argument goes like this. Consider a prospect A that gives the agent an infinite payoff with probability $p > 0$, let B be a prospect that gives the agent some finite payoff y with that same probability p , and let C be a prospect that gives the agent some other finite payoff x for sure. We have to show that A is better than C . According to recklessness, there's some value of y that would make B better than C . But an infinite payoff is better than a finite payoff, so, by outcome-outcome dominance, A is at least as good as B . Therefore, by transitivity, A is better than C .

5.2 Other forms of obsession

We've formulated infinity obsession in a way that ties it very closely to recklessness, as the preceding argument shows. But slightly stronger assumptions lead from recklessness to other forms of obsession, which may be even more troubling. To illustrate, suppose we accept

²²It's true that, in section 4, we argued that recklessness requires one to deny either outcome-outcome dominance or prospect-outcome dominance. We think it would be quite remarkable to reject the particular application of outcome-outcome dominance that follows. Be that as it may, the logical point is that if one rejects *only* the slightly less compelling condition of prospect-outcome dominance, then one still falls prey to infinity obsession.

The sure thing principle: If, on the supposition that some proposition E is true, prospect A is at least as good as B , and, on the supposition that E is false, A is better than B , then A is better than B .

Although the sure thing principle in its full generality is closely associated with expected utility theory, the particular applications we will make of it are relatively uncontroversial; for example, they are compatible with the risk-weighted expected utility theory we described in section 2.²³ Here, then, are two further types of obsession that will ensnare the reckless if they accept the sure thing principle.

First, infinity obsession compares a prospect with *some* chance of resulting in an infinite payoff to a prospect with *no* such chance: it says the former is always better. The following condition suggests, more generally, that a *higher* chance of an infinite payoff is always better.

Generalized infinity obsession: Getting an infinite payoff with some probability p and nothing otherwise is better than getting the same infinite payoff with any smaller probability q and a finite payoff x otherwise—no matter how good x may be.

A bit roughly, someone who is infinity-obsessed in this generalized sense is obsessed with increasing the probability of infinite payoffs, completely heedless of finite considerations.²⁴

Although generalized infinity obsession is stronger than plain old infinity obsession, one can argue for the stronger claim from the weaker. Suppose that A is the prospect of getting the infinite payoff with probability p and nothing otherwise, and B is the prospect of getting the infinite payoff with probability q and x otherwise. First, it seems harmless to assume that, p being higher than q , A results in the infinite payoff whenever B does. Now suppose that A and B both result in the infinite payoff. Then they turn out equally well; on that supposition, A and B are equally good.

²³That is, we will at most use a simple case of the *comonotonic* sure thing principle discussed by Buchak (2013). We won't go into the details, since one can judge our application of the principle on its own terms.

²⁴Hájek (2003) entertains the view that, contrary to generalized infinity obsession, all prospects involving infinite payoffs are equally good. But this can't be right if one accepts outcome-outcome dominance in the relevant cases and thinks an infinite payoff is better than a finite one.

On the contrary supposition, B results in the finite payoff x , and A either results in the infinite payoff or in nothing. By infinity obsession, A is better than B , on this supposition. Therefore, by the sure thing principle, A is better than B .

Second, here is a troubling kind of obsession that involves only finite payoffs. In section 4, we showed how to construct a ‘generalized St Petersburg gamble’ using recklessness and only finite payoffs. An argument very similar to the one for infinity obsession, but using the sure thing principle in place of outcome-outcome dominance, shows that reckless theories are likely to be ‘St Petersburg-obsessed’:

St Petersburg obsession: Any non-zero probability, no matter how small, of facing a generalized St Petersburg gamble, is better than any finite payoff for sure.

We’ll omit the argument for the sake of space.²⁵

Of course, one can try to argue in similar ways for forms of obsession that apply in a wider range of cases, using similar or stronger dominance principles. The intention of our current arguments is simply to show that there are a range of troubling obsessions that a reckless theory will have difficulty avoiding in any plausible and principled way.

5.3 Does infinity obsession have practical consequences?

Even if infinity obsession and its kin are troubling as theoretical consequences of recklessness, their practical implications remain unclear. For one thing, our formulation of infinity obsession only involves a class of very simple prospects (each involving at most two outcomes), and realistic cases are bound to be much more complicated. We haven’t, moreover, considered negatively infinite payoffs, and the related question of how to weigh chances of positive infinities against chances of negative infinities; how infinite payoffs compare with generalized St Petersburg gambles; and so on. But all that this suggests is that even harder thinking about infinities might be needed to handle realistic cases.

