In Defense of Timidity

There are many moral issues but limited resources to address them, we must carefully choose which one to prioritize. Effective Altruism (EA), a social and philosophical movement created at the turn of the 2010s, focuses on this prioritization challenge. EA attempts to identify the most effective ways to use our resources to maximize welfare and is committed to supporting the identified most effective initiatives (MacAskill, 2019, 14). Historically, EA has prioritized helping the global poor (e.g., MacAskill, 2019; MacAskill, 2015; Singer, 2015).

Recently, EA extended its scope to include and prioritize far-future people (e.g., Ord, 2020; MacAskill, 2022; Greaves et al., 2022; Greaves and MacAskill, 2021; Tarsney, 2023; based on the early work of Beckstead, 2013; Bostrom, 2013; Bostrom, 2002). This extension is known as 'longtermism': 'The view that we should be particularly concerned with ensuring that the long-term future goes well' (Greaves et al., 2021, 20). The central argument for longtermism is what I call the 'longtermist comparative argument', or LCA for short.[1]

The LCA compares the cost-effectiveness of interventions having most of their consequences in the long term (in 100 years onward) with the most cost-effective interventions having most of their impact in the short term (in the next 100 years). Greaves and MacAskill show that long-term interventions are, by far, the most cost-effective interventions. They then argue that because long-term interventions are much more cost-effective than short-term ones, consequentialist and plausible non-consequentialist views, i.e., moderately stake-sensitive views, morally ought to prioritize long-term interventions (Greaves & MacAskill, 2021).

Short-term and long-term interventions are different in many regards. One of these differences relates to an issue in decision theory of uncertain prospects. The issue arises when dealing with prospects having tiny probabilities of huge payoffs. For example, how should we compare a prospect having a probability of one-trillionth to save a quintillion people with a prospect saving 1000 people for sure (Greaves & MacAskill, 2021, 25)? According to standard expected-value

[1] This argument supports strong longtermism the view that long-term consequences are the most important aspect of our decisions today. However, the argument can merely support longtermism if we refuse some views necessary for strong longtermism. For instance, if one refuses a zero rate of pure time preference but rather defends a tiny positive rate, the argument supports longtermism rather than strong longtermism.

theory, the risky prospect is a thousand times better. However, its probability is so low that choosing it will likely result in 1000 people dying for nothing. The LCA is based on comparisons of similar prospects. The prospect with a tiny probability of a large payoff is similar to the long-term intervention, while the sure prospect is similar to the short-term intervention.

There are two main views on this issue: recklessness and timidity. Each view favours a different type of prospect. On the one hand, recklessness (or fanaticism, both refer to the same view) holds that no matter how small the probability of an outcome is, if the payoff is large enough, the outcome is better than any outcome with a great payoff for sure. The reckless view favours the long-term intervention in the LCA. On the other hand, the most plausible version of timidity holds that when an outcome has a probability below a certain threshold, this outcome cannot be better than an outcome with a probability above the threshold, regardless of the size of its payoff. This timid view tends to favour the short-term intervention in the LCA. Thus, the compellingness of the longtermist comparative argument amounts to identifying which view is most plausible: recklessness or timidity. I used the phrase ‘tends to favour’ rather than ‘favours' because, as Greaves and MacAskill (2021, 25) state, timidity might undermine their argument for individual agents but not for societies. I agree. Affluent societies’ contributions to long-term interventions can undoubtedly be above a reasonable probability threshold because of the size of their contribution.

For simplicity’s sake, I assume that long-term interventions have a probability below the probability threshold for individual interventions. To give an order of magnitude, the numbers used in the LCA entail that an individual donation of $4,000 has a probability of 1 in 125 trillion to avoid extinction from a pandemic in the next 100 years (Kosonen, 2022, 261). The dependency of longtermism on recklessness is also unequivocally acknowledged by Greaves and MacAskill:

> We regard this [non-fanatical view] as one of the most plausible ways in which the argument for strong longtermism might fail. Our view is that at present, the question cannot be confidently settled, since research into the possibility of a non-fanatical decision theory is currently embryonic. However, initial results suggest that avoiding fanaticism might come at too high a price (Greaves & MacAskill, 2021, 25).

Greaves and MacAskill defend longtermism because timidity might face worse issues than recklessness. I will argue for an opposite conclusion: a specific timid view faces fewer issues than recklessness. More precisely, I argue that a *lexical discounting* view with a *range threshold* is more plausible than recklessness and, thus, that it is morally permissible for individuals not to prioritize long-term interventions. In other words, the moral duty entailed by the LCA can only apply to societies.

In part 1, I set the table of the issue and specify the timid view I defend, a discounting view with a range threshold. In part 2, I present the five most important objections to recklessness, objections that reckless people have, at least not yet, responded to. In part 3, I present the three most important objections to discounting views and show that the view I defend responds to two of them. I also present a mathematical solution to third objection.

In part 4, I summarize, conclude, and present the implications of my argument for longtermism.

1. Setting the Table

The issue I address in this paper requires some technical clarifications. In section 1.1, I present some background clarifications. In section 1.2, I present the paradox of recklessness to set the issue at hand. In section 1.3, I define the two competing views addressing the paradox: recklessness and timidity. In section 1.4, I present a novel discounting view using a range threshold instead of the standard fixed threshold. I ground this novelty in a contextualist view addressing the vagueness of the concept of 'tiny probability'.

1.1. Clarifications

Let me provide five background clarifications necessary for our discussion. First, here are some definitions essential for my paper. By 'prospect', I mean a set of outcomes with some probability p ($0 < p < 1$) of getting a gain or a loss. The gains and losses addressed in this paper are restricted to payoffs, i.e., quantifiable gains or losses that are potentially arbitrarily large. These payoffs are also restricted to the kind of good longtermists are interested in: additional happy lives.

Second, as most of the arguments I present involve an ordering of prospects, let me define the operators I use. Let us assume that X > Y means that X is better than Y, and that X ⋡ Y means that X is not worse than Y.

Third, the theory of the value of uncertain prospects assumed in this paper is the standard expected value theory (EVT). According to EVT, an outcome is equal to the product of the value of its expected gain or loss by the probability that it will occur. The value of the prospect is the weighted sum of the outcomes. Moreover, according to EVT, a prospect is better than another if and only if it has a higher expected value. This is the most common normative model used to evaluate simple uncertain prospects and the model used by longtermists.

Third, recklessness and timidity can be positive or negative. The positive version addresses good prospects (e.g., people having a good life), and the negative version addresses bad prospects (e.g., lives of suffering). For the sake of brevity and simplicity, I only address the positive version, but the discussion can also apply to negative versions of the views with some adjustments.

