# How to Discount Small Probabilities\*

### Petra Kosonen<sup>†</sup>

#### November, 2023

ABSTRACT: Maximizing expected value leads to counterintuitive recommendations in cases that involve tiny probabilities of huge payoffs. In response, some have argued that we ought to discount very small probabilities down to zero. There are many ways of doing this. In this paper, I discuss how, exactly, this view should be formulated and what problems the different versions of this view face. I will argue that the most straightforward ways of discounting small probabilities face significant problems. Among these problems are, for example, violations of dominance and acyclicity. I conclude by discussing more plausible versions of discounting small probabilities that avoid these violations.

Orthodox decision theory states that a rational agent always maximizes expected utility. However, this seems to lead to counterintuitive recommendations in cases that involve very small probabilities of huge payoffs. In these cases, an option can be great in expectation even if the probability of obtaining a valuable outcome is tiny, as long as the valuable outcome is great enough. One example of such a case is the *St. Petersburg game*:<sup>1</sup>

<sup>\*</sup>I wish to thank Jean Baccelli, Tomi Francis, Andreas Mogensen, Jake Nebel, Teruji Thomas and the audience of the Global Priorities Institute seminar for valuable feedback and discussions.

<sup>&</sup>lt;sup>†</sup>Population Wellbeing Initiative, University of Texas at Austin. I would be grateful for comments: kosonenpetra@gmail.com

<sup>&</sup>lt;sup>1</sup>A version of this game was originally proposed by Nicolaus Bernoulli in 1713. The game was simplified by Gabriel Cramer in 1728 and published by Daniel Bernoulli in 1738. See Pulskamp (n.d.) and Bernoulli (1954).

**St. Petersburg game:** A fair coin is flipped until it lands on heads.

The prize is  $2^n$ , where *n* is the number of coin flips.

The St. Petersburg game gives a 1/2 probability of \$2, a 1/4 probability of \$4, a 1/8 probability of \$8, and so on. Consequently, its expected monetary value is infinite:

$$2 \cdot \frac{1}{2} + 4 \cdot \frac{1}{4} + 8 \cdot \frac{1}{8} + \dots = 1 + 1 + 1 + \dots = \infty.$$

Therefore, agents who maximize expected monetary value would pay any finite amount to play the game—but this seems counterintuitive.<sup>2</sup> Furthermore, if this game's (monetary) value is infinite, one would value it higher than any of its possible finite payoffs, which seems irrational.<sup>3</sup>

Another case that involves tiny probabilities of huge payoffs is Pascal's Hell:<sup>4</sup>

**Pascal's Hell:** Satan offers Pascal a deal: If a coin lands on heads, Satan will create a million Graham's number of happy Earth-like planets. But if the coin lands on tails, then everyone on Earth will suffer excruciating pain until life on Earth is no longer possible. The probability of heads happening is one-in-Graham's-number.

Should Pascal accept the offer? The probability of the payoff is tiny. However, as the possible payoff is enormous, Pascal is forced to conclude that the expected value of accepting the offer is positive.<sup>5</sup> Consequently, humanity will almost certainly suffer excruciating pain until life on Earth is no longer possible.

<sup>&</sup>lt;sup>2</sup>Pulskamp (n.d., p. 6). Daniel Bernoulli argues that, due to the diminishing marginal utility of money, one should not pay any finite sum to play the St. Petersburg game. See Bernoulli (1954). However, Menger (1967, pp. 217–218) shows that if utilities are unbounded, one can always create a *Super St-Petersburg game*, in which the payoffs grow sufficiently fast so that the expected utility of the game is infinite. See also Samuelson (1977, §2).

<sup>&</sup>lt;sup>3</sup>Huemer (2016, pp. 34–35) and Russell and Isaacs (2021). There are variants of the St. Petersburg game that do not seem to make any sense by the lights of Expected Utility Theory because they have no unique expected utility. See, for example, Nover and Hájek (2004) on the Pasadena game.

<sup>&</sup>lt;sup>4</sup>This case is slightly modified from Kosonen (2022, pp. 2-4). It is based on Bostrom's (2009) *Pascal's Mugging*, which in turn is based on informal discussions by various people, including Eliezer Yudkowsky (2007*b*). For criticism of Pascal's Mugging, see Hanson (2007), Yudkowsky (2007*a*) and Baumann (2009).

<sup>&</sup>lt;sup>5</sup>This may not hold if Pascal maximizes expected utility and utilities are bounded as standard

In response to cases like this, some have argued that we ought to discount very small probabilities down to zero—let's call this *Probability Discounting*.<sup>6</sup> Nicolaus Bernoulli first proposed this idea in response to the St. Petersburg game. He writes: "[T]he cases which have a very small probability must be neglected and counted for nulls, although they can give a very great expectation. [...] This is a remark which merits to be well examined."<sup>7</sup> Recently, Smith (2014) and Monton (2019) have also defended Probability Discounting. Monton argues that one ought to discount very small probabilities down to zero, while Smith argues that it is rationally permissible—but not required—to do so.<sup>8</sup> However, we do not yet have a well-specified and plausible theory that tells us how to discount small probabilities. As Monton writes: "I don't have a perfectly rational, reasonable decision theory to hand you just yet (sorry)."

This paper discusses how Probability Discounting can be formulated and what the most plausible version of it might look like. §1 discusses a simple version of Probability Discounting on which one should conditionalize on outcomes associated with tiny probabilities not occurring. I show that this view faces a problem with individuating outcomes, and it also violates dominance. §2 discusses a version of Probability Discounting that considers very-small-probability outcomes as tiebreakers when prospects would otherwise be equally good. I show that this

**Probability Fanaticism:** For any probability p > 0 and any (finitely) good outcome o, there is some great enough outcome O such that probability p of O (and otherwise nothing) is better than certainty of o.

<sup>7</sup>Pulskamp (n.d., p. 2). Other proponents of Probability Discounting include, for example, Buffon and Condorcet. See Hey et al. (2010) and Monton (2019, pp. 16–17).

<sup>8</sup>Smith argues that discounting small probabilities down to zero is a way of getting a unique expected value for the Pasadena game. See Nover and Hájek (2004).

<sup>9</sup>Monton (2019, p. 15).

axiomatizations of expected utility maximization (such as the von Neumann-Morgenstern utility theorem) require. See Kreps (1988, p. 63).

<sup>&</sup>lt;sup>6</sup>See Hájek (2014), Isaacs (2016), Kosonen (2022, pp. 137–239) and Cibinel (forthcoming) for criticism of discounting small probabilities. Also, see Beckstead (2013, ch 6), Goodsell (2021), Russell and Isaacs (2021), Wilkinson (2022), Russell (forthcoming) and Beckstead and Thomas (forthcoming) for discussions of issues related to *Probability Fanaticism*:

view also violates dominance. §3 discusses a version of Probability Discounting on which one should conditionalize on very-small-probability *states* not occurring. I discuss three ways of specifying this view. I show that they violate either dominance or acyclicity (or both). §4 discusses more plausible versions of Probability Discounting that avoid the earlier violations of dominance and acyclicity.

## **1** Naive Discounting

This section discusses a version of Probability Discounting on which one should conditionalize on outcomes associated with tiny probabilities not occurring. I show that this view faces a problem with individuating outcomes and it also violates dominance.

According to Probability Discounting, an agent is rationally required or permitted to discount very small probabilities down to zero. On this view, there is some discounting threshold t such that probabilities below this threshold are discounted down to zero.<sup>10</sup> But when are probabilities small enough to be discounted? Or, as Buffon writes: "[O]ne can feel that it is a certain number of probabilities that equals the moral certainty, but what number is it?"<sup>11</sup> Some possible discounting thresholds have been suggested. For Buffon and Condorcet, the discounting thresholds were 1/10,000 and 1/144,768 (respectively), while for Monton, this threshold is approximately 1 in 2 quadrillion.<sup>12</sup> As Monton argues, the discounting threshold is plausibly subjective. There is no objective answer to Buffon's question.

<sup>&</sup>lt;sup>10</sup>Alternatively, one could discount probabilities up to and including the threshold t. Note that this threshold might be vague.

