

CHAPTER 3

*Probability Discounting Solves the Intrapersonal Addition Paradox**

ABSTRACT: Nebel (2019) argues for the Repugnant Conclusion via the “Intrapersonal Repugnant Conclusion”, on which certainty of a mediocre life is better for individuals than a sufficiently small chance of an excellent life. In this chapter, I deny that accepting the Intrapersonal Repugnant Conclusion leads us to the Repugnant Conclusion. I point out that on many views which avoid the Repugnant Conclusion, we should discount small probabilities down to zero to avoid an implausibly reckless decision theory. But if we do, then Nebel’s crucial premise of Ex Ante Pareto fails because discounting at the individual level can fail to match up with discounting at the population level.

*I wish to thank Tomi Francis, Johan Gustafsson, Andreas Mogensen, Teruji Thomas, two anonymous reviewers of *Ethics* and the audience of the Slippery Slope Normativity Summit 2020 for valuable feedback and discussions. A version of this chapter was published in *Ethics*. See Kosonen (2021).

The Repugnant Conclusion, introduced by Parfit, states:¹

“For any possible population of at least ten billion people, all with a very high quality of life, there must be some much larger imaginable population whose existence, if other things are equal, would be better even though its members have lives that are barely worth living.”

More generally, the Repugnant Conclusion is that for any population wherein each individual has very high positive welfare, there is some much larger population wherein each individual has very low positive welfare, which is better.² The Repugnant Conclusion is a consequence of Total Utilitarianism, which states that one population is better than another just in case the total quantity of welfare it contains is greater. The Repugnant Conclusion strikes many as an unacceptable consequence, and various attempts at constructing an alternative population axiology to Total Utilitarianism have been made.³ However, a series of impossibility theorems have shown that no axiology can satisfy simultaneously all intuitively compelling principles that have been identified.⁴

In *An Intrapersonal Addition Paradox*, Nebel argues for the Repugnant Conclusion via an intrapersonal analogue of it.⁵ In this chapter, I deny that accepting the Intrapersonal Repugnant Conclusion leads us to the Repugnant Conclusion. My argument is that on many views which avoid the Repugnant Conclusion, we should discount small probabilities down to zero to avoid an implausibly reckless decision theory. But if we do, then Nebel’s

¹Parfit (1984, p. 388).

²Arrhenius (2000, p. 248).

³For an overview, see Greaves (2017).

⁴Arrhenius (2000).

⁵Nebel (2019, pp. 309–343).

crucial premise of Ex Ante Pareto fails because discounting at the individual level can fail to match up with discounting at the population level. The structure of this chapter is as follows: §1 presents Nebel’s argument. §2 explores the idea that we should discount small probabilities down to zero. §3 responds to Nebel’s argument. §4 argues that those who discount small probabilities must reject Ex Ante Pareto. §5 concludes.

1 Nebel’s argument

This section presents Nebel’s argument. His argument proceeds in two stages. The first stage is an argument for the Intrapersonal Repugnant Conclusion, and the second stage is an argument from the Intrapersonal Repugnant Conclusion to the Repugnant Conclusion. I will begin by discussing the first stage of the argument.

1.1 The Intrapersonal Repugnant Conclusion

Nebel’s *Intrapersonal Repugnant Conclusion* states:⁶

Intrapersonal Repugnant Conclusion: For any person S , there exists some probability p such that any prospect in which S would have a wonderful life with probability p or less, and would otherwise never exist, is worse for S than certainty of a life that is barely worth living.⁷

⁶Nebel (2019, p. 314). There is another—distinct—claim that could also be naturally described as the Intrapersonal Repugnant Conclusion: For any life lived by S at a very high welfare level for n years, there is some much longer life that would be better for S in which her welfare is barely above the zero level at each point in time. Temkin (2012, p. 119) calls this the Single Life Repugnant Conclusion. See also McTaggart (1927, pp. 452-453).

⁷The sense of ‘worse’ at issue here is ‘*ex ante* worse’. And, the sense of ‘betterness’ that is now used belongs to *ex ante* axiology, as opposed to *ex post* axiology, which is at issue in standard discussions of population

When arguing for the Intrapersonal Repugnant Conclusion, Nebel considers a couple that is planning to conceive a child by injecting a single sperm into a single egg.⁸ Suppose that only one person could possibly originate from this pair of gametes—call her Sally. Sally’s parents have three different ways they can do this injection: \mathcal{A} , \mathcal{Z} and $\mathcal{A}+$ (see table 1). \mathcal{A} will give Sally a very happy life at welfare level a if state 1 obtains, but she will not exist if state 2 obtains. \mathcal{Z} will give Sally a low positive life at welfare level z in both outcomes. Lastly, $\mathcal{A}+$ will give Sally welfare level $a+$ (slightly above a) if state 1 obtains and welfare level $z-$ (slightly below z) if state 2 obtains.