Finally, the issue I address in this paper has recently received increased attention because of its importance for longtermism. Three influential publications were recently published on the topic by longtermists or longtermist-adjacents researchers (Beckstead & Thomas, 2023; Wilkinson, 2022; Kosonen, 2022). Wilkinson fully embraces recklessness and strongly rejects timidity. Beckstead and Thomas reject both recklessness and timidity. Kosonen also rejects both views, but she seems, *inter alia*, to consider discounting views to be the most plausible view. The present paper explicitly rejects recklessness and defends a specific timid view discounting tiny probabilities. Let us now clarify the problem we address in the paper.

1.2. The Paradox of Recklessness

The best presentation of the problem facing prospects with a tiny probability takes the form of a spectrum argument presented by Beckstead and Thomas (2023, 1-2). A person meets God once dead. God offers her a ticket for a year of additional life. On her way to exchange the ticket for an additional year, the devil stops her and offers to replace the ticket with another one with a slight decrease in probability (a probability of 0.999) but a significant increase in additional years (10 years). This new deal is much more appealing to her. She exchanges God's ticket for the devil's ticket. The devil then repeats fifty thousand times the offering of a subsequent deal with a slightly lower probability but a much larger payoff than the previous deal. More precisely, for each deal *n*, the probability of getting the payoff is 0.999^n. Thus, the probability of God's deal

(deal 0) is 0.999^0, i.e., 1; the probability of the devil's deal 1 is 0.999^1, i.e., 0.999; the probability of the second devil's deal is 0.999^2, i.e., 0.998, and so forth. The payoff is 10^n, nothing otherwise. Deal 0 has a payoff of 10^0; Deal 1 has a payoff of 10^1, and so forth. Here is the table representing the deals:

Table 1: The Devil's Deals

Deal	0	1	2	3	…	n	…	50,000
Payoff	1	10	10^2	10^3	…	10^n	…	$10^{50,000}$
Probability	1	0.999	0.998	0.997	…	0.999^n	…	$<10^{-21}$

The expected value of the deals always clearly goes up: 1, 9.99; 99.8; 997; and so forth. Subsequent deals always have a much higher expected value than previous ones. Because of this increase in expected value, the person chooses the last deal. However, the last deal probability is 10^{-21}, a probability only slightly higher than winning the first prize of the 6/49 lottery three times in a row.[2] As you might expect, she died the same night with no additional year of life. The paradox is that each subsequent deal seems better than the previous one, but the last deals seem worse than the first ones (Beckstead and Thomas, 2023, 2).

There are three possible ways of dealing with this paradox. Roughly, one can hold that subsequent prospects are always better than previous ones: this is recklessness. One can be critical of the ordering and consider that, at a certain point, the prospects' value is worse: this is timidity. One can also consider that each prospect is better than the previous one but that when prospects are far apart, especially on opposite sides of the spectrum, previous ones might be better than prospects on the other far end: this is the rejection of transitivity (e.g., Temkin, 2012; Temkin, 2015). I leave the rejection of transitivity aside in this paper. Let us now define recklessness and timidity.

[2] The 6/49 lottery consists in choosing six numbers between 1 and 49. The probability of winning the first prize is just above 1 in 14 million.

1.3. The Definitions of Recklessness and Timidity

Beckstead and Thomas (2023, 4) define recklessness as follows:

> ***Recklessness***: 1) For any finite payoff *x*, no matter how good, and any probability *p*, no matter how tiny, 2) there's a finite payoff *y*, such that getting *y* with probability *p* is better than 3) getting *x* for sure. (My numeration)

To make the definition more concrete, let me give an example with arbitrary numbers. For the first portion of the definition, let's consider a prospect X that is expected to give a large payoff *x* of 10^{40} additional lives with a tiny probability *p* of 10^{-20}. For the third portion of the definition, let's consider a prospect X_1 who has the same payoff of $x=10^{40}$, but a probability of 1. X_1 is a fantastic prospect. However, for the second portion of the definition, recklessness entails that a prospect *Y* with the same probability *p* of 10^{-20}, if it has a large enough payoff *y,* makes prospect Y better than prospect X_1. Put simply, no matter how tiny the probability of a prospect is, if the payoff is large enough, the prospect can be better than a sure prospect with an arbitrarily large payoff.

Beckstead and Thomas (2023, 4) then define timidity as follows:

> ***Timidity***: By any standard of closeness, there's a finite payoff *x*, and close-together, positive probability *p* and *q*, such that getting *x* with the slightly higher probability *p* is no worse than getting any other finite payoff, no matter how good, with the slightly lower probability *q*.

The key idea of the definition is that there is a probability threshold (*p* in the definition). A prospect with a probability ever so slightly lower than the threshold (probability *q* in the definition) cannot be considered better than a prospect with a probability *p* or higher, regardless of the size of the payoff of the prospect with probability *q*.

There are two general approaches to timidity. The first approach is to bind *the value* of the payoff above. According to this approach, the increase in the payoff adds an ever-lowering marginal value, limiting the value that a payoff can have, leading to a trivial marginal value for huge payoffs. This view represents well payoffs like money, a good with a clear decreasing marginal value. However, for the type of good in which longtermists are interested in, additional lives, this approach seems implausible.[3] It seems morally problematic to give a diminishing value to additional lives. If we were to learn that many more people existed in the past than we thought, would it mean that our lives now have less value? This seems highly implausible. I set this view and its objections aside in this paper.

The second approach is to discount the value of prospects with tiny probabilities. There are two basic features of such a view: the kind of threshold used and the kind of discounting applied. Some refinements on the application of the discounting of outcomes avoid some troubling issues faced by the standard view, Nicolausian discounting, which simply gives no weight to discounted outcomes. A more refined approach to discount, tail discounting, gives infinitesimal weight to discounted outcomes. In cases where the non-discounted outcomes are tied, tail discounting uses discounted outcomes as tiebreakers, by applying a normal expected-value analysis of the discounted outcomes to settle the ordering. This approach to discounting avoids some objections (for some examples, see Beckstead & Thomas, 2023, 9-10). I assume tail discounting for the present paper. I focus on objections based on distortions close to the threshold, which can be responded to by using a range threshold.

1.4. Range Threshold

Most objections to discounting views involve outcomes close to the threshold to exploit the strong cutting power of a fixed probability threshold. A range threshold softens this cutting power and permits the avoidance of inconsistencies or irrational orderings by discounting views.