<sup>&</sup>lt;sup>11</sup>Hey et al. (2010, p. 256).

<sup>&</sup>lt;sup>12</sup>Buffon's discounting threshold was the probability of a 56-year-old man dying in 24 hours—an outcome reasonable people typically ignore (Monton, 2019, pp. 8–9). Condorcet had a similar justification for his threshold (Monton, 2019, pp. 16–17). Monton's (2019, p. 17) discounting threshold is between  $1/2^{50}$  and  $1/2^{51}$ , as he treats the probability of getting tails at least 50 times in a row (with a fair coin) as rationally negligible.

Instead, it is up to each individual where the discounting threshold is.<sup>13,14</sup>

So, on this view, one should discount small probabilities—but small probabilities of *what*? This paper discusses versions of Probability Discounting that ignore very-small-probability outcomes or states.<sup>15</sup> I will begin with the former views. There are many ways of ignoring outcomes associated with small probabilities. One way to ignore the very-small-probability outcomes of some prospect A would be to treat A as interchangeable with a prospect B, which really does assign probability zero to these outcomes.<sup>16</sup> However, B cannot assign the same probabilities as A to the remaining outcomes; otherwise, the sum of all the probabilities assigned to outcomes of B would be less than one.<sup>17</sup> Instead, the probabilities assigned by B can be obtained from those assigned by A by conditionalizing on the supposition that some outcome of non-negligible probability occurs, where 'non-negligible' means a probability that is at least as great as the discounting threshold.<sup>18</sup>

Let  $EU(X)_{pd}$  denote the expected utility of prospect X when tiny probabilities have been discounted down to zero (read as 'the probability-discounted expected utility of X'). A *prospect* is taken to be a situation that may result in different outcomes with different probabilities. One of the simplest versions of Probability Discounting—let's call it *Naive Discounting*—states:

**Naive Discounting:** Prospect X is at least as preferred as prospect Y if and only if  $EU(X)_{pd} \ge EU(Y)_{pd}$ , where  $EU(X)_{pd}$  and  $EU(Y)_{pd}$ 

<sup>&</sup>lt;sup>13</sup>The subjectivity of the threshold may be reasonable for individuals' rational preferences. But it seems less so in the context of ethics when we are asking which prospects are better or worse.

<sup>&</sup>lt;sup>14</sup>Smith (2014) holds that the threshold might not apply to simple prospects, that is, prospects that assign a non-zero probability to only finitely many outcomes. Also, Smith maintains that this threshold may be different in different situations.

<sup>&</sup>lt;sup>15</sup>Whether one ignores very-small-probability outcomes or states makes a difference in some cases. A very-small-probability state might result in an outcome that overall has a non-negligible probability (when one also considers the other states). In that case, the state is associated with a negligible probability but the outcome is not.

<sup>&</sup>lt;sup>16</sup>Smith (2014, p. 478).

<sup>&</sup>lt;sup>17</sup>Ibid.

<sup>&</sup>lt;sup>18</sup>Ibid.

are obtained by conditionalizing on the supposition that some outcome of non-negligible probability occurs.<sup>19</sup>

OUTCOME INDIVIDUATION PROBLEM. So, on Naive Discounting, one should conditionalize on very-small-probability outcomes not occurring—but what counts as an 'outcome'? In particular, Naive Discounting faces the following problem:<sup>20</sup>

**Outcome Individuation Problem:** If we individuate outcomes with too much detail, all outcomes have negligible probabilities. Is there a privileged way of individuating outcomes that avoids this?

The most obvious non-arbitrary way of individuating outcomes is by their utilities:<sup>21</sup>

**Individuation by Preference:** Outcomes should be distinguished as different if and only if one has a preference between them.

Following this principle, each final utility level that a prospect might result in is considered a distinct outcome, and the possibilities of these outcomes are ignored if their associated probabilities are below the discounting threshold.

However, individuating outcomes by their utilities might result in ignoring all possible outcomes of some prospect if all its final utility levels are very unlikely. In response to such cases, agents might lower their discounting thresholds until at least some outcomes have non-negligible probabilities. However, in cases where all outcomes have a zero probability, it is not possible to do so (except, of course, by

 $<sup>^{19}</sup>$  Note that  $EU(X)_{pd}$  and  $EU(Y)_{pd}$  are obtained by conditionalizing, potentially, on different events not occurring.

<sup>&</sup>lt;sup>20</sup>See also Beckstead and Thomas (forthcoming, p. 13).

<sup>&</sup>lt;sup>21</sup>Contrast Individuation by Preference with a similar principle presented by Broome (1991, p. 103):

**Principle of Individuation by Justifiers:** Outcomes should be distinguished as different if and only if they differ in a way that makes it rational to have a preference between them.

not discounting at all).<sup>22</sup> Imagine, for example, an ideally shaped dart thrown on a dartboard, where each point results in a different utility. The probability that the dart hits a particular point may be zero. But one should not ignore every possible outcome of throwing the dart. Nevertheless, one might argue that we need not worry about cases where all outcomes have a zero probability because they are rare in practice. In all (or near all) cases we care about, some outcomes have non-zero probabilities.

STATEWISE DOMINANCE. Some might be satisfied with the above solution to the Outcome Individuation Problem. However, besides this problem, Naive Discounting also violates dominance. More precisely, Naive Discounting violates the following dominance principle:<sup>23</sup>

**Statewise Dominance:** If the outcome of prospect X is at least as preferred as the outcome of prospect Y in all states, then X is at least as preferred as Y(Weak Statewise Dominance). Furthermore, if in addition the outcome of X is strictly preferred to the outcome of Y in some possible state, then X is strictly preferred to Y(Strong Statewise Dominance).

Statewise Dominance is very plausible. If some prospect is sure to turn out at least as well as another prospect, but it might turn out better, then that prospect should be better.<sup>24</sup>

To see how Naive Discounting violates Statewise Dominance, consider the following prospects (see table 1):<sup>25</sup>

<sup>&</sup>lt;sup>22</sup>Beckstead and Thomas (forthcoming, pp. 12–13).

<sup>&</sup>lt;sup>23</sup>Savage (1951, p. 58) and Luce and Raiffa (1957, p. 287).

<sup>&</sup>lt;sup>24</sup>Russell (forthcoming, p. 9) writes on Strong Statewise Dominance: "What if Statewise Dominance fails? In that case, I'm not sure what we're doing when we compare how good prospects are. [...] [W]hat we ultimately care about is how well things turn out; choosing better prospects is supposed to guide us toward achieving better outcomes. In light of this, if dominance reasoning is wrong, then I don't want to be right. If A is sure to turn out better than B, then this tells us precisely the thing that betterness-of-prospects is supposed to be a guide to."

<sup>&</sup>lt;sup>25</sup>On discounting small probabilities and dominance violations, see Isaacs (2016), Smith (2016),

#### Naive Statewise Dominance Violation:

*Prospect A* Gives \$1 million in state 1 and nothing in state 2.

*Prospect B* Gives nothing in both states.

Suppose the probability of state 1 is below the discounting threshold. After conditionalizing on the supposition that some outcome of non-negligible probability occurs, A is substituted by B. One would then be indifferent between A and B, even though the outcomes of A and B are equally good in state 2, but the outcome of A is better than the outcome of B in state 1.

TABLE 1	
NAIVE STATEWISE DOMINANCE	VIOLATION

I State 2
hold $1-p$
on \$0
\$0

To summarize, Naive Discounting states that one should conditionalize on not obtaining very-small-probability outcomes. This view faces the Outcome Individuation Problem, which can be solved by individuating outcomes by their utilities (except in cases where all outcomes have a zero probability). However, Naive Discounting also faces another problem: It violates Statewise Dominance.<sup>26,27</sup>

Monton (2019, pp. 20-21), Lundgren and Stefánsson (2020, pp. 912-914) and Beckstead and Thomas (forthcoming, §2.3).

<sup>&</sup>lt;sup>26</sup>One might object that probability discounters need not worry about violating Statewise Dominance because the whole point of the view is to treat some nonzero probabilities as though they were zero; so if some good thing comes with such a probability, it should contribute nothing to the overall value of the prospect.