TABLE 1
THE INTRAPERSONAL ARGUMENT

	State 1 (p)	State 2 ($1 - p$)
\mathcal{A}	a	
$\mathcal{A}+$	$a+$	$z-$
\mathcal{Z}	z	z

Nebel argues that $\mathcal{A}+$ is better than \mathcal{A} for Sally because Sally’s welfare is higher if state 1 obtains, and her life would be worth living if state 2 obtains. This is supported by the following principle:⁹

Probable Addition Principle: If, in every state of the world in which a person S would exist in Y , S would be better off in X , and if, in every other state, S ’s life would be worth living in X , then X is better than Y for S .¹⁰

axiology.

⁸Nebel (2019, p. 313).

⁹Nebel (2019, p. 315).

¹⁰One could object that if \mathcal{A} is chosen and state 2 obtains, then Sally does not exist, and therefore there

Next, Nebel argues that \mathcal{Z} must be better than $\mathcal{A}+$ when the probability of state 1 is very small. Suppose that it is one-in-a-googolplex. Then, it would be irresponsible for Sally's parents to choose $\mathcal{A}+$ instead of \mathcal{Z} , as $\mathcal{A}+$ has such a small probability of resulting in a better outcome ($a+$ instead of z) and a very high probability of resulting in a worse outcome ($z-$ instead of z). This is supported by the following principle:¹¹

Minimal Prudence: No matter how good some life would be¹², there is some small probability and some pair of mediocre lives such that certainty of the better mediocre life would be better for some person S than a gamble that might yield the very good life but would almost certainly yield the worse mediocre life.

Next, if $\mathcal{A}+$ is better than \mathcal{A} (by the Probable Addition Principle), and \mathcal{Z} is better than $\mathcal{A}+$ (by Minimal Prudence), it follows by transitivity that \mathcal{Z} must be better than \mathcal{A} for Sally—which is the Intrapersonal Repugnant Conclusion.¹³ So, accepting the Probable Addition Principle, Minimal Prudence and the transitivity of *better than* leads to the Intrapersonal Repugnant Conclusion. This conclusion is not repugnant, but Nebel argues that accepting it leads to the Repugnant Conclusion. Next, I will summarize the second stage of his argument.

is no one for whom it would have been better had $\mathcal{A}+$ been chosen instead. Thus, it is not the case that $\mathcal{A}+$ must be better than \mathcal{A} for Sally. Nebel (2019, §V) discusses similar concerns, but I will not address them here as my argument does not rely on them.

¹¹Nebel (2019, p. 316).

¹²An exception here could be an infinitely good life. An agent who maximizes expected value would prefer a gamble with any probability of an infinitely good outcome.

¹³This argument is structurally analogous to Parfit's Mere Addition Paradox and Huemer's Benign Addition Paradox. See Parfit (1984, ch 19) and Huemer (2008, pp. 899–933).

1.2 From intrapersonal to interpersonal Repugnant Conclusion

Nebel considers a simplified case to show the derivation of the Repugnant Conclusion from its intrapersonal analogue. Consider these two outcomes (see table 2):

The Repugnant Conclusion:

A₀: Ann has a very happy life at welfare level a .

Z: Bob, Cat and Dan have mediocre lives at welfare level z .

Suppose that the Intrapersonal Repugnant Conclusion is true and that (unrealistically) a 1/3 chance of having a very happy life at welfare level a is worse for a person than certainly having a mediocre life at welfare level z .¹⁴ Nebel argues that this will lead to the conclusion that **Z** is better than **A₀** (i.e., the Repugnant Conclusion).

TABLE 2
THE REPUGNANT CONCLUSION

	Ann	Bob	Cat	Dan
A₀	a			
Z		z	z	z

Nebel considers three prospects: A , A^* and Z (see table 3). Prospect A certainly results in outcome **A₀** (Ann has a very happy life at welfare level a), while prospect A^* results in either Bob, Cat or Dan getting a very happy life at welfare level a , each with a 1/3 probability. Nebel argues that A and A^* are equally good, given that they are egalitarian prospects that guarantee equally good outcomes.¹⁵ Both prospects result in one person existing at welfare

¹⁴Nebel's argument can be generalized to more realistic instances of the Intrapersonal Repugnant Conclusion with probabilities smaller than 1/3. Nebel (2019, p. 323).

¹⁵Nebel (2019, p. 319).

level a , and the probabilities in A^* do not favor anyone over the others.¹⁶ If the outcomes are equally good, and no one is unfairly advantaged in either prospect, then the prospects must be equally good.