The core justification for using a range threshold is briefly presented by Peterson (2002). He argues that concepts like ‘tiny probability’ are vague and cannot accept a clear distinction. Peterson makes his point with the more concrete but still vague concepts of ‘small city’ and ‘big city’. If a person moved to a small city, the city would still not be big. Even if we repeat the process of adding one person thousands of times, there will never be a point where we can say that the additional person transforms the small city into a big one. A concept like ‘tiny probability’ is similarly vague and cannot accept a fixed threshold to properly deal with borderline cases. This issue is

---

[3] See, e.g., (Williams, 1973) for a defence of the diminishing value of additional years of life.

addressed in the extensive literature on the philosophy of language and the philosophy of sciences regarding how to deal with vagueness.[4]

The main views to address vagueness are the many-valued logic views, supervaluationism, subvaluationism, and contextualism.[5] As a rough overview, many-valued logic views hold that borderline cases have a truth value between full falsehood and full truth. Supervaluationism holds that borderline cases are neither true nor false. Subvaluationism holds that borderline cases are both true and false. Contextualism holds that the set of objects to which a term applies changes according to the context.

I hold the latter view but will assume it to be the most plausible for brevity. What contextualists refer to as being part of the context can refer to many aspects: who is talking or making a decision, what objects are in the set, external conditions, etc. I will emphasize the importance of accounting for the objects of the sets to clarify where to put the threshold for a particular situation.

For instance, Tom is six feet tall. Is he above the threshold for tallness? Being tall is vague and depends on contextual elements. For the general population, six feet is tall, but in the NBA, this is short. Being six feet tall is a borderline case that requires knowing the set of objects to set the threshold for tallness correctly. A way to see it is that we can use a range threshold for tallness: the lower bound is the limit below which nobody can be considered tall regardless of the context, and the upper bound is the limit above which everybody is tall, regardless of the context. Let's say the range threshold for tallness is from 5''5 to 6''5. The specific threshold for a specific case depends on the objects present in the sets. In the general population, the threshold for tallness would be below 6 feet in Canada and above 5''5 feet, but it would be above 6 feet and below 6''5 for NBA players. The same ideas apply to dealing with tiny probabilities.

Practically speaking, setting the range of the threshold for probability requires three steps.

---

4 For a comprehensive overview of the history of the views on the concept of vagueness, see (Williamson, 1996).

5 There is a huge literature on vagueness that I cannot account for in this paper. Here are some important papers for the four main views: the many-valued logic view (e.g. Smith, 2008; Halldén, 1949; Körner, 1955; Kaplan, A. & Schott, H.F, 1951), supervaluationism (e.g., Mehlberg, 1958; Lewis, 1983), subvaluationism (e.g., Hyde, 1997; Varsi, 1997), and contextualism (e.g., Kamp, 1981; Raffman, 1996; Shapiro, 2006, Blome-Tillmann, 2013).

The first step is to set a fixed threshold as a starting point for handling unambiguous cases. For instance, suppose the threshold is set at a probability of one-trillionth, and the range is between one-billionth and one-quadrillionth. If prospect A has a probability of one-millionth and prospect B of one-hundred-quadrillionth, both cases are outside of the range threshold. We should then simply use the fixed threshold.

The second step is to set the width of the threshold with a lower limit below which prospects would have to be discounted regardless of the context and an upper limit above which no prospect could be discounted regardless of the context.

The third step is to set up the rules to adjust the threshold for specific cases. Three rules for adjustments apply to the cases I address.

Rule 1: We can adjust the fixed threshold when the expected value of a prospect above the fixed threshold is comparatively *much smaller* than one of the prospects below the fixed threshold but above the range threshold. This avoids inconsistent ordering when huge payoffs are barely below the threshold compared to tiny payoffs.

Rule 2: We can make adjustments when two prospects of similar payoffs have just a tiny difference in probability, one being slightly above and the other slightly below the fixed threshold. This avoids differential treatment of otherwise similar outcomes.

Rule 3: We should discount prospects with negligible payoff to avoid a prospect with a high probability of getting virtually nothing being preferred to a tiny probability of something of impact.

The objections presented in part 3 are prime examples of when the range threshold ought to be used. Before testing the range threshold, let me justify the need to discount tiny probabilities by showing how a non-discounting approach, recklessness, ought to be discarded as a plausible view to deal with prospects with tiny probabilities.

2. Objections to Recklessness

Recklessness faces five main objections. First, recklessness gives an unreasonable response on the entry price to pay to enter a St-Petersburg gamble. Second, the reckless interpretation of the St-Petersburg gamble contradicts some axiomatic dominance principles. Third, recklessness is obsessed with infinite/huge payoffs. Fourth, infinite/huge outcomes swamp all decisions, even when large losses are almost guaranteed. Finally, recklessness is insensitive to the law of large numbers. The last two objections are novel objections I add to the literature.

2.1. The Entry Price of the St-Petersburg Gamble

The St-Petersburg gamble is a classical objection to unbounded EVT (Bernoulli, 1732). A person is offered a gamble in which a coin is tossed several times. A first head means earning two dollars, and each subsequent consecutive head doubles the price. When the coin ends on tails, the game is over, and the gambler takes the amount of money earned by the previous flip.

For the first toss, getting a head means earning at least \$2 with a probability of 1/2, a second head means getting at least \$4 with a probability of 1/4, and so forth. More formally, the gamble can be represented for *n* number of coin flips by a gain of $\$2^n$ and a probability of $1/2^n$:

E: (1/2 * \$2) + (1/4 * \$4) + (1/8 * \$8) + (1/16 * \$16) …
E: \$1 + \$1 + \$1 + \$1 …

The expected value is the sum of the multiplication of each payoff by its probability. As we can get an infinite number of heads in a row, the expected value of a reckless interpretation of the St-Petersburg prospect is the sum of an infinite number of outcomes with a value of \$1, i.e., an infinite expected value.

The issue with recklessness arises when we ask: how much is one ready to pay to enter the gamble? The rational answer is paying less than the expected value. As the expected value of a reckless interpretation of a St-Petersburg gamble is infinite, reckless people should rationally be ready to pay any amount less than infinite. This is a very counterintuitive and unreasonable answer, as "yet few of us would pay even \$25 to enter such a game" (Hacking, 1980, 563).

Answers compatible with recklessness have been proposed. For instance, Feller argues that a casino must guarantee to have access to an infinite amount of money in case the gambler gets lucky. However, no casino can access an infinite amount of money, so the gamble is unrealistic (Feller, 1968). Samuelson (1961) argues that even if a casino had access to enough money, it would not offer such a gamble because it could lose an infinite amount of money. These solutions are not available for the kind of good longtermism is interested in, e.g., additional lives. There may be an infinite quantity of good accessible in the universe (Bostrom, 2011), and nobody would 'lose' anything if we created an infinite quantity of happy people. Recklessness is helpless against this basic kind of gamble. A discounting view does not face this objection. For instance, if the probability threshold is at one-billionth, the series of possible outcomes stops at around 30 outcomes. As each outcome has a value of $1, the expected value of the gamble is around $30. One should be ready to pay less than $30. That is a reasonable answer.