<sup>&</sup>lt;sup>27</sup>Hájek (2014) shows that Expected Utility Theory also violates Statewise Dominance in cases that involve possible states of zero probability. Monton (2019, §7) argues that violations of Statewise Dominance should not count against Probability Discounting, given that Expected Utility Theory violates Statewise Dominance too. Easwaran (2014, p. 14) and Hájek (2014, p. 557) suggest that

## 2 Lexical Discounting

This section discusses a version of Probability Discounting that treats very-smallprobability outcomes as tiebreakers when prospects would otherwise be equally good. This view avoids the previous violation of Statewise Dominance. However, as I will show, it violates Statewise Dominance in another case.

There is a straightforward solution to the previous case: Treat outcomes whose probabilities are below the discounting threshold as tiebreakers. Then, A is better than B because A and B have equal probability-discounted expected utility but, in addition, A gives a negligible probability of a positive outcome (while B does not). More generally, in tied cases, prospects can be compared by their expected utilities without any discounting (like Expected Utility Theory would do).

On this proposal, prospects are first ranked by their probability-discounted expected utilities. Then, in cases of ties, these prospects are ranked by their expected utilities without discounting small probabilities. Formally this view—let's call it *Lexical Discounting*—states the following:

**Lexical Discounting:** Prospect X is at least as preferred as prospect Y if and only if

- $\succ EU(X)_{pd} > EU(Y)_{pd}$  or
- $\succ EU(X)_{pd} = EU(Y)_{pd}$  and  $EU(X) \ge EU(Y)$ ,

where  $EU(X)_{pd}$  and  $EU(Y)_{pd}$  are obtained by conditionalizing on the supposition that some outcome of non-negligible probability occurs.

It is slightly misleading to say that Lexical Discounting is a form of discounting small probabilities down to zero because small probabilities and their associated utilities are considered in cases of ties. The outcomes whose probabilities are (at

Expected Utility Theory can be supplemented with dominance reasoning. However, this would result in a violation of the Continuity axiom of Expected Utility Theory. See Kosonen (2022, §2). Later in §4, I discuss versions of Probability Discounting that do not violate Statewise Dominance.

and) above the discounting threshold just take lexical priority over the very-smallprobability outcomes.<sup>28</sup>

STATEWISE DOMINANCE. However, Lexical Discounting violates Statewise Dominance as well. Consider the following case (table 2):

#### Lexical Statewise Dominance Violation:

*Prospect A* Gives \$10 in states 1 and 2, \$100 in state 3, and nothing in state 4.

*Prospect B* Gives \$10 in state 1, \$100 in states 2 and 3, and nothing in state 4.

The probability of states 1 and 4 is 0.49, and the probability of states 2 and 3 is 0.01. For simplicity, let the discounting threshold be (implausibly) 0.03. Let's also assume that the utility of money equals the monetary amount.

LEXICAL STATEWISE DOMINANCE VIOLATION					
	State 1	State 2	State 3	State 4	
p	0.49	0.01	0.01	0.49	
A	\$10	\$10	\$100	\$0	
B	\$10	\$100	\$100	\$0	

Table 2 Lexical Statewise Dominance Violation

After conditionalizing on not obtaining \$100 with A (as its associated probability is below the discounting threshold), A's probability-discounted expected utility is  $EU(A)_{pd} \approx 5.05$ .<sup>29</sup> And, after conditionalizing on not obtaining \$100 with B, B's probability-discounted expected utility is  $EU(B)_{pd} = 5$ .<sup>30</sup> Given that the former is greater than the latter, A is better than B according to Lexical Discounting.

$$^{29}0.5/(1-0.01)\cdot 10 \approx 5.05$$

 $^{30}0.49/(1-0.02) \cdot 10 = 5.$ 

<sup>&</sup>lt;sup>28</sup>It might be argued that because some small probabilities are much smaller than others, one should have multiple discounting thresholds that form probability ranges, where higher probability ranges take lexical priority over the lower ones.

However, the only difference between A and B is that A gives \$10 in state 2, while B gives \$100 in that same state. Therefore, Lexical Discounting—too—violates Statewise Dominance.

This violation of Statewise Dominance happens because when one conditionalizes on not obtaining \$100 with A (state 3), the probability of state 3 is divided between states 1, 2 and 4. However, when one conditionalizes on not obtaining \$100 with B (states 2 and 3), the probability of states 2 and 3 is divided between states 1 and 4. Therefore, after ignoring the possibility of obtaining \$100, the probability of obtaining nothing is greater with B than with A.

To summarize, Lexical Discounting states that outcomes whose probabilities are (at or) above the discounting threshold take lexical priority over very-smallprobability outcomes in determining prospects' betterness ranking—very-smallprobability outcomes are only treated as tiebreakers. However, like Naive Discounting, Lexical Discounting also violates Statewise Dominance.

## **3** State Discounting

This section discusses a version of Probability Discounting on which one should conditionalize on very-small-probability states not occurring. Three versions of this view are presented. I will show that they all violate dominance or acyclicity (or both).

#### 3.1 Pairwise and Set-Dependent State Discounting

Again, there is a straightforward solution to the previous violation of Statewise Dominance. Earlier it was assumed that one should ignore (except in cases of ties) the possibility of obtaining outcomes associated with tiny probabilities. However, one might ignore very-small-probability *states* instead—let's call this view *State Discounting*. One can also make a lexical version of this view:

**State Discounting** Prospect X is at least as preferred as prospect Y if and only if

- $\succ EU(X)_{pd} > EU(Y)_{pd}$  or
- $\succ \ EU(X)_{pd} = EU(Y)_{pd} \text{ and } EU(X) \geq EU(Y),$

where  $EU(X)_{pd}$  and  $EU(Y)_{pd}$  are obtained by conditionalizing on the supposition that no state of negligible probability occurs.

In the previous violation of Statewise Dominance, State Discounting tells one to ignore states 2 and 3 as their associated probabilities are below the discounting threshold. Consequently, A and B have equal probability-discounted expected utility (as they give the same outcomes in states 1 and 4). And B has greater expected utility without discounting, so it is better than A (assuming a lexical version of State Discounting). Thus, State Discounting avoids violating Statewise Dominance in the previous case.<sup>31</sup> However, next I will present some problems for State Discounting.

STATE INDIVIDUATION PROBLEM. State Discounting faces an analogous problem to the Outcome Individuation Problem, namely, the

**State Individuation Problem:** If one individuates states with too much detail, all states have negligible probabilities. Is there a privileged way of individuating states that avoids this?

As before, a possible solution is to individuate states by the utilities of their outcomes.<sup>32</sup>

However, there are different views about how exactly this should be done. On one version of State Discounting, prospects are always compared two at a time,

<sup>&</sup>lt;sup>31</sup>However, as I will show later, one version of State Discounting violates Statewise Dominance in this case.

<sup>&</sup>lt;sup>32</sup>As before, one problem with this is that, in some cases, all states might have probabilities below the discounting threshold. One could lower the threshold in such cases. However, this will not solve the problem in cases where all states have a zero probability of occurring.

and the possible states of the world are partitioned for every pairwise comparison separately. Alternatively, one could compare all available options at once and partition the states for every choice set separately. Let's call these views *Pairwise State Discounting* and *Set-Dependent State Discounting*, respectively (the following acyclicity violation illustrates the difference between these views).

**Pairwise State Discounting:** States are partitioned by comparing two prospects at a time.

**Set-Dependent State Discounting:** States are partitioned by comparing all available prospects at once.

ACYCLICITY. Although both versions of State Discounting avoid the earlier violations of Statewise Dominance (in the way explained before), they violate *Acyclicity* instead. Acyclicity states the following:

Acyclicity: If prospect  $X_1$  is preferred to prospect  $X_2$ , which is preferred to prospect  $X_3, \ldots$ , and prospect  $X_{n-1}$  is preferred to prospect  $X_n$ , then it is not the case that  $X_n$  is preferred to  $X_1$ .

According to Acyclicity, if there is a sequence of prospects such that each is preferred to the next, then it is not the case that the last prospect is preferred to the first.