On the other hand, it may be tempting to think that questions of infinity are practically irrelevant, in the sense that they do not actually arise

²⁵In the context of expected utility theory, the point that an unbounded utility function leads to violations of the continuity axiom through St Petersburg gambles goes back at least to Arrow (1965).

when it comes to evaluating the prospects that we are likely to face any time soon. That would arguably reduce the sting of infinity obsession. In the rest of this section, we'll consider two attempts to justify such a claim. We think both warrant further investigation, but even if they happen to work in our current circumstances, there are still foreseeable cases in which infinity obsession leads to bizarre recommendations.

CANCELLATION OF INFINITIES AND 'SIMPLE CLUELESSNESS'. At a first pass, one could avoid worrying about infinities by simply giving zero probability to any scenario in which you achieve an infinite payoff. However, we don't think this is tenable: it seems irrational to assign zero (or even infinitesimal) probability to all the infinite scenarios that have been contemplated by physicists, futurists, and theists. As Bostrom (2011) points out, many astrophysicists now believe that the universe is infinite, and this would make it likely that there are infinitely many valuable lives. In turn, it is hard to be absolutely certain that one will not happen to achieve an infinite payoff in terms of lives improved, or perhaps in terms of good lives created.²⁶

However, there is a variation on this strategy that seems more promising, and which anyway clarifies our predicament. When it comes to comparing the prospects of fairly mundane actions like buying cake versus buying ice cream, it seems intuitive that any infinitarian considerations should cancel out. While, as above, it is hard to entirely rule out the possibility that buying cake will indirectly lead to an infinite payoff, it might be reasonable to give *the same* probability to an infinite payoff from cake as from ice cream. Then the infinitarian considerations might not favour either option on balance. Indeed, the forms of infinity obsession we have discussed concern cases in which the probability of an infinite payoff varies. So it is not obvious that they have any implications for how we should think

²⁶A familiar point is that certainty is an implausibly high standard in that it is immune to new evidence, at least by standard Bayesian lights. Granted, there are broadly Bayesian views on which certainty is relatively easy to come by and easy to lose. According to a view defended by Williamson (2000, ch. 10), Bayesians should conditionalize on everything they know. Insofar as we know all kinds of everyday propositions, like what we'll have for dinner, we can be certain of many things. But even on this view, the point remains: we just *aren't* in a position to know that an infinite payoff is completely off the table. We'd have to know a lot more about physics (let alone theology!) for that to be true, and what we'd have to know would be in tension with what we currently suspect. One might also worry whether high stakes in themselves can undermine epistemic justification.

about mundane cases.

How well this story works will depend on the details of our theory of evaluating prospects, but it also depends on a strong type of indifference principle. Why should the probabilities be *exactly* the same in the cases of cake and ice cream? If one of the probabilities is ever so slightly larger, then generalized infinity obsession suggests that *that* fact should be decisive.²⁷ We agree that this is a serious challenge, and only note that an indifference principle of the relevant form has been defended by Greaves (2016) in cases of what she calls ‘simple cluelessness’. The present example would involve simple cluelessness insofar as there is no reason to think that one of the options leads to an infinite payoff that doesn’t apply, *mutatis mutandis*, to the other.

EMPIRICAL STABILIZING ASSUMPTIONS. Another way to try to limit the practical consequences of infinity obsession (and its generalisations) is to invoke what Bostrom (2011) calls an *empirical stabilizing assumption*—some empirical assumption which, if true, means that what’s best conditional on achieving a finite payoff and what’s best overall are broadly equivalent. One can imagine other domains where distinct considerations generally lead to the same recommendations. For example, it may be that, for many children, doing what their mother says and doing what’s in their own best interest generally amount to the same thing, because (and here’s the analogue of an empirical stabilizing assumption) mothers are predictably motivated to promote their children’s interests.

What kind of empirical stabilizing assumption might hold? Someone might argue that whatever best promotes the long-term survival and flourishing of our descendants in the distant future will generally be what’s best in terms of finite considerations *and* what’s best in terms of infinite considerations. It will do well with respect to finite payoffs because of the enormous (even if finite) number of potentially good lives in the future.²⁸ And it will do well with respect to infinite considerations because,

²⁷A different thought is that the probability of achieving an infinite payoff will necessarily be ‘imprecise’ for either option. Then the theory of decision-making with imprecise probabilities will come to the fore. Here a serious worry is that imprecision about the probability of an infinite payoff will tend to mean that almost any option is permissible—see Mogensen (2020) for a version of this worry for high-stakes finitary cases.