2.2. Inconsistency with Prospect-Outcome Dominance

Recklessness conflicts with two principles that are "assumed in axiomatic approaches to expected utility theory" (Beckstead & Thomas, 2023, 16). These two principles are *outcome-outcome dominance* and *prospect-outcome dominance*. For the sake of brevity, I only present the latter.

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

Prospect-outcome dominance is a conditional. The principle is proved to be false if the antecedent is true and the consequent is false, for at least one case. A violation of the principle happens when prospects A and B are reckless interpretations of St-Petersburg gambles. Both prospects have an infinite expected value as shown in the first objection, and each outcome of the prospects has necessarily a finite expected value.

Regarding the antecedent, prospect A, having an infinite expected value, is strictly better than each possible outcome of B, which all have a finite expected value. The antecedent is true.

Regarding the consequent, the principle then states that it should be true that prospect A (which has an infinite expected value) is strictly better than prospect B (which has an infinite expected value). However, both are equal, they have an infinite expected value. The consequent is false.

As a conditional with a true antecedent and a false consequent is false, the reckless interpretation of the St-Petersburg gamble violates Prospect-outcome dominance. If the expected value of the St-Petersburg gamble is not infinite, as is the case for timid interpretations, some outcomes of prospect B, the later ones, will always be better than prospect A. The antecedent is thus always false, making the conditional necessarily true for all consequents. In sum, recklessness is in conflict with Prospect-outcome dominance in some instances, but timidity is never.

2.3. Infinite/Large Payoffs Obsession

The third objection shows how recklessness favours prospects with infinite or huge payoffs to the point of obsession. An infinite payoff with a non-zero probability promises an infinite expected value. In this case, any scenario involving an outcome with an infinite payoff that cannot be ruled out as impossible becomes a priority over any finite payoff, regardless of how uncertain the prospect is. This obsession with infinite payoffs makes reckless views prefer unreasonable prospects. The objection is softer for a huge finite payoff, but recklessness still succumbs to it.

This excessive interest in huge payoffs is clearly expressed in Bostrom's (2009) Pascal's mugger thought experiment. An unarmed mugger accosts Blaise Pascal (a reckless person). The mugger tells Pascal to hand over his wallet. Pascal refuses. The mugger argues that she will give Pascal twice the money in his wallet tomorrow if he hands his wallet now. Pascal refuses as the probability of it happening is lower than 0.5. The mugger ramps up the numbers and affirms that she can offer Pascal a huge finite amount of good. As there is a non-zero chance that this is true (we cannot completely rule out this possibility), and because the gain can be so significant, Pascal takes the most rational decision according to recklessness: giving his wallet in exchange for a tiny probability of getting a vast payoff. Pascal received nothing the following day; he got scammed. The main problem with huge and infinite payoffs is that we can rarely be entirely sure that something is impossible: giving zero credence would be overly confident (MacAskill, Bykvist, & Ord, 2020, 152).

This problem shows our epistemic limitations. Is the probability that the mugger respects her promise 10^{-100},100^{-200} or 0? Smith (2014) advises discounting outcomes with tiny probabilities because of our epistemic limitations. When probabilities are very small, how can we distinguish it from zero? Reckless people are ready to support interventions with a great potential payoff on the basis that it is not theoretically impossible that we get a great but unlikely outcome. For instance, Wilkinson argues that it is preferable to let people die now to invest the money that would have saved them in:

> Speculative research into how to do computations using 'positronium'- a form of matter which will be *ubiquitous* in the far future of our *universe*. If our universe has the right structure *(which it probably does not)*, then in *the distant future* we may be able to use positronium to *instantiate all of the operations of human minds* living blissful lives (Wilkinson, 2022, 446) (my emphasis).

Each italic expression refers to a very contentious speculation: it is not theoretically impossible that positronium is compatible with the structure of the universe; it is not theoretically impossible that positronium can become ubiquitous in the universe if we can create it for cheap enough; it is not impossible that we will settle in the reachable universe so the positronium will spread everywhere in it; it is not theoretically impossible to create digital consciousness so positronium can instantiate blissful lives; and it is not impossible that research today will have an impact on the distant future, i.e., if we do not face societal collapse leading to losing the knowledge developed in the year 2023. A decision based on five speculations being merely not impossible does not account sufficiently for our epistemic limitations. All of the conditions have a tiny probability almost indistinguishable from zero for some of them, and a probability of zero for some of them would make investing in positronium a complete waste. If we had epistemic conditions such that we could distinguish between tiny probability and probability zero, expected-value analysis would be a better guide, and recklessness would be a much more plausible view. However, our epistemic limitations will remain, so recklessness will continue to face this obsession with large and infinite payoffs.

2.4. Decision Swamping and Lack of Risk-Aversiveness

I will add to the list of objections to recklessness that it leads to extreme risk-seeking by having outcomes with infinite or huge payoffs swamping decisions, even if all the other outcomes are extremely harmful and overwhelmingly probable. In other words, recklessness is not enough risk-averse.

Let us imagine two prospects, A and B. Prospect A has two potential outcomes, one with an estimated probability *p* of 10^{-30} of deriving an infinite quantity of good, and the other has a probability of 1—*p* of deriving 10^{-20} lives of misery. Prospect B can derive 10^{100} happy lives for sure.

Table 2: Extreme Risk-Seeking

Prospects/probabilities	**p (10^{-30})**	**(1—p)**
Prospect A	Infinite happy lives	10^{-20} lives of suffering
Prospect B	10^{100} happy lives	10^{100} happy lives

Intuitively, prospect B is much better than prospect A because an astronomical quantity of good can be derived for sure. By contrast, prospect A will almost certainly derive a lot of suffering. However, recklessness prefers prospect A, because the tiny probability of an infinite payoff swamps the decision. Epistemic limitations have nothing to do with this case contrarily with the last one. According to recklessness, prospect A has an infinite expected value, while prospect B can only have a finite expected value. Even with an almost certainty of tremendous suffering, recklessness recommends preferring the risky bet.

2.5 Insensitivity to the Law of Large Numbers

The final objection I present is that recklessness is insensitivity to the law of large numbers. Expected-value analyses are a good guide when the expected value of outcomes is representative of the actual outcomes. However, this representativity is only possible when the condition of the law of large numbers is respected. The law of large numbers stipulates that the expected value becomes closer to the actual value as more trials are done (Dekking, 2005).