To see why both versions of State Discounting violate Acyclicity, consider the following case:

**Acyclicity Violation:** A random number generator returns a number between 1 and 100.

*Prospect A* Gives \$1000 with numbers 1 and 2 (probability 0.02); otherwise it gives nothing.

*Prospect B* Certainly gives \$10 no matter what number comes up.

*Prospect* C Gives \$1000 with number 1 (probability 0.01) and otherwise it gives \$1.

Let the discounting threshold be 0.02. First, compare A and B. Individuating states by the utilities of their outcomes results in two states as shown in table 3. A is better than B because neither state has a non-negligible probability, and A's expected utility is greater than that of B.<sup>33</sup> Next, compare B and C. In this case, individuating states by the utilities of their outcomes results in states shown in table 4. As the probability of state 1\* is below the discounting threshold, one should ignore the possibility of state 1\* occurring. Once one does that, B is better than C, as it gives a better outcome in state 2\* (\$10 vs. \$1).

	TABLE 3	3		Тав	LE 4
	A is better than $B$		B is better than $C$		er than $C$
	State 1	State 2	State 1* State 2*		State 2*
Output	1 or 2 ( <i>p</i> =0.02)	3 to 100 ( <i>p</i> =0.98)	Output	1 (p=0.01)	2 to 100 ( <i>p</i> =0.99)
A	\$1000	\$0	В	\$10	\$10
В	\$10	\$10	C	\$1000	\$1

Now we have that A is better than B, which is better than C. It follows by Acyclicity that C is not better than A. However, when we compare A and C pairwise, we notice that C is better than A. In this case, individuating states by the utilities of their outcomes results in states shown in table 5. As states  $1^{**}$  and  $2^{**}$ have probabilities below the discounting threshold, the agent should ignore the possibilities of these states. Then, C is better than A because it gives a better outcome in state  $3^{**}$  (\$1 vs. \$0). So, we have a violation of Acyclicity: A is better than B, which is better than C, which is better than A.<sup>34</sup>

 ${}^{33}EU(A)_{pd} = 0.02 \cdot 1000 = 20$  and  $EU(B)_{pd} = 10$ .

<sup>&</sup>lt;sup>34</sup>Cibinel (forthcoming) shows that either probability discounters display cyclic preferences or they are committed to verdicts that are fanatical by their own lights.

U IS BEITER THAN A				
State 1** State 2** State 3**				
Output	1 ( <i>p</i> =0.01)	2 ( <i>p</i> =0.01)	3 to 100 ( <i>p</i> =0.98)	
A	\$1000	\$1000	\$0	
C	\$1000	\$1	\$1	

TABLE 5 C is detted than A

Let's now go back to Pairwise and Set-Dependent State Discounting. If we partition states for each pair of options in a way that depends on the particular two options being compared (in line with Pairwise State Discounting), then State Discounting violates Acyclicity within choice sets. Consequently, it is not clear what one ought to choose when all A, B and C are available, as there is no most-preferred alternative.35

However, if we partition states in a way that depends on the overall choice set (in line with Set-Dependent State Discounting), then there is no violation of Acyclicity within choice sets (see table 6). In this case, states 1\*\*\* and 2\*\*\* have probabilities below the discounting threshold, so one should ignore the possibilities of these states. Consequently, B is the best prospect as it gives the best outcome in state  $3^{***}$ , and C is the second-best prospect as it gives a better outcome than A in that state.

NO VIOLATION OF ACYCLICITY				
State 1*** State 2*** State 3***				
Output	1 ( <i>p</i> =0.01)	2 ( <i>p</i> =0.01)	3 to 100 ( <i>p</i> =0.98)	
A	\$1000	\$1000	\$0	
B	\$10	\$10	\$10	
C	\$1000	\$1	\$1	

TABLE 6

<sup>35</sup>Fishburn (1991, p. 116).

However, Set-Dependent State Discounting violates Acyclicity across choice sets (as shown in tables 3, 4 and 5). In particular, it was shown that Set-Dependent State Discounting violates Pairwise Acyclicity, that is, it violates Acyclicity when we compare two options at a time.

It is odd that adding or removing options can influence which events one ignores. For example, when comparing A and B, Set-Dependent State Discounting does not ignore the possibility of the random number generator outputting number 1 or 2. However, when C is also available, Set-Dependent State Discounting ignores these possibilities. Consequently, the value of A decreases significantly when C is also available, as one then ignores the possibility of obtaining \$1000 with A.<sup>36</sup> STOCHASTIC DOMINANCE. Finally, I will show that both versions of State Discounting violate the following dominance principle:<sup>37</sup>

**Stochastic Dominance:** If, for each outcome o, the total probability of getting an outcome at least as preferred as o is at least as high with prospect X as it is with prospect Y, then X is at least as preferred as Y (*Weak Stochastic Dominance*). If, in addition, for some outcome u, the total probability of getting an outcome at least as preferred as u is greater with X than with Y, then X is strictly preferred to Y (*Strong Stochastic Dominance*).

A violation of Stochastic Dominance happens if, for all outcomes, some prospect X gives an at least as high probability of an at least as great outcome as some other prospect Y does, and for some outcome, X gives a greater probability of an at least as great outcome as Y does—yet Y is judged better than or equally as good as X.

<sup>&</sup>lt;sup>36</sup>This case shows that Set-Dependent State Discounting violates *Contraction Consistency* and *Strong Expansion Consistency*. See Sen (1977, pp. 63–66). More generally, Contraction Consistency implies Acyclicity. See Sen (1977, p. 67).

<sup>&</sup>lt;sup>37</sup>More precisely, the definition given here is for *first-order stochastic dominance*, an idea that was introduced to statistics by Mann and Whitney (1947) and Lehmann (1955), and to economics by Quirk and Saposnik (1962). The name 'first-degree stochastic dominance' is due to Hadar and Russell (1969, p. 27).

To see why both versions of State Discounting violate Stochastic Dominance, consider the following case:

#### **Two Coins:**

*Prospect A* Gives \$10 if a coin lands on heads (p = 0.5), nothing if it lands on tails (p = 0.49), and \$100 if it lands on the edge (p = 0.01).

**Prospect** B Gives \$10 if another coin lands on heads (p = 0.49), nothing if it lands on tails (p = 0.49), and \$100 if it lands on the edge (p = 0.02).

Let the discounting threshold be 0.03. These prospects give the same probabilities of the same outcomes as the prospects in *Lexical Statewise Dominance Violation* (table 2). But instead of four states, we now have nine different states due to having two coins. Let 'H' stand for 'heads', 'T' for 'tails' and 'E' for 'edge'. Also, let '(X, Y)' stand for the first coin landing on 'X' and the second one on 'Y'. Only four states have probabilities above the discounting threshold: (H, H), (H, T), (T, H) and (T, T) (see table 7).

Two Coins				
v	H, H	H, T	T, H	T, T
	0.253	0.253	0.247	0.247
A B	\$10 \$10	\$10 \$0	\$0 \$10	\$0 \$0

TADLE 7

 $p^*$ =probability conditional on one of these states occurring.

After conditionalizing on one of these four states occurring, the probabilitydiscounted expected utility of A is greater than that of B. Now the only difference between these prospects is that A gives \$10 in state (H, T) and nothing in state (T, H), while B gives \$10 in the latter state and nothing in the former one—and the former state has a greater probability.<sup>38</sup> Thus, A is better than B according

 $<sup>{}^{38}</sup>EU(A)_{pd} \approx 5.05 \text{ and } EU(B)_{pd} = 5.$ 

to both versions of State Discounting. However, this is a violation of Stochastic Dominance. Before discounting, both A and B give a 0.51 probability of at least \$10, but B gives a greater probability of at least \$100 (0.02 vs. 0.01). So, for all outcomes, B gives an at least as high probability of an at least as great outcome as A does, and for some outcome, B gives a greater probability of an at least as great as great outcome as A does. Thus, B stochastically dominates A—and both versions of State Discounting violate Stochastic Dominance.