TABLE 3
TO THE REPUGNANT CONCLUSION

	State 1 (1/3)				State 2 (1/3)				State 3 (1/3)			
	Ann	Bob	Cat	Dan	Ann	Bob	Cat	Dan	Ann	Bob	Cat	Dan
A	a				a				a			
A^*		a					a					a
Z		z	z	z		z	z	z		z	z	z

Nebel argues that A and A^* are equally good, and that Z is better than A^* .

Next, compare prospect A^* to prospect Z , which guarantees outcome Z (Bob, Cat and Dan all have mediocre lives at welfare level z). As we have already assumed that a 1/3 chance of existing with a very happy life is worse for a person than certainly existing with a mediocre life, it follows that A^* is worse than Z for each Bob, Cat and Dan. Consequently, it must be overall worse, Nebel argues. This is supported by the following principle:

Weak Pareto for Equal Risk: For any egalitarian prospects X and Y , if X is better than Y for each person who might exist in either prospect, then X is better than Y .¹⁷

¹⁶According to some person-affecting views, these options might be incomparable in value because different people exist in their outcomes. For a discussion of the narrow person-affecting principle, see Parfit (1984, ch 16 and 18).

¹⁷Fleurbaey (2010, p. 656) and Nebel (2019, p. 320). ‘Egalitarian’ here should be understood in either *ex ante* or *ex post* sense amongst those who will exist—otherwise A^* would not be egalitarian.

This principle states that a prospect is overall better than another prospect if it is better for everyone who might exist in either prospect—at least when there is no risk of unfairness.

Nebel offers an inductive argument in favor of this Pareto principle.¹⁸ First, he argues that if only a single person might exist, and X is better than Y for that person, then X must be overall better than Y . Because X is better than Y for that person, we ought to prefer X for the sake of that person. And because we ought to prefer X for the sake of the only person who might exist, from an impartial perspective, we ought to prefer X . Thus, the Pareto principle is true for all prospects in which only a single person might exist.

Then Nebel considers egalitarian prospects X and Y in which any number n of people might exist. He argues that if X is better than Y for n number of people, then X must also be better than Y for $n + 1$ number of people. It cannot be the case that the addition of one more person (for whom X is also better than Y) suddenly reverses the betterness relation between X and Y —if X is better than Y for the additional person (and everyone else), then X must remain better than Y . Nebel argues that if this inductive step was unjustified, it should be for some reason having to do with some relation between the $(n + 1)$ th person and the others. If there was any risk of inequality, then the relational fact that some might be worse off than others could be blamed—but there is no such risk as the prospects are egalitarian. Therefore, Nebel argues that it is hard to see why the principle should be true for n but not for $n + 1$.¹⁹ So, Weak Pareto for Equal Risk is true when only a single person might exist, and by the inductive step, it is also true when two people might exist, and again by the inductive step, it is true when three people might exist, and so on. No matter how many people might exist, if X is better than Y for each of those people, then X is better than

¹⁸Nebel (2019, p. 321).

¹⁹Nebel (2019, p. 322).

Y.

Nebel's inductive argument for Weak Pareto for Equal Risk requires us to think that the principle fails even when only a single person might exist or that the difference between its true and false instances lies in the addition of only a single possible person whose existence will not generate a trade-off between different people's interests.²⁰ Nebel argues that neither of these possibilities seems very plausible, and thus, he concludes that we ought to accept Weak Pareto for Equal Risk.

The derivation of the Repugnant Conclusion from the Intrapersonal Repugnant Conclusion thus goes like this: Certainly existing with a mediocre life is better for a person than a 1/3 chance of existing with a very happy life (by the Intrapersonal Repugnant Conclusion). Thus, Z is better than A^* for each person who might exist in either prospect. By Weak Pareto for Equal Risk, Z is therefore overall better than A^* . Furthermore, A^* is equally good as A because they are egalitarian prospects that guarantee equally good outcomes. Consequently, as A and A^* are equally good, and Z is better than A^* , Z must be better than A . Lastly, because Z guarantees outcome Z and A guarantees outcome A_0 , and Z is better than A , outcome Z must be better than outcome A_0 —and we have arrived at the Repugnant Conclusion.²¹

Nebel considers the second stage of his argument to be more compelling than the argument for the Intrapersonal Repugnant Conclusion, so he focuses on possible responses to the latter.²² He discusses how Parfit's *Perfectionism* could respond. Perfectionism states that "even if some change brings a great net benefit to those who are affected, it is a change

²⁰Nebel (2019, p. 322).