Let us take the example of a run of rolls of a fair six-sided die to get a basic intuition of this law. The theoretical average value we are expected to get from rolling the die is 3.5. In a run with very few rolls of die, let us say four rolls, the actual average could easily range from 6 to 1. However, as more and more rolls are made, the actual and theoretical average (the expected value) will converge to 3.5. After a couple of hundred dice rolls, the actual average value closely matches the theoretical average of 3.5. Here is a typical representation of the average of dice rolls:

figure 1: Average dice rolls by number of rolls

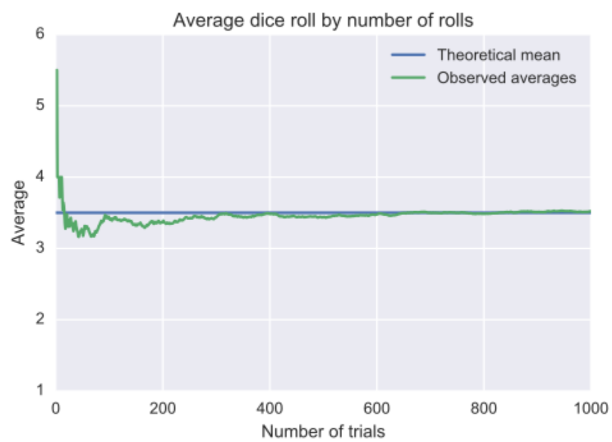

If rolls of dice give us some intuitive sense of the law of large numbers, the difference in the actual value and the expected value is not that problematic in this case. The reason is that dice rolls are not very sensitive to the law of large numbers for two reasons. First, it has a very small number of faces. The smaller the number of faces, the smaller the number of trials to average out losses. Second, each roll gives a positive value rather than a negative value. With negative values, the importance of averaging out the losses becomes ever more important.

The cases in which longtermism is interested are much more sensitive to the law of large numbers. In these cases, the die has hundreds of trillions of faces, and all faces but one have a negative value. For instance, with a donation of $4,000, we have a 1-in-125 trillion of deriving a large quantity of good by avoiding human extinction from a pandemic (Kosonen, 2022, 261). However, this $4,000 could have saved a life for sure. In this case, each roll of dice not hitting the jackpot is synonymous with the death of someone. A die with hundreds of trillions of faces requires hundreds of trillions of rolls (1 roll being a donation of $4,000) for the expected value to get close to the theoretical average. This means that, on average, before hitting the potential jackpot (that might never occur), hundreds of trillions of lives have to be sacrificed. This is especially troubling if we think of our epistemic conditions. For instance, how do we know that the project will not be aborted before we average out the losses? Millions of lives would have been wasted to get no gain. However, there is a lot of uncertainty, leaving room for a reasonable reply by longtermists to this objection.

However, for the Devil's deals thought experiment, we clearly see the irrationality of recklessness. There is a probability of 10^{-21} to get the last Devil's deal, the preferred option by recklessness. Respecting the conditions that the law of large numbers entails for the last deal to be worth the shot would require sextillions of trials to average out the losses. However, recklessness advises going for the last deal with *a single trial*! This is, in my opinion, the worst objection to recklessness, it is insensitive to the number of trials, an essential aspect to account for.

The law of large numbers best explains the paradox of recklessness. The first deals meet the law of large numbers. If one has a probability of 99.7% of getting 1000 years of life, a single trial is sufficient to average the losses. For the first hundreds of deals, a single trial is sufficient, and recklessness is thus a good guide, as is timidity. However, as the probability gets lower, more and more trials are necessary to average the losses and have the expected value not to count on luck. With a probability around 0.1, a couple of trials are more than enough, and for the last deals, sextillions of trials would be necessary. The problem is that the Devil's deal precludes meeting this condition, a single ticket has to be chosen. By contrast, discounting views avoid this issue by discounting the deals that do not meet the necessary conditions to average out the losses. The number of trials is a contextual aspect that the contextualist view I hold accounts for.

3. Objections to Discounting views

As we saw, recklessness entails some very disturbing conclusions. To avoid this conclusion, we have to bind the function in some ways, as all the problems are derived from the function not being bounded. Discounting tiny probabilities is the most promising approach to binding functions. However, as we will see, discounting tiny probabilities is done by putting in place a threshold. However, some important distortions occur around the threshold because of the strong cutting power of a threshold. As we will see, using a range threshold allows us to soften this strong cutting power and avoid these issues.

3.1. Sensitivity to Tiny Difference in Probability Close to the Threshold

The first objection to discounting views is presented by Wilkinson (2022, 460-466). It uses four prospects: P_1, P_2, P_3, and P_4. Probability p is right on the threshold, probability q is just below the threshold by an arbitrarily small margin, and payoff x is minimal (e.g., avoiding a mosquito bite). P_1 derives nothing for sure; P_2 has a probability q of deriving a payoff of 10^{100}, P_3 has a probability p of a payoff x, and P_4 has a probability p of a payoff of 10^{100}. P_2, P_3, and P_4 derive nothing otherwise.

Table 3: Sensitivity to Tiny Difference in Probability

Prospects	Probability	Payoff
P_1	1	0
P_2	q	10^{100}
P_3	p	x
P_4	p	10^{100}

The objection shows two sensitivity issues around the threshold. First, tiny changes in probabilities around the threshold make similar prospects treated arbitrarily very differently. P_2 and P_4 both have the same great payoff, and their probabilities are as close as can be. However, one is discounted, and the other is not. These prospects should be treated relatively similarly because of the tiny difference in probability.

Second, the worst issue is that a prospect with an insignificant payoff just above the threshold can be preferred to a prospect with a huge payoff with an insignificantly lower difference in probability. According to discounting views, P_2 is not better than P_3, even if P_3 has a negligible payoff and P_2 has a huge payoff. The minimal probability difference has too much influence on the ordering.

Response

The leading cause of these issues in the orderings is that the discounting threshold has a too strong cutting power. The objection exploits this strong cutting power in two ways. Prospects below the threshold are treated very differently from prospects above the threshold, even if the difference in probability is negligible or the difference in payoff is immense. A solution to this objection should address the too-strong cutting power of the threshold to make it softer. This is where the range threshold comes into play.

The objection does not involve numerical probability. However, we know that the threshold is p, and q is below and arbitrarily close to p. We can thus safely assume that q is within the range threshold. We can also assume that the threshold extends a bit further on both sides of the fixed threshold.

First, regarding the first ordering issue between P_2 and P_4, we can adjust the threshold based on rule 2; the two prospects have the payoff and an arbitrarily small difference in probability. If we include the slightly lower probability q below the threshold to be above the threshold, a minuscule change, both P_2 and P_4 can now be treated similarly, and P_4 would remain better than P_2, but by a slight margin, as it should be.

Second, regarding the other ordering issue, we can apply rule 3 as P_3 has a negligible payoff, so it can be discounted. Then P_2 becomes preferable to P_3. We can also apply rule 1, the payoff of P_2 is much larger than the payoff of P_3, and the probability of P_2 is below the fixed threshold. We can thus make an adjustment by including probability q to be above the threshold, making P_2 above the threshold. P_2 becomes preferable to P_3, the correct ordering. The change to the threshold is minuscule, and it solves the two issues simultaneously.