#### 3.2 **Baseline State Discounting**

According to the previous versions of State Discounting, states might be partitioned differently depending on what other options are available. This leads to violations of Acyclicity. However, states might also be partitioned in a way that does not depend on the other available options. This can be done by comparing each prospect to some baseline or status quo prospect—let's call this *Baseline State Discounting*:

**Baseline State Discounting:** States are partitioned by comparing every prospect to a status quo prospect (each separately).<sup>39</sup>

STATEWISE DOMINANCE. However, Baseline State Discounting violates Statewise Dominance in the same way as Lexical Discounting does. Consider again *Lexical Statewise Dominance Violation* (table 2). This time, let's specify the events that result in each outcome:

**Random Number:** A random number generator returns a number between 1 and 100.

*Prospect* A Gives \$10 with numbers 1 to 50, \$100 with number 51 and nothing with numbers 52 to 100.

<sup>&</sup>lt;sup>39</sup>Note that, on Baseline State Discounting, one might sometimes ignore some events  $e_1$  and  $e_2$  when comparing some prospect X to the status quo prospect, but not ignore them when comparing another prospect Y to the status quo prospect.

*Prospect* B Gives \$10 with numbers 1 to 49, \$100 with numbers 50 and 51 and nothing with numbers 52 to 100.

In this case, the baseline prospect is (presumably) certainly getting nothing. When A is compared to this baseline prospect, state individuation by utilities results in three states as shown in table 8. As the probability of state 2 is below the discounting threshold of 0.03, the possibility of this state is ignored. Consequently, as before, the probability-discounted expected utility of A is  $EU(A)_{pd} \approx 5.05$ .<sup>40</sup>

A vs. the Baseline			
Output	<b>State 1</b>	<b>State 2</b>	<b>State 3</b>
	1–50 ( <i>p</i> =0.5)	51 ( <i>p</i> =0.01)	52–100 ( <i>p</i> =0.49)
<i>A</i>	\$10	\$100	\$0
Baseline	\$0	\$0	\$0

Next, compare *B* to the baseline prospect. This time state individuation by utilities results in the three states shown in table 9. Again, the probability of state  $2^*$  is below the discounting threshold, so the possibility of this state is ignored. Then, the probability-discounted expected utility of *B* is  $EU(B)_{pd} = 5.^{41}$ 

TABLE 9

B vs. the Baseline				
State 1* State 2* State 3*				
Output	1–49 ( <i>p</i> =0.49)	50 or 51 (p=0.02)	52–100 ( <i>p</i> =0.49)	
В	\$10	\$100	\$0	
Baseline	\$0	\$0	\$0	

 ${}^{40}EU(A)_{pd} = 0.5/0.99 \cdot 10 \approx 5.05.$ 

 $^{41}EU(B)_{pd}=0.49/0.98\cdot 10=5.$ 

As A's probability-discounted expected utility is greater than that of B (5.05 vs. 5), A is better than B. However, B statewise dominates A because the only difference between these prospects is that A gives \$10 if the random number generator returns number 50, while B gives \$100 in that case. Consequently, Baseline State Discounting violates Statewise Dominance when states are partitioned in the usual way corresponding to possible states of the world (such as 'number 50 is returned'). This violation of Statewise Dominance happens because the possible states of the world that Baseline State Discounting ignores are not the same for every prospect. For example, when comparing A to the baseline prospect, the possibility of the random number generator returning number 50 is not ignored, but this possibility is ignored when B is compared to the baseline prospect.

To summarize, instead of ignoring very-small-probability outcomes, Probability Discounting might ignore very-small-probability states. State Discounting faces the State Individuation Problem, which can be solved by individuating states by the utilities of their outcomes. I have discussed three ways of formulating State Discounting. Pairwise State Discounting always compares two options at a time and ignores very-small-probability states in every pairwise comparison. However, Pairwise State Discounting violates Stochastic Dominance and Acyclicity within choice sets. Set-Dependent State Discounting compares all available options simultaneously and ignores very-small-probability states in every choice set. This view violates Pairwise Acyclicity and Stochastic Dominance. Finally, Baseline State Discounting ignores very-small-probability states of some prospect when states are partitioned by comparing this prospect to a baseline prospect. This view violates Statewise (and hence also Stochastic) Dominance. To conclude, all three versions of State Discounting violate plausible principles of rationality.<sup>42</sup>

<sup>&</sup>lt;sup>42</sup>Someone might adopt a view on which one should first filter one's options by Statewise and Stochastic Dominance and then choose following some version of Probability Discounting from amongst the remaining options. This view avoids the dominance violations, but it also seems *ad hoc*. However, some may find the benefit of a greater fit with our intuitions worth the cost in terms of simplicity.

### 4 Stochastic and Tail Discounting

This section discusses more plausible versions of Probability Discounting that avoid the earlier violations of dominance (and Acyclicity).

### 4.1 Stochastic Discounting

One version of Probability Discounting—let's call it *Absolutist Stochastic Discount-ing*—works like this: To obtain the probability-discounted expected utility of a prospect, first add the lowest possible (positive) utility, weighted by the probability of getting at least that much utility. Next, add the difference between the lowest utility and the next lowest utility, weighted by the probability of getting at least the higher amount of utility. Then, add the difference between this utility and the next lowest utility of getting at least that much utility of getting at least that much utility. Then, add the difference between this utility and the next lowest utility, weighted by the probability of getting at least that much utility, and so on until the next probability is below the discounting threshold.<sup>43</sup> Then, ignore the rest of the utility levels (whose probabilities are below the discounting threshold). Negative utilities are then treated similarly, and their expectation is summed with the expectation of positive utilities to obtain the value of a prospect.

The above approach is mathematically equivalent to calculating the probabilitydiscounted expected utility of a prospect in the same way as Expected Utility Theory would do with the following exception: The greatest positive and lowest negative utilities (whose utility levels have negligible cumulative probability) have been replaced with the greatest positive or lowest negative utility whose utility level has a non-negligible cumulative probability (respectively for positive and negative utilities).

According to Absolutist Stochastic Discounting, there is an objective neutral level. On this view, one should ignore the possibility of very high or very low utility levels when the probability of ending up with at least or at most that much utility is negligible (respectively for positive and negative utilities). This view recommends

<sup>&</sup>lt;sup>43</sup>This is similar to an alternative way of calculating the expected utility of a prospect discussed by Buchak (2014, p. 1100). See also Goodsell (forthcoming).

against accepting the offer in Pascal's Hell if one has only a tiny probability of obtaining an outcome at least as good as a million Graham's number of happy Earthlike planets. However, it does not recommend against accepting the offer if there is a non-negligible probability of obtaining an outcome at least that good for some reason unrelated to Satan's offer.<sup>44</sup>

Consider a prospect that has possible outcomes whose values are normally distributed with a mean of zero when the outcomes are ordered in terms of betterness. Absolutist Stochastic Discounting tells one to substitute the values of the outcomes in the grey areas (see the graph below) with the values of u and -u (respectively for positive and negative outcomes), where u and -u are the best positive and the worst negative utility levels that have non-negligible cumulative probabilities. So, utilities greater than u are replaced by u in the utility calculations; utilities less than -u are replaced by -u in the utility calculations.



The following prospects Y and Z can only result in positive or negative outcomes, respectively. Consequently, Absolutist Stochastic Discounting tells one to substitute the values of the best positive outcomes of Y (the grey area in the top image below) with u and the values of the worst negative outcomes of Z (the grey area in the bottom image below) with -u where u and -u are the best positive and the worst negative utility levels of Y and Z (respectively) that have non-negligible cumulative probabilities.

<sup>&</sup>lt;sup>44</sup>For example, an agent who thinks there is a non-negligible probability of going to Heaven would not ignore the possibility of a great payoff in Pascal's Hell. More generally, such an agent would not discount small probabilities very often (if ever); the non-negligible probability of going to Heaven makes it the case that there is a non-negligible probability of ending up with at least uamount of utility for all positive values of u.