²⁸Parfit (2011, chapter 36), Bostrom (2003, 2013) and Beckstead (2019) all argue for this conclusion.

if humans manage to survive and flourish into the distant future, they'll be reasonably likely to find and execute any available strategies for achieving infinite payoffs (and perhaps also for avoiding negatively infinite payoffs). It is plausible, after all, that future generations will have *some* interest in producing infinite payoffs (especially if recklessness is true), and that flourishing and long-lasting future civilisations will be better placed than we are to do so, if it's possible at all. Conversely, it's unclear how we could realistically promote infinite payoffs without enabling our descendants in more general ways.²⁹ While there is much more to think about here, it does seem that, even if recklessness is true, there might be significant overlap between what's best with respect to finitary considerations and what's best overall, in our current situation.

TWO CAVEATS. The above considerations suggest that infinity obsession and its cousins might not be extremely revisionary for current practical purposes. Perhaps this makes recklessness somewhat less unattractive. We end, however, with two serious caveats.

First, even if some empirical stabilizing assumption meant that taking infinite payoffs into account would not greatly change our evaluation of the prospects open to us, it would significantly change the reasons behind these evaluations, and in strange ways. For someone who suffers from generalized infinity obsession, the main reason that it would be good to prevent climate change would be that it would (for example) increase the chance that future people find some way of achieving an infinite amount of value, rather than the fact that lots of future generations would be better off in some more likely, merely finite way. Likewise, prudentially, the main reason that it would be important to avoid smoking would be that it might increase your odds of living long enough for some method of giving people infinite value to be discovered, and 'lung cancer is an awful experience' would be a relatively minor consideration. Infinity obsession would remain revisionary in this way.

Second, there are situations that might arise in the distant future where the relevant probabilities wouldn't cancel out and no empirical stabilizing assumption would apply. In some of those cases, infinity obsession still seems deeply troubling. Here is a stylized illustration.

Infinite research vs. utopia: Our descendants reach the lim-

²⁹For a similar line of thought, more fully developed, see Williams (2013).

its of technological progress and become very confident, but not certain, that it's impossible to achieve an infinite payoff. (And, indeed, it is impossible). They must decide how some vast amount of resources should be allocated between two projects: creating an extremely good (but only finite) utopia, or researching possible methods of achieving an infinite payoff.

It is likely that if these people were infinity-obsessed, they would spend nearly all of the resources on the infinite research; they would keep becoming more and more certain that their research would bear no fruit; and they would keep going at it as long as they didn't become *completely* certain that achieving an infinitely good outcome was impossible—which perhaps they never would.

5.4 Wrapping up

In this section we've argued for the following claims. A reckless theory of value must be infinity-obsessed if (i) it says that infinite payoffs are better than finite ones and (ii) it satisfies transitivity and outcome-outcome dominance. It will be infinity-obsessed in the generalized sense, as well as St Petersburg-obsessed, if it allows certain applications of the sure thing principle. These types of obsession are implausible, and are a major cost of recklessness. It's less clear what the practical consequences of infinity obsession and its kin would be; it's at least arguable that infinitarian considerations cancel out in mundane cases, and that empirical stabilizing assumptions are relevant in others. But there would probably be some major differences in some cases we could encounter, and infinity obsession would remain highly revisionary at the level of justification.

6 Recklessness and Moral Uncertainty

Some philosophers have argued that we should treat uncertainty about moral or evaluative matters in essentially the same way that we should treat empirical uncertainty, and specifically that we should use expected utility theory in both cases.³⁰ After all, the axiomatic basis of expected

³⁰For early examples, see Lockhart (2000), Ross (2006), Sepielli (2010); see MacAskill et al. (2020) for a recent book-length treatment, and Riedener (2020) for a version of the axiomatic approach mentioned in the next sentence.

utility theory is *prima facie* plausible regardless of what type of uncertainty is at stake. In this section, we raise an objection to this type of view, based on the possibility of reckless and infinity-obsessed theories of value. We then comment on how this fits into the general challenge of constructing an adequate theory of evaluative uncertainty.

Let us step back for a moment to make sure the topic is clear. In this context, a natural way to understand the question is: how should your preferences reflect your credences about value? Here is an example.

Suppose that you have some credence in utilitarianism and some credence in a pluralistic form of egalitarianism, as theories of value. Should you then prefer a more equal outcome to a less equal one when egalitarianism says that it is better but utilitarianism says that they are equally good?