One might think that we could change the values to make the objection still relevant to account for my solution. However, no modification of the values of the objection would save it. The most difficult to answer modification is the comparison between P_2 and P_3. The two modifications with the most potential are to reduce the value of probability q, to make it outside of the range threshold so no adjustment would make P_2 above the threshold, and to make the payoff of P_3 non-negligible to avoid it being discounted. With these two modifications, the two solutions I used are out of reach. This is right, with these modifications, my view could not avoid concluding that $P_3 > P_2$. However, I believe that the two modifications make this ordering correct. The first modification makes P_3 much better by making its payoff non-negligible (e.g., saving a life) and P_2 much worse by making its probability much lower, beyond the lower bound of the range threshold. Saving a life with a non-discounted probability is better than saving humanity with a probability so low that it is lower than the low range of the discounting threshold. Reckless people certainly disagree with this ordering. However, this timid ordering is not irrational or unreasonable as was the ordering of the not modified version of the objection. This timid ordering is just a reasonable timid ordering, as is the reasonable ordering of recklessness of the same case. The difference in ordering does not mean unreasonableness or irrationality for one of the views. This modified version is the best version of this objection, making this kind of argument toothless for my view.

3.2. Inconsistency with Background Independence

The second objection resembles the Egyptology objection addressed to the average view and the variable value theory in population ethics.[6] In the average view, when we wonder whether adding a person to a population is good, her well-being has to be above the average. To identify the average, we must know how well-off people are and have been. This includes how well-off people were in Ancient Egypt (hence the name of the objection), everywhere on Earth at all times, and in corners of the universe. The objection is that knowledge about the well-being of people we cannot influence (past people and individuals in unreachable corners of the universe) should not

[6] The original version of this criticism was made by MacMahan (1981) and the acknowledgement of this resemblance has been introduced by Wilkinson (2022).

influence our moral decisions today. Should it matter for our ethical decision today if the hunter-gatherers were much happier than we thought?

The same objection is addressed to timid views bounding the value of the payoff above. According to this view, the marginal value of creating additional people diminishes. To know how valuable additional lives are, we must know how far we are from the upper limit. Thus, a view bounding the value of the payoff above requires knowing how well-off and how many people and aliens have been and still exist (Beckstead and Thomas, 2023, 11). These views do not respect the separability of persons, the principle according to which a person can be valued independently of other people's conditions (See Broome, 2004, chap 8).

The Objection To Discounting Views

According to Beckstead and Thomas (2023, 11-12) and Wilkinson (2022, 466-473), we can reach a similar objection to discounting views with a more refined argument. The objection is based on the incompatibility of discounting views with the following principle:

> **Background Independence**: For any prospects P_a and P_b and any outcomes O', if $P_a \succcurlyeq P_b$, then $P_a - O' \succcurlyeq P_b - O'$.[7] (Wilkinson, 2022, 467)

Background Independence entails that if we subtract (or add) the same quantity of good from both P_a and P_b, the ordering before the manipulation ($P_a \succcurlyeq P_b$) should remain the same as after the removal ($P_a - O' \succcurlyeq P_b - O$). To show how discounting views violate this principle, let us take the version of the objection proposed by Beckstead and Thomas (2023, 11-12). The objection involves three prospects: A, B, and C. Payoff N is a trillion times the size of *n*. Probability *p* is slightly below the threshold, and *q* is an extremely tiny probability, less than a trillionth the size of *p*. A key feature of the objection is that the probability *p+q* is precisely on the probability threshold. Prospect A has a probability *p* of deriving *n* lives. Prospect B has a probability of *p+q* of deriving *n* lives. Prospect C has a probability *p* of deriving *n+N* lives:

[7] I changed the formulation made by Wilkinson of the background independence principle from an addition to a subtraction as the example with a subtraction is more intuitive to understand and Beckstead & Thomas use an example with a subtraction that I use. I also use the term 'prospect' instead of 'lottery' for uniformity purposes in this paper.

Table 4: Payoffs for Different Prospects: Original Background Independence

Options/ Probabilities	**p**	**q**	**1-p-q**
A	n	0	0
B	n	n	0
C	n + N	0	0

According to timidity, prospect B is at least as good as prospect C because the probability of C is below the threshold while the probability of B is above, and prospect A is inferior to C because of the payoff size. The ordering is thus B ≽ C ≻ A. Beckstead and Thomas then modify the case and say that the outcomes of *n* with probability *p* were, all along, in another galaxy where we have no influence, so payoff *n* with probability *p* has to be subtracted from all the prospects. Here is the modified table.

Table 5: Payoffs for Different Prospects: Modified Background Independence

Options/ Probabilities	**p**	**q**	**1-p-q**
A	0	0	0
B	0	n	0
C	N	0	0

The same amount of good is removed from all prospects, so the ordering of these prospects should remain the same according to the Background Independence principle. However, if we keep the ordering, we have an irrational ordering. In the initial ordering, B cannot be worse than C. However, after the modification, B is a trillion times less probable and produces a trillion times less value. C should now be clearly better than B. Moreover, the objection allegedly also shows that the ordering of discounting views depends on what happens in faraway galaxies, a very counterintuitive implication for any views.

Response to The Objection

There are two issues raised by this objection: the dependence on what happens in other galaxies and the issue related to the ordering of prospects.

First, the dependence on what happens in other galaxies is fallacious for the version addressed to discounting views. Discounting views respect the separability of persons: how many people existed, exist, or will exist, or how happy they are has no influence on probabilities. Discounting views do not depend on what happens in other galaxies, they depend on the subtraction of probabilities of some outcomes making a prospect above the threshold be below the threshold after the subtraction. Thus, what is removed can be anything like something happening in another galaxy in the objection. However, they could have said anything about the n lives created with probability p. They could have been new characters in Harry Potter. Would this mean that discounting views depends on fictional characters? No, it just means that the ordering of prospects by discounting views is sensitive to the removal of some consequences on all prospects[8].

Second, to avoid violating Background Independence[9], we have two opposite strategies to avoid the issues. On the one hand, we can make the threshold higher by a value q (a very tiny probability). This way, all prospects are discounted in both versions. If you remember, when the non-discounted outcomes are equal, discounted outcomes are given an infinitesimal weight to be used as tiebreakers. A standard expected-value analysis is applied to the discounted outcomes. As all the prospects in both versions are below the threshold, tail discounting would use the tiebreak procedure for the ordering by applying a normal expected-value analysis. This leads to the ordering of C > B > A for both versions. On the other hand, we could also make the threshold a bit lower by a value q so probability p alone is above the threshold. Again, the ordering would be C > B > A before and after the modification. In sum, with a tiny change in the threshold, either by making it slightly lower or higher, the issue evaporates.

Modifications of the objection would not save it. Modifications of the value of the payoffs would not change anything. The essence of the objection is that $p+q$ is equal to the threshold. Thus, ef-

[8] Russell (2023) gives another explanation of why the argument is implausible.