Call the versions of Probability Discounting that have the same structure as Absolutist Stochastic Discounting *Stochastic Discounting*. There is another way of understanding Stochastic Discounting. This view is similar to Baseline State Discounting, as it compares each prospect to a baseline prospect—let's call it *Baseline Stochastic Discounting*. On this view, one calculates the amount by which the baseline/status quo utility level is increased or decreased by the different possible outcomes of a prospect. Then, to obtain the probability-discounted expected utility of a prospect, one first adds the lowest possible gain (i.e., positive change to the baseline), weighted by the probability of gaining at least that much. Next, one adds the difference between the lowest gain and the next lowest gain, weighted by the probability is below the discounting threshold. Then, one ignores the rest of the possible gains. Losses (i.e., negative changes to the baseline) are treated similarly, and their expectation is summed with the expectation of gains to obtain the value of a prospect.<sup>45</sup>

Similarly as before, this approach is mathematically equivalent to calculating the probability-discounted expected utility of a prospect by substituting the values

<sup>&</sup>lt;sup>45</sup> One can also make a version of Stochastic Discounting that is analogous to Pairwise State Discounting in that it compares prospects to other available prospects pairwise—call this *Pairwise Stochastic Discounting*. On this view, one considers the utility difference in each state between two prospects and ignores the largest differences when the cumulative probability of states with differences at least that large is negligible.

of the greatest gains and losses with the values of g and l (respectively for gains and losses), where g and l are the greatest gains and losses that have non-negligible cumulative probabilities.

Absolutist Stochastic Discounting has the (possible) disadvantage of requiring an objective neutral utility level. Baseline Stochastic Discounting does not require one because it ignores very large changes to the baseline—the baseline serves the same purpose as the objective neutral level on the absolutist view.

Also, Absolutist Stochastic Discounting implies that sometimes one might not ignore the possibility of a huge gain (or loss) even if there is only a tiny probability of it occurring. This can happen if one would end up with a negative outcome regardless of the gain, and the probability of obtaining an outcome that is at most as good as that is non-negligible. Baseline Stochastic Discounting, in contrast, ignores a tiny probability of a great gain even in that case.

Similarly, on the absolutist view one should not ignore a tiny probability of a huge gain if the probability of ending up with a higher utility level is non-negligible for some reason unrelated to the prospect in question. As mentioned before, this view does not recommend against accepting the offer in Pascal's Hell if there is a non-negligible probability of gaining a greater payoff for some reason unrelated to the offer. Baseline Stochastic Discounting, in contrast, recommends against accepting the offer even in that case. This is so because once one has 'subtracted' the baseline prospect from the offer, there is only a tiny probability of obtaining an outcome at least as good as a million Graham's number of happy Earth-like planets. So, unlike Baseline Stochastic Discounting, Absolutist Stochastic Discounting sometimes lets tiny probabilities of huge gains or losses dictate one's course of action. Therefore, it does not capture the motivation behind Probability Discounting as well as Baseline Stochastic Discounting does.

On both versions of Stochastic Discounting, the probability-discounted expected utility of prospect X is obtained by summing the probability-discounted expected utilities of its positive and negative outcomes (here 'outcomes' are either final utilities if one accepts Absolutist Stochastic Discounting or gains/losses if one

accepts Baseline Stochastic Discounting). Let these be denoted by  $EU(X)_{pd, pos}$ and  $EU(X)_{pd, neg}$  (respectively). As before, Stochastic Discounting can use verysmall-probability outcomes as tiebreakers to rank prospects with equal probabilitydiscounted expected utility. Stochastic Discounting can then be stated as follows:

**Stochastic Discounting:** Prospect X is at least as preferred as prospect Y if and only if

> 
$$EU(X)_{pd}$$
 >  $EU(Y)_{pd}$  or
 >  $EU(X)_{pd} = EU(Y)_{pd}$  and  $EU(X) ≥ EU(Y)$ ,

where, for all prospects X, it holds that

$$EU(X)_{pd} = EU(X)_{pd, pos} + EU(X)_{pd, neg}.$$
<sup>46</sup>

Now, recall the earlier violations of Statewise and Stochastic Dominance (*Lexical Statewise Dominance Violation, Random Number* and *Two Coins*). Unlike the earlier versions of Probability Discounting, both versions of Stochastic Discounting imply that *B* is better than *A*. Both *A* and *B* give a 0.51 probability of at least \$10. In addition, *A* gives a 0.01 probability of at least \$100, while *B* gives a 0.02 probability of at least \$100. After ignoring the possibility of obtaining \$100, the probability-discounted expected utility of *A* and *B* is  $EU(A)_{pd} = EU(B)_{pd} = 5.1.^{47}$  As *A* and *B* have equal probability-discounted expected utility, these prospects are then compared by their expected utilities without discounting. Consequently, *B* is bet-

$$EU(X)_{pd,\ pos} = \sum_{i=1}^n p(E_i)u(x_i) + \Big(\sum_{i=n+1}^\infty p(E_i)\Big)u(x_n),$$

where  $x_n$  is the greatest positive utility that has a non-negligible cumulative probability. The probability-discounted expected utility of negative outcomes (i.e.  $EU(X)_{pd, neg}$ ) can be calculated in a similar way as that of positive outcomes (changing what needs to be changed).

$${}^{47}EU(A)_{pd} = EU(B)_{pd} = 0.51 \cdot 10 = 5.1.$$

<sup>&</sup>lt;sup>46</sup>The probability-discounted expected utility of positive outcomes can be calculated as follows:

ter than A, and both versions of Stochastic Discounting avoid the earlier violations of Statewise and Stochastic Dominance.<sup>48</sup>

#### 4.2 Tail Discounting

There is a similar view to Absolutist Stochastic Discounting called *Tail Discounting*.<sup>49</sup> According to Tail Discounting, one should ignore both the left and the right 'tails' of the distribution of possible outcomes of some prospect X when these outcomes are ordered by one's preference. Suppose the possible outcomes of some prospect are normally distributed when they are ordered from the least to the most preferred. Then, Tail Discounting advises one to ignore the grey areas under the curve:



Call the outcomes that fall in the middle of the distribution of possible outcomes 'normal outcomes'. An outcome is normal if and only if there is a nonnegligible probability of getting at least and at most as good an outcome. Tail Discounting then states the following:

**Tail Discounting:** Prospect X is at least as preferred as prospect Y if and only if

 $\succ EU(X)_{pd} > EU(Y)_{pd}$  or

<sup>&</sup>lt;sup>48</sup>Wilkinson (2022, §6) shows that views that reject fanaticism must violate separability or Stochastic Dominance. Absolutist Stochastic Discounting violates the former. Given that Baseline Stochastic Discounting ignores background uncertainty (and thus satisfies separability), it must violate Stochastic Dominance.

<sup>&</sup>lt;sup>49</sup>Tail Discounting is from Beckstead and Thomas (forthcoming, §2.3).

 $\succ EU(X)_{pd} = EU(Y)_{pd}$  and  $EU(X) \ge EU(Y)$ ,

where  $EU(X)_{pd}$  and  $EU(Y)_{pd}$  are obtained by conditionalizing on the supposition that some normal outcome occurs.<sup>50</sup>

Tail Discounting has the advantage over Absolutist Stochastic Discounting that it does not require an objective neutral level. However, similarly as Absolutist Stochastic Discounting, Tail Discounting recommends accepting the offer in Pascal's Hell if, for some reason unrelated to the offer, there is a non-negligible probability of obtaining a greater outcome.<sup>51</sup> This is because then a million Graham's number of happy Earth-like planets falls in the middle part of the distribution of possible outcomes, which is not ignored.<sup>52</sup>

- i The probability of obtaining an outcome that is at least as preferred as *o* is above the discounting threshold and
- ii the probability of obtaining an outcome that is at most as good as *o* is above the discounting threshold.

If some outcome o fulfills the above necessary conditions, and

- the probability of obtaining an outcome that is better than *o* is below the discounting threshold, then decrease the probability of obtaining *o* until the total discounted probability of outcomes that are at least as good as *o* equals the discounting threshold (and conditionalize to make sure the remaining probabilities add up to 1), and
- if the probability of obtaining an outcome that is worse than *o* is below the discounting threshold, then decrease the probability of obtaining *o* until the total discounted probability of outcomes that are at most as good as *o* equals the discounting threshold (and conditionalize to make sure the remaining probabilities add up to 1).

<sup>51</sup>Cibinel (forthcoming) independently shows that Tail Discounting sometimes does not ignore tiny probabilities of extreme values.