For our purposes, one can understand the ‘should’ as a matter of coherence between your beliefs and your desires. The question is whether you could coherently have the credences described while also (for example) being entirely indifferent about equality. A theory of evaluative uncertainty would answer this question and others like it.

Here is the problem. Suppose that each of the theories T_i in which you have some credence ranks simple prospects by expected utility—each one with respect to a different utility function u_i . Suppose your preferences also rank simple prospects by expected utility, with respect to yet another utility function u . Finally, suppose that your preferences over simple prospects satisfy the following

Pareto assumption: If every theory T_i in which you have some credence says that A is at least as good as B , then you weakly prefer A to B (that is, you either strictly prefer A to B , or you are indifferent between them).

If, in addition, one of the T_i says that A is better than B , then you strictly prefer A to B .

(So, in the preceding example, you *do* prefer the more equal outcome to the less equal one.) Then, one can show that, if any of the T_i are reckless, your preferences must be reckless as well.³¹ Indeed, it follows from Harsanyi’s

³¹We’re continuing to assume here that all the theories under consideration recognize

aggregation theorem that u is a weighted sum of the u_i :³²

$$u = \alpha_1 u_1 + \alpha_2 u_2 + \dots + \alpha_n u_n \quad (1)$$

for some positive real numbers α_i . If any one of the T_i is reckless, or equivalently if any one of the u_i is unbounded, then u must be unbounded as well.

Despite the drawbacks of recklessness, it is hard to deny that we should have *some* credence in reckless theories of value, given the challenges faced by the alternatives, and given that on many standard theories, the value of finite payoffs is unbounded. It is therefore hard to deny that, if empirical and evaluative uncertainty are both governed by expected utility theory, then one's preferences should be reckless. Moreover, since advocates of expected utility theory typically endorse outcome–outcome dominance and the sure thing principle as core commitments, we seem bound to fall into generalized infinity obsession and St Petersburg obsession as well.

This could reasonably be seen as an objection to the use of expected utility theory in this context. It is true, after all, that some quite different approaches to evaluative uncertainty will avoid the problem altogether. Gustafsson and Torpman (2014) defend 'My Favourite Theory', which in this context would be the view that your preferences should match the judgments of the evaluative theory in which you have highest credence; they reject both expected utility theory and the Pareto condition. And normative externalists, like Weatherson (2019), might claim that there simply aren't interesting norms in this area (let alone ones that would require reckless preferences). So our argument could be taken to lend support to such alternative approaches, or simply as raising an important problem for those attracted to expected utility theory and similar views.

that payoffs of the relevant kind are good to some extent (or at least limited in how bad they are). This avoids some strange and exceptional cases where different theories perfectly cancel one another out. For example, if total utilitarianism and the opposite of total utilitarianism were the only theories in which you had credence, and you gave them exactly the same weight, then your preferences would not be reckless; you would be indifferent about everything. But even without this assumption, it would take a remarkable coincidence for different theories to cancel out exactly in a way that avoids recklessness.

³²See Harsanyi (1955, 1977, §4.8). In the traditional interpretation of Harsanyi's framework, each u_i is a utility function representing the preferences of some person, and u is a utility function representing the preferences of a social planner. The result is that the social planner's utility function should be a weighted sum of those of individuals. The idea of using Harsanyi's theorem in the context of moral uncertainty is found in Beckstead (2013) and independently in Riedener (2015, 2020).

This argument is closely related to, but importantly different from, a common observation in the literature. Suppose there were a universal notion of ‘moral value’ such that each theory T_i assigned a degree $v_i(x)$ of moral value to each outcome x . We could then make intertheoretic comparisons of moral value, i.e. we could ask whether T_i gives greater moral value to x than T_j does.³³ And suppose you formed your preferences in line with expected moral value, i.e.

$$u = \text{Cr}(T_1)v_1 + \text{Cr}(T_2)v_2 + \dots + \text{Cr}(T_n)v_n \quad (2)$$

where $\text{Cr}(T_i)$ is your credence in T_i ; the similarity between (1) and (2) should be clear. Then the common observation is that your preferences will tend to be determined by high-stakes theories, i.e. ones that claim there are very large differences in moral value between relevant outcomes; moreover, this can be true even if you have only tiny credence in those high-stakes theories. The low-stakes theories are ‘swamped’ by the high-stakes ones. Again, this may seem objectionable.³⁴