[9] The case meets all the criteria to use the range threshold.

fective modifications to attack my view would make the distribution of the probability between *p* and *q* different, making the values either closer or farther apart. On the one hand, if one makes *p* larger and *q* smaller, it would make the adjustment of lowering the threshold easier as we would remove a smaller *q* for the first strategy I used. On the other hand, if we make *p* smaller and *q* larger, this changes nothing for the second strategy, as we just have to make *p+q* below the threshold slightly higher by a micro adjustment to discount all prospects and use the tail discounting. Because we can go both ways, making the threshold slightly lower or higher, no modification of the values of *p* and *q* saves the objection.

2.3. Objection 3: Continuity Axiom

The last objection to discounting views is presented by (Kosonen, 2022, 200-201) and involves the violation of the Continuity axiom. Let us assume that for all prospects X and Z, XpZ represents the prospect of a lottery where X has a probability *p* of realizing and Z has a probability 1 — *p* of realizing. Continuity is defined as follows:

> ***Continuity***: If X >Y > Z, then there are probabilities p and q ∈ (0, 1), such that
> XpZ > Y > XqZ.

The principle is a conditional. Thus, it can be proved to be false if the antecedent is true and the consequent is false, for at least one case. Let's start by creating a true antecedent. Suppose a prospect A with a probability *t,* right on the probability threshold. Let us call it A_t. With A_t, we can now create a new case representing the antecedent of Continuity by replacing X with A_t: A_t >Y > Z. We can easily imagine a case where it is true: A_t has a tiny but not discounted probability of a huge payoff, Y has a higher probability but a much lower expected value making it worse than A_t, and Z does no good for sure.

To violate the principle, we have to show the consequent to be false, i.e., show that XpZ > Y > XqZ is false or, in our case in which X is replaced by A_t, that A_tpZ > Y > A_tqZ. As A_t has a probability on the probability threshold, multiplying it by any possible value of *p* outside of 1 makes A_tpZ discounted. As Y is not discounted (the discount rate can always be adjusted so that), it is false that A_tpZ > Y. In this case, the consequent A_tpZ > Y > A_tqZ is false for the true antecedent A_t >Y > Z.

Response

This violation only occurs with the use of a clear-cut discount. It can be easily avoided by using a continuous discount function.

A clear-cut discount may be formalized mathematically by multiplying the utility function of a prospect X by a discontinuous « step-function », widely called the Heaviside function:

$$H(p-t) := \begin{cases} 0 & \text{si} \quad p < t \\ 1 & \text{si} \quad p \geq t \end{cases} \qquad \text{Equation (1)}$$

We then obtain a discounted prospect X_d, whose utility is a function of probability p:

$$X_d(p) = X * H(p-t) \qquad \text{Equation (2)}$$

As a result, Xd=0 on the domain p ∈ (0,t) and Xd=X for p ∈ (t,1]. Xd is effectively discontinuous and, therefore, violates the continuity axiom, which is necessary to appropriately use expected value theory.

Discontinuous jumps rarely occur in nature. For example, a car never accelerates instantly. Introducing them in equations often leads to artificial mathematical glitches. For these reasons, it is common among physicists and other scientists to approximate the Heaviside step function by a similar continuous function. For example, the step function may be approximated with a polynomial or an exponential function. By adjusting the parameters, we can create a smooth step function arbitrarily similar to a perfect Heaviside step function while respecting continuity (Figure X). This function can then be used instead of the Heaviside function in equation 2, creating a utility curve X_d(p) that respects the continuity axiom, for any continuous prospect X.

Figure 2: Continuous functions approaching the Heaviside step function, for increasingly narrow step width

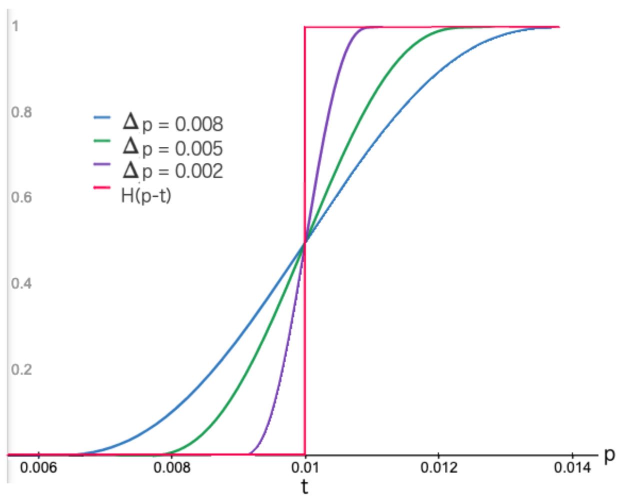

$y = \left(\left(p - t - \frac{\Delta p}{2}\right)/\Delta p\right)^3 \cdot \left[10 - 15\left(\left(p - t - \frac{\Delta p}{2}\right)/\Delta p\right) + 6 \cdot \left(\left(p - t - \frac{\Delta p}{2}\right)/\Delta p\right)^2\right]$, probability threshold $t = 0.001$, step function defined within the range Δp. In red: the Heaviside step function.

As the discount function can be adjusted to be arbitrarily close to its extremely discontinuous version, the "width of the step" can be arbitrarily narrowed, for example, to fit between two prospects. Then, in the context of calculations not involving this sensitive zone, the results will be the same as for using a perfect step function, which can then be used for simplicity's sake. Note that it can be shown that the continuity axiom would also be preserved if using any other continuous discount functions such as an exponential discount.

4. Conclusion

In conclusion, I start by summarizing the objections to timidity and its answers and then summarize the objections to recklessness. After this summary, I present a concluding analysis before presenting the implications of my argument for longtermism.

In part 2, we saw that recklessness faces at least five objections. First, recklessness entails paying any finite amount of money to participate in St-Petersburg-type gambles, an unreasonable conclusion. The St-Petersburg gamble has nothing special about it, so the issue that recklessness faces generalizes to a large variety of decision situations. Second, recklessness violates Prospect-outcome dominance, a principle assumed in axiomatic approaches to EVT. Third, Reckless people are obsessed with infinite/huge payoffs, leading to the unreasonable prioritization of prospects with promises of huge/infinite gains with a probability that cannot be ruled out. Fourth, recklessness entails making very risky decisions when outcomes have a large enough payoff, even if the probability of the gain is infinitesimal and all the other payoffs are extremely bad and almost certain. Fifth, recklessness is insensitive to the law of large numbers. When there are not enough trials, outcomes with a too small probability of getting a positive payoff do not have their losses averaged out, leading, most likely, to a clear loss, unless the gambler is very lucky. Counting on luck is not a good approach for an effective decision theory approach.