<sup>52</sup>One can also make a version of Tail Discounting similar to Baseline Stochastic Discounting. On this view—let's call it *Baseline Tail Discounting*—one compares every prospect to a baseline prospect as follows: First, calculate the difference in utilities a prospect makes in each state (com-

<sup>&</sup>lt;sup>50</sup>Formally this view states the following:

**Tail Discounting (formal):** In order to determine  $EU(X)_{pd}$ , first order the possible outcomes of some prospect X from the least to the most preferred. Then, conditionalize on obtaining some outcome in the middle part of the distribution such that the following necessary conditions hold for all outcomes *o* that are not ignored:

Stochastic Discounting and Tail Discounting deal with tiny probabilities differently. Stochastic Discounting, in effect, replaces the highest positive and lowest negative utility levels that are associated with tiny probabilities with the highest positive and lowest negative utility levels that are non-negligible—let's call this the *Replacing Method*. In contrast, Tail Discounting conditionalizes on the extreme outcomes not happening—let's call this the *Conditionalization Method*. However, these methods can be switched. Tail Discounting could also use the Replacing Method, that is, replace the extreme outcomes with the highest and lowest utility levels that are non-negligible. Similarly, Stochastic Discounting could use the Conditionalization Method, that is, conditionalize on the highest positive and lowest negative utility levels not occurring. The difference between Tail Discounting and Stochastic Discounting then boils down to whether one wishes to only ignore very high positive and very low negative utility levels (Stochastic Discounting) or also the highest and lowest utility levels of a prospect, even when these are not far from the neutral utility level (if there is a neutral level).

Recall the earlier violations of Statewise and Stochastic Dominance (*Lexical Statewise Dominance Violation, Random Number* and *Two Coins*). Tail Discounting also implies that *B* is better than *A*. Again, *A* and *B* have equal probabilitydiscounted expected utility:  $EU(A)_{pd} = EU(B)_{pd} \approx 5.1.^{53}$  These prospects are then compared by their expected utilities without discounting. Thus, *B* is bet-

pared to the baseline prospect). Then, order these differences from the greatest loss to the greatest gain. Next, ignore the right and left tails of this distribution by conditionalization. Baseline Tail Discounting differs from Baseline Stochastic Discounting because the former uses the Conditionalization Method to ignore huge gains and losses while the latter uses the Replacing Method. Also, one can make a version of Tail Discounting similar to Pairwise Stochastic Discounting (i.e., *Pairwise Tail Discounting*). On this view, one compares prospects pairwise instead of comparing every prospect to a baseline prospect. Similarly, Pairwise Tail Discounting differs from Pairwise Stochastic Discounting because the former uses the Conditionalization Method while the latter uses the Replacing Method.

 $<sup>{}^{53}</sup>EU(A)_{pd} = (0.5-0.02)/0.94 \cdot 10 \approx 5.1$  and  $EU(B)_{pd} = (0.49-0.01)/0.94 \cdot 10 \approx 5.1$ . The divisor '0.94' comes from subtracting the discounting threshold of 0.03 from both tails of the distribution. '0.02' and '0.01' are subtracted to make sure that the full discounting threshold of 0.03 is ignored in the right tail. In general, on Tail Discounting, one discounts a little bit of each 'tail' with every prospect (until the discounting threshold is ignored from both tails).

ter than A—and Tail Discounting avoids violating Statewise and Stochastic Dominance in the earlier cases.

To summarize, I have discussed three versions of Probability Discounting in this section. Absolutist Stochastic Discounting states that one should ignore the possibility of a very high (very low) utility level in cases where the probability of obtaining at least (at most) that much utility is below the discounting threshold. Baseline Stochastic Discounting works similarly, but it operates on gains and losses instead of final utilities. Lastly, Tail Discounting states that one should ignore the 'tails' of the distribution of possible outcomes of some prospect. All these views avoid the earlier violations of Statewise and Stochastic Dominance (and Acyclicity).

### 4.3 Violating the axioms of Expected Utility Theory

As Probability Discounting differs from Expected Utility Theory, it must violate at least one of the axioms that together entail Expected Utility Theory: Completeness, Transitivity, Independence and Continuity.<sup>54</sup> In fact, all versions of Probability Discounting violate both Independence and Continuity.<sup>55</sup> As a result of these violations, these views are vulnerable to exploitation by a money pump.<sup>56</sup> A money-pump argument intends to show that agents who violate some alleged requirement of rationality are vulnerable to making a combination of choices that leads to a sure loss. If vulnerability to this kind of exploitation is a sign of irrationality, then Stochastic and Tail Discounting are untenable as theories of instrumental rationality.<sup>57</sup>

<sup>&</sup>lt;sup>54</sup>Von Neumann and Morgenstern (1947), Jensen (1967, pp. 172–182) and Hammond (1998, pp. 152–164).

<sup>&</sup>lt;sup>55</sup>See Kosonen (2022, §5). See also Russell (forthcoming, p. 18n24).

<sup>&</sup>lt;sup>56</sup>For the Independence Money Pump, see Hammond (1988*b*, pp. 292–293), Hammond (1988*a*, pp. 43–45), Gustafsson (2021, p. 31 n21) and Gustafsson (2022, §5). For the Continuity Money Pump, see Gustafsson (2022, §6).

<sup>&</sup>lt;sup>57</sup>For possible ways of avoiding exploitation in these money pumps, see Kosonen (2022, §5). However, note that, independently of Probability Discounting, agents with unbounded utilities are also vulnerable to money pumps because they violate countable generalizations of the Independence axiom. See Russell and Isaacs (2021).

## 5 Conclusion

Expected utility maximization with unbounded utilities implies counterintuitive recommendations in cases that involve tiny probabilities of huge payoffs. In response to such cases, some have argued that we should discount small probabilities down to zero. I have discussed how exactly this view can be formulated. First, I showed that less plausible versions of Probability Discounting violate dominance. More specifically, I showed that Naive Discounting, Lexical Discounting and Baseline State Discounting violate Statewise Dominance. I also showed that Pairwise State Discounting violates Stochastic Dominance and Acyclicity within choice sets and that Set-Dependent State Discounting violates Pairwise Acyclicity and Stochastic Dominance.

Then, I showed that more plausible versions of Probability Discounting, namely Stochastic Discounting and Tail Discounting, avoid these dominance violations. However, they must—of course—violate at least one of the axioms of Expected Utility Theory. As a result of this violation, those who accept these views can be exploited by a money pump.

What should one do now? One could, for example, bite the bullet and accept a version of Probability Discounting discussed in this paper (see table 10 below), find a more plausible version of Probability Discounting, bound utilities<sup>58</sup>, conditionalize on one's knowledge before maximizing expected utility<sup>59</sup> or accept Probability Fanaticism.<sup>60</sup>

<sup>&</sup>lt;sup>58</sup>See Beckstead and Thomas (forthcoming, §2.1) and Kosonen (2022, pp. 32–38, 60–111).

<sup>&</sup>lt;sup>59</sup>See, for example, Francis and Kosonen (n.d.).

<sup>&</sup>lt;sup>60</sup>See Beckstead and Thomas (forthcoming) and Wilkinson (2022).