Because our argument focusses on the specific phenomenon of recklessness, it is able to rely on a much more minimal view about evaluative uncertainty. First, we don’t need to say anything about any cardinal notion of ‘moral value’ or ‘high stakes’; our use of utility functions is compatible with the view that utility is just a formal device based on the structural axioms of expected utility theory, such as the sure thing principle. Second, relatedly, we don’t need to say anything about the meaningfulness of any sort of intertheoretic comparisons, a deep problem for moral uncertainty that goes back at least to Hudson (1989).³⁵ And, finally, far from *assuming* that the overall utility function u is a weighted sum of the u_i , in analogy

³³Strictly speaking, in this discussion we only need to make intertheoretic comparisons between *differences* in moral value, but we set this aside to simplify the exposition.

³⁴See e.g. Ross (2006), Greaves and Ord (2017), MacAskill et al. (2020, ch. 6) for discussions of this problem. We note that the recent literature on normative uncertainty is highly sophisticated, and the comments in the next paragraph are not intended as objections.

³⁵It would be reasonable, though, to suggest that the agent’s preferences must end up reflecting an implicit view about such comparisons. Where would this implicit view come from, if intertheoretic comparisons are meaningless? Perhaps it could be purely subjective and pragmatic—that is, not an implicit view *about* intertheoretic comparisons, but a practical stance towards reconciling competing demands. At any rate, the logical point is that, insofar as you do, on whatever basis, satisfy the stated conditions, your preferences will be reckless. One could also try to generalize the argument to allow for *incomparability* in the agent’s preferences, contrary to the most orthodox version of expected utility

to (2), we derive it, via Harsanyi's theorem, from a small number of more basic principles.

7 Conclusion

In summary, as far as the evaluation of prospects goes, we must be willing to pass up finite but arbitrarily great gains to prevent a small increase in risk (timidity), be willing to risk arbitrarily great gains at arbitrarily long odds for the sake of enormous potential (recklessness), or be willing to rank prospects in a non-transitive way. All options seem deeply unpalatable, so we are left with a paradox. Of course, as we argued in section 6, the paradox may effectively resolve itself: on some views about normative uncertainty, we are bound to have reckless preferences, even if our credence in reckless theories is vanishingly small. But this result is unpalatable in itself.

If we accept timidity, we must think that sometimes even a small reduction in an unlikely payoff cannot be outweighed by a sufficiently large reduction in a more likely one. We must care about seemingly irrelevant details involving the distant reaches of space and time. We are likely to endorse extreme risk aversion in some cases and, when it comes to negative timidity, extreme risk seeking in others. And, for all that, it is not clear we can avoid favouring problematic long-shot bets.

If we accept recklessness, then we must deny some compelling dominance principles, like prospect-outcome dominance. And we are likely to be infinity-obsessed, pursuing any chance of an infinite payoff, or even any chance of a generalized St Petersburg gamble, at any finite expense.

Some may see this as another argument for ranking prospects non-transitively. We are not inclined toward this resolution, but we think that increasing one's confidence in this position is a fair reaction to the arguments, given that new challenges have been presented for other approaches, but not for this one.

theory. We note that some generalizations of Harsanyi's theorem allow for incomparability (see McCarthy et al., 2019), but if incomparability is sufficiently widespread (as one might guess if intertheoretic comparisons make *no sense at all*) then the argument for recklessness will fail; see MacAskill (2013) for a discussion of similar issues.

References

- Arrow, K. (1965). *Aspects of the Theory of Risk-Bearing*. Yrjö Jahnssoinin Säätiö, Helsinki.
- Baron, J. and Greene, J. (1996). Determinants of insensitivity to quantity in valuation of public goods: Contribution, warm glow, budget constraints, availability, and prominence. *Journal of Experimental Psychology: Applied*, 2(2):107–125.
- Beckstead, N. (2013). *On the Overwhelming Importance of Shaping the Far Future*. PhD thesis, Rutgers.
- Beckstead, N. (2019). A brief argument for the overwhelming importance of shaping the far future. In Greaves, H. and Pummer, T., editors, *Effective Altruism: Philosophical Issues*, pages 80–99. Oxford University Press.
- Bostrom, N. (2003). Astronomical waste: The opportunity cost of delayed technological development. *Utilitas*, 15(3):308–314.
- Bostrom, N. (2011). Infinite ethics. *Analysis and Metaphysics*, 10:9–59.
- Bostrom, N. (2013). Existential risk prevention as global priority. *Global Policy*, 4(1).
- Broome, J. (2004). *Weighing Lives*. Oxford University Press.
- Buchak, L. (2013). *Risk and rationality*. Oxford University Press.
- Chalmers, D. J. (2002). The St. Petersburg two-envelope paradox. *Analysis*, 62(2):155–157.
- Colyvan, M. (2008). Relative expectation theory. *Journal of Philosophy*, 105(1):37–44.
- Greaves, H. (2016). Cluelessness. *Proceedings of the Aristotelian Society*, 116(3):311–339.
- Greaves, H. and Ord, T. (2017). Moral uncertainty about population axiology. *Journal of Ethics and Social Philosophy*, 12(2):135–67.