In part 3, we saw that discounting views face three major objections. The first objection shows that discounting views are too sensitive to tiny differences in probabilities close to the threshold and not enough sensitive to large differences in payoffs. The second objection involves prospects from which the same consequence is subtracted. The idea is that the ordering before the subtraction should remain the same after the subtraction. However, by removing the same consequence from all prospects, some prospects above the threshold become below the threshold. This leads to inconsistency in the ordering. For both objections, a range threshold avoids the inconsistencies by accounting for the specific objects of the set of the objections. The third crucial objection is that discounting views violate the Continuity axiom. I have shown that discontinuity can easily be avoided by approximating the step function with a polynomial function, smoothing the step.

In sum, Recklessness faces five objections, and I have shown that the three main objections to discounting views can be avoided. There are other plausible objections addressed to timidity that I do not have the space to respond to. For instance, here are two of them: the Indology objection[10] (Wilkinson, 2022, 473-475) and the objection that small decreases in unlikely payoffs cannot be outweighed (Beckstead & Thomas, 2022, 10). Both of these objections can be answered with a similar response to what I gave to the Background assumption objection. New objections that my view cannot answer might be developed. However, discounting views have a lot of flex-

[10] Russell (2023) shows some other issues with this objection to timidity.

ibility to address such challenges. Anyhow, in the current state of research, recklessness seems to obviously be a worse view than a refined discounting view.

As stated in the introduction, the longtermist comparative argument supporting the moral obligation of individuals to support long-term interventions depends on recklessness. However, as we have seen in this paper, recklessness is less plausible than the discounting view I defended. The argument on which Greaves and MacAskill base strong longtermism is thus in jeopardy for individuals. The debate on recklessness leads to the intuitive conclusion that societal projects like protection against existential risks are a duty for societies and that it should be morally permissible for individuals to support other existing individuals effectively with their donations or by choosing careers based on their preference rather than having the duty to maximize good for the future. In sum, even consequentialist views do not entail longtermism for individuals.

Bibliography

Beckstead, N. (2013). On the overwhelming importance of shaping the far future. Dissertation, Rutgers University.

Bernoulli, N. (1732/1954). 'Letter #18 to Daniel Bernoulli'. In Van der Waerden, 566–8.

Blome-Tillmann, M. (2013). Pacific Philosophical Quarterly, vol. 94, no. 1, 89–100

Bostrom, N. (2013). Existential Risk Prevention as Global Priority. Global Policy, vol. 4, no. 1, 15–31

Bostrom, N. (2011). Infinite Ethics. *Analysis and Metaphysics*, vol. 10, 9–59.

Bostrom, N. (2009). Pascal's mugging. *Analysis*, vol. 69, no. 3, 443–45.

Bostrom, N. (2002). Existential Risks: Analyzing Human Extinction Scenarios and Related Hazards. Journal of Evolution and Technology, vol. 9.

Broome J. (2004). Weighting lives. Oxford University Press.

Feller, W. (1968). *An Introduction to Probability Theory and its Applications Volume I, II.* Wiley.

Greaves, H., MacAskill, W., & Thornley, B. (2021). The Moral Case for Long-term Thinking. In Cargill & Tyler (Eds.), *The Long View: Essays on Policy, Philanthropy, and The Long-Term Future*. First Edition.

Greaves, H. & MacAskill, W. (2021). ‘The Case for Strong Longtermism’. Working paper. The Future of Humanity Institute.

Hacking, I. (1980). Strange Expectations. *Philosophy of Science*, 47, 562–7.

Halldén, S. (1949). The Logic of Nonsense. Upsala Universitets Arsskrift.

Kamp, H. (1981). “The Paradox of the Heap”, Uwe Münnich (ed.), 1996, *Aspects of Philosophical Logic*, Cambridge: Cambridge University Press, 225–77.

Kaplan, A. & Schott, H.F., (1951). ‘A calculus for empirical classes’, Methodos 3, 165–88.

Körner, S. (1955). Conceptual Thinking. Cambridge, Cambridge University Press.

Kosonen, P. (2022). Tiny probabilities of Vast Value. Ph. D. Thesis, Worcester College, Oxford University.

Lewis, D.K. (1983). Philosophical Papers. Vol. 1, Oxford, Oxford University Press.

MacAskill, W., Bykvist, K., & Ord, T. (2020). *Moral Uncertainty*. Oxford University Press.

MacAskill, W. (2022). *What We Owe The Future: basic* Books/Oneworld Publications.

MacAskill, W. (2019). ‘The Definition of Effective Altruism’. In *Philosophical Issues: Effective Altruism* (Engaging Philosophy). Oxford University Press, 10–28.

MacAskill, W. (2015). *Doing Good Better: Effective Altruism and a Radical New Way to Make a Difference.* Random House.

McMahan, J. (1981). 'Problems of population theory'. *Ethics*, vol. 92, no. 1, 96–127. https://doi.org/10.1086/292301

Mehlberg, H. (1958). The Reach of Science. Toronto. Toronto University Press.

Mogensen, A. L. (2022). 'The only ethical argument for positive d?'. *Philosophical Studies*, vol. 179, 2731–50.

Ord, T. (2020). The Precipice: Existential Risk and the Future of Humanity. Hatchette, 480 pages.

Peterson, M. (2002). 'What is a de Minimis risk?'. *Risk Management*, vol. 4, 47–55.

Raffman, D. (1996) 'Vagueness and context-sensitivity'. *Philosophical Studies*, vol. 81, 175–92.

Russel, J. S. (2023). On two arguments for Fanaticism. Noûs, online version before inclusion in an issue.

Samuelson, P. (1960). 'The St. Petersburg Paradox as a Divergent Double Limit'. *International Economic Review*, vol. 1, no. 1, 31–7.

Shapiro, S. (2006). *Vagueness in Context.* Oxford: Oxford University Press.

Singer, P. (2015). The Most Good You Can Do: How Effective Altruism is Changing Ideas About Living Ethically. Yale University Press.

Singer, P. (1972). Famine, affluence, and morality. *Philosophy and Public Affairs*, vol. 1, 229–243.

Smith, N. J.J. (2014). Is evaluative compositionality a requirement of rationality? *Mind*, 457–502.

Smith, N. J. J. (2008). *Vagueness and Degrees of Truth*, New York: Oxford University Press.

Tarsney, C. (2023). The epistemic challenge to longtermism. *Synthese, vol.* 201, no. 6, 1-37

Temkin, L. (2015). Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning: A Précis. *Journal of Moral Philosophy*, vol. 12, no. 4, 363–92.

Temkin, L. S. (2012). Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning. Oxford University Press.

Wilkinson, H. (2022). In Defense of Fanaticism. *Ethics*, 132, no. 2, 445–77.

Williams, B. (1973). Problems of the Self. Cambridge University Press.

Williamson, T. (1996). Vagueness. Routledge.