	Outcome/State Discounting	Stochastic Discounting	Tail Discounting			
Absolutist	Absolutist Discounting (e.g. Naive/Lexical Discounting): Ignores outcomes of tiny probabilities. (Violates Statewise Dominance. See §1 and §2.)	Absolutist Stochastic Discounting: Ignores highest positive and lowest negative utility levels. Neutral utility level needed. Uses Replacing Method. (See §4.1.)	Absolutist Tail Discounting: Ignores extreme utilities. No neutral level needed. Uses Conditionalization Method. (See §4.2.)			
Baseline	Baseline State Discounting: Ignores tiny-probability states when comparing prospects to a baseline. (Violates Statewise Dominance. See §3.2.)	Baseline Stochastic Discounting: Ignores greatest gains and losses when comparing prospects to baseline. Uses Replacing Method. (See §4.1.)	Baseline Tail Discounting: Ignores greatest gains and losses when comparing prospects to baseline. Uses Conditionalization Method. (See footnote 52.)			
Pairwise	Pairwise State Discounting: Ignores tiny-probability states in pairwise comparisons of prospects. (Violates Stochastic Dominance and Acyclicity within choice sets. See §3.1.)	Pairwise Stochastic Discounting: Ignores greatest differences in utilities in pairwise comparisons of prospects. Uses Replacing Method. (See footnote 45.)	Pairwise Tail Discounting: Ignores greatest differences in utilities in pairwise comparisons of prospects. Uses Conditionalization Method. (See footnote 52.)			
Set-Dependent	Set-Dependent State Discounting: Ignores tiny-probability states in every choice set. (Violates Stochastic Dominance, Pairwise Acyclicity and Contraction and Expansion Consistency. See §3.1.)	Set-Dependent Stochastic Discounting: Ignores greatest differences in utilities in every choice set. Uses Replacing Method.	Set-Dependent Tail Discounting: Ignores greatest differences in utilities in every choice set. Uses Conditionalization Method.			

### Table 10 Some Possible Versions of Probability Discounting

# References

- Baumann, P. (2009), 'Counting on numbers', Analysis 69(3), 446-448.
- Beckstead, N. (2013), On the overwhelming importance of shaping the far future, PhD thesis, Rutgers, the State University of New Jersey.
- Beckstead, N. and Thomas, T. (forthcoming), 'A paradox for tiny probabilities and enormous values', *Nous*.
- Bernoulli, D. (1954), 'Exposition of a new theory on the measurement of risk', *Econometrica* **22**(1), 23–36.
- Bostrom, N. (2009), 'Pascal's Mugging', Analysis 69(3), 443-445.
- Broome, J. (1991), *Weighing Goods: Equality, Uncertainty and Time*, Blackwell, Oxford.
- Buchak, L. (2014), 'Risk and tradeoffs', Erkenntnis 79(6), 1091-1117.
- Cibinel, P. (forthcoming), 'A dilemma for Nicolausian discounting', Analysis .
- Easwaran, K. (2014), 'Regularity and hyperreal credences', *Philosophical Review* **123**(1), 1–41.
- Fishburn, P. C. (1991), 'Nontransitive preferences in decision theory', *Journal of Risk and Uncertainty* 4(2), 113–134.
- Francis, T. and Kosonen, P. (n.d.), Ignore outlandish possibilities. Unpublished manuscript.
- Goodsell, Z. (2021), 'A St Petersburg paradox for risky welfare aggregation', *Analysis* **81**(3), 420–426.

Goodsell, Z. (forthcoming), 'Decision theory unbound', Nous.

- Gustafsson, J. E. (2021), 'The sequential dominance argument for the Independence axiom of Expected Utility Theory', *Philosophy & Phenomenological Research* **103**(1), 21–39.
- Gustafsson, J. E. (2022), *Money-Pump Arguments*, Cambridge University Press, Cambridge.
- Hadar, J. and Russell, W. R. (1969), 'Rules for ordering uncertain prospects', *The American Economic Review* **59**(1), 25–34.
- Hájek, A. (2014), 'Unexpected expectations', Mind 123(490), 533-567.
- Hammond, P. J. (1988*a*), 'Consequentialist foundations for expected utility', *Theory and Decision* **25**(1), 25–78.
- Hammond, P. J. (1988*b*), 'Orderly decision theory: A comment on professor Seidenfeld', *Economics and Philosophy* **4**(2), 292–297.
- Hammond, P. J. (1998), Objective expected utility: A consequentialist perspective, *in* S. Barberà, P. J. Hammond and C. Seidl, eds, 'Handbook of Utility Theory Volume 1: Principles', Kluwer, Dordrecht, pp. 143–211.
- Hanson, R. (2007), 'Pascal's Mugging: Tiny probabilities of vast utilities'.
  URL: https://www.lesswrong.com/posts/a5JAiTdytou3Jg749/pascal-s-muggingtiny-probabilities-of-vast-utilities?commentId=Q4ACkdYFEThA6EE9P
- Hey, J. D., Neugebauer, T. M. and Pasca, C. M. (2010), Georges-Louis Leclerc de Buffon's 'Essays on moral arithmetic', *in* A. Sadrieh and A. Ockenfels, eds, 'The Selten School of Behavioral Economics: A Collection of Essays in Honor of Reinhard Selten', Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 245–282.

Huemer, M. (2016), Approaching Infinity, Palgrave Macmillan, New York.

- Isaacs, Y. (2016), 'Probabilities cannot be rationally neglected', *Mind* **125**(499), 759–762.
- Jensen, N. E. (1967), 'An introduction to Bernoullian utility theory: I. Utility functions', *The Swedish Journal of Economics* **69**(3), 163–183.
- Kosonen, P. (2022), Tiny Probabilities of Vast Value, PhD thesis, University of Oxford.
- Kreps, D. M. (1988), Notes on the Theory of Choice, Westview Press, Boulder.
- Lehmann, E. L. (1955), 'Ordered families of distributions', *The Annals of Mathematical Statistics* **26**(3), 399–419.
- Luce, R. D. and Raiffa, H. (1957), *Games and Decisions: Introduction and Critical Survey*, Wiley, New York.
- Lundgren, B. and Stefánsson, H. O. (2020), 'Against the *De Minimis* principle', *Risk Analysis* **40**(5), 908–914.
- Mann, H. B. and Whitney, D. R. (1947), 'On a test of whether one of two random variables is stochastically larger than the other', *The Annals of Mathematical Statistics* **18**(1), 50–60.
- Menger, K. (1967), The role of uncertainty in economics, *in* M. Shubik, ed., 'Essays in Mathematical Economics: In Honor of Oskar Morgenstern', Princeton University Press, Princeton, pp. 211–231.
- Monton, B. (2019), 'How to avoid maximizing expected utility', *Philosophers' Imprint* **19**(18), 1–24.
- Nover, H. and Hájek, A. (2004), 'Vexing expectations', *Mind* 113(450), 237–249.
- Pulskamp, R. J. (n.d.), Correspondence of Nicolas Bernoulli concerning the St. Petersburg game. Unpublished manuscript. Accessed through:

https://web.archive.org/.

**URL:** *http://cerebro.xu.edu/math/Sources/NBernoulli/correspondence\_peters-burg\_game.pdf* 

- Quirk, J. P. and Saposnik, R. (1962), 'Admissibility and measurable utility functions', *The Review of Economic Studies* **29**(2), 140–146.
- Russell, J. S. (forthcoming), 'On two arguments for fanaticism', Nous .
- Russell, J. S. and Isaacs, Y. (2021), 'Infinite prospects', *Philosophy and Phenomeno-logical Research* **103**(1), 178–198.
- Samuelson, P. A. (1977), 'St. Petersburg paradoxes: Defanged, dissected, and historically described', *Journal of Economic Literature* **15**(1), 24–55.
- Savage, L. J. (1951), 'The theory of statistical decision', *Journal of the American Statistical Association* **46**(253), 55–67.
- Sen, A. (1977), 'Social choice theory: A re-examination', *Econometrica* **45**(1), 53–88.
- Smith, N. J. J. (2014), 'Is evaluative compositionality a requirement of rationality?', *Mind* **123**(490), 457–502.
- Smith, N. J. J. (2016), 'Infinite decisions and rationally negligible probabilities', Mind 125(500), 1199–1212.
- von Neumann, J. and Morgenstern, O. (1947), *Theory of Games and Economic Behavior*, 2 edn, Princeton University Press, Princeton.
- Wilkinson, H. (2022), 'In defence of fanaticism', Ethics 132(2), 445–477.
- Yudkowsky, E. (2007*a*), 'A comment on Pascal's Mugging: Tiny probabilities of vast utilities'.

**URL:** *https://www.lesswrong.com/posts/a5JAiTdytou3Jg749/pascal-s-mugging-tiny-probabilities-of-vast-utilities?commentId=kqAKXskjohx4SSyp4* 

Yudkowsky, E. (2007*b*), 'Pascal's Mugging: Tiny probabilities of vast utilities'. **URL:** *http://www.overcomingbias.com/2007/10/pascals-mugging.html*