- Gustafsson, J. E. and Torpman, O. (2014). In defence of my favourite theory. *Pacific Philosophical Quarterly*, 95(2):159–174.
- Hammond, P. J. (1998). Objective expected utility: A consequentialist perspective. In Baberà, S., Hammond, P. J., and Seidl, C., editors, *Handbook of utility theory, Volume 1*, pages 143–211. Kluwer.
- Harsanyi, J. C. (1955). Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, 63(4):309–321.
- Harsanyi, J. C. (1977). *Rational behavior and bargaining equilibrium in games and social situations*. Cambridge University Press.
- Hudson, J. L. (1989). Subjectivization in ethics. *American Philosophical Quarterly*, 26(3):221–229.
- Hájek, A. (2003). Waging war on Pascal’s wager. *The Philosophical Review*, 112(1):27–56.
- Lockhart, T. (2000). *Moral Uncertainty and Its Consequences*. Oxford University Press, New York.
- MacAskill, W. (2013). The infectiousness of nihilism. *Ethics*, 123(3):508–520.
- MacAskill, W., Bykvist, K., and Ord, T. (2020). *Moral Uncertainty*. Oxford University Press.
- McCarthy, D., Mikkola, K., and Thomas, T. (2019). Aggregation for potentially infinite populations without continuity or completeness. Manuscript.
- McCarthy, D., Mikkola, K., and Thomas, T. (2020). Utilitarianism with and without expected utility. *Journal of Mathematical Economics*, 87:77–113.
- Menger, K. (1934). Das Unisicherheitsmoment in der Wertlehre. *Z. Nationalökonomie*, 51:459–485. Translated as ‘The role of uncertainty in economics’, in Shubik, M. (ed.), *Essays in mathematical economics in honor of Oskar Morgenstern*, Princeton University Press, 1967.

- Mogensen, A. (2020). Maximal cluelessness. Manuscript.
- Monton, B. (2019). How to avoid maximizing expected utility. *Philosophers' Imprint*, 19.
- Parfit, D. (2011). *On What Matters: Volume Two*. Oxford University Press, Oxford.
- Peterson, M. (2002). What is a *de Minimis* risk? *Risk Management*, 2(4):47–55.
- Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior & Organization*, pages 323–343.
- Riedener, S. (2015). *Maximising Expected Value Under Axiological Uncertainty: An Axiomatic Approach*. PhD thesis, Oxford University.
- Riedener, S. (2020). An axiomatic approach to axiological uncertainty. *Philosophical Studies*, 177:483–504.
- Ross, J. (2006). Rejecting Ethical Deflationism. *Ethics*, 116(4):742–768.
- Sepielli, A. (2010). *Along an Imperfectly-Lighted Path: Practical Rationality and Normative Uncertainty*. PhD Thesis, Rutgers University.
- Smith, N. J. J. (2014). Is evaluative compositionality a requirement of rationality? *Mind*, 123(490):457–502.
- Stefánsson, H. O. and Bradley, R. (2015). How valuable are chances? *Philosophy of Science*, 82(4):602–625.
- Temkin, L. S. (2012). *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford University Press, Oxford.
- Thomas, T. (2020a). Are spectrum arguments defused by vagueness? Manuscript.
- Thomas, T. (2020b). Separability in population ethics. In *The Oxford Handbook of Population Ethics*. Oxford University Press. Forthcoming.
- Weatherson, B. (2019). *Normative Externalism*. Oxford University Press.

Williams, B. (1973). The Makropulos Case: Reflections on the Tedium of Immortality. In *Problems of the Self: Philosophical Papers 1956-1972*, pages 82-100. Cambridge University Press, Cambridge.

Williams, E. G. (2013). Promoting Value As Such. *Philosophy and Phenomenological Research*, 87(2):392-416.