²¹Nebel (2019, p. 323).

²²Nebel (2019, p. 324).

for the worse if it involves the loss of one of the best things in life.”²³ Perfectionism could respond that, even if some prospect brought a great net benefit to a person, it is worse for her if it lowers her probability of enjoying the best things in life. Perfectionism would therefore deny that \mathcal{Z} is better than $\mathcal{A}+$ for Sally. However, Nebel argues that it is irrational to prefer prospects that will almost certainly be worse for us in the pursuit of arbitrarily small chances of enjoying the best things in life—this would be an absurdly reckless decision theory.²⁴

In this chapter, I will argue that Perfectionism (and other population axiologies) can respond to Nebel’s challenge without being absurdly reckless by discounting small probabilities.^{25,26} However, first I will discuss an independent motivation for this kind of discounting.

2 Discounting small probabilities

According to orthodox decision theory, a rational agent always maximizes expected utility. However, doing this would lead one to make highly counter-intuitive choices when presented with options that have a small probability of a huge payoff. One such case is the St. Petersburg paradox, a version of which was originally proposed by Nicolaus Bernoulli in

²³Parfit (2004, p. 19).

²⁴Nebel (2019, p. 324).

²⁵As noted by an anonymous reviewer, someone might argue that we should not draw axiological conclusions from normative arguments about how we should treat very small probabilities. However, if this is true, then Nebel’s argument never gets off the ground because we should not draw axiological conclusions from the value of prospects when such normative arguments are relevant to their value.

²⁶Parfit regarded ignoring tiny chances as one of the five mistakes in ‘moral mathematics’. He (1984, p. 75) writes: “When the stakes are very high, no chance, however small, should be ignored. The same is true when each chance will be taken very many times. In both these kinds of cases, each tiny chance should be taken to be just what it is, and included in the calculation of the expected benefit. We can usually ignore a very small chance. But we should not do so when we may affect a very large number of people, or when the chance will be taken a very large number of times.” See Parfit (1984, pp. 73–75).

1713.^{27,28} The modern version of the game is played by flipping a fair coin until it lands on heads. The prize is then $\$2^n$, where n is the number of coin flips. This game has infinite expected monetary value, so an agent who maximizes expected monetary value would pay any finite amount to play it. However, this seems counter-intuitive. As Nicolaus Bernoulli, agreeing with his friend Gabriel Cramer, writes: “[T]here is no person of good sense who wished to give merely 20 coins.”²⁹

Daniel Bernoulli (cousin of Nicolaus Bernoulli) argues that we should not be willing to pay any finite sum to play the St. Petersburg game because of the diminishing marginal utility of money.³⁰ He argues that the expected utility of the game is finite even though it has infinite expected monetary value. However, one can change the game slightly to bypass this objection by changing the prize from money to something that has no diminishing marginal utility, such as (possibly) days of life.^{31,32} When the payoffs are utilities, the game

²⁷The game was then simplified by Gabriel Cramer in 1728 and published by Daniel Bernoulli in 1738. See Pulkamp (2013) and Bernoulli (1954).

²⁸Bostrom (2009) presents another case that involves a very small probability of a huge payoff.

²⁹Pulkamp (2013, p. 6).

³⁰More specifically, he argues that the utility of money equals the logarithm of the monetary value. See Bernoulli (1954). Cramer (Pulkamp, 2013, p. 4) also came close to suggesting that money has diminishing marginal utility: “One asks the reason for the difference between the mathematical calculation and the common value. I believe that it comes from this that the mathematicians value money in proportion to its quantity, and men of good sense in proportion to the usage that they may make of it.”

³¹Monton (2019, p. 2). Relatedly, 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).

³²Although one could argue that the longer one has lived, the less valuable extra days of life are. Temkin writes: “I believe that in many cases, though certainly not all, once people have experienced certain kinds of events ‘enough’ times in their lives, there will be a diminishing marginal value to subsequent similar experiences.” See Temkin (2008, p. 208). One could also add that at some point, it is not even possible to have new valuable experiences that are different enough from one’s past experiences such that this diminishing marginal utility does not happen—at some point everything worth experiencing has been experienced.

has infinite expected utility.³³ Nevertheless, few would sacrifice the rest of one's days as a payment to play the game in the hopes of living longer (though almost certainly dying soon), and this reluctance seems rational. Furthermore, if the game has infinite expected utility and we value gambles at their expected utilities, then we value the St. Petersburg game more than any of its possible (finite) payoffs—which seems clearly irrational.³⁴

Nicolaus Bernoulli, in turn, argues that in order to solve this paradox, we ought to discount very small probabilities down to zero—let's call this *Probability Discounting*. 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.”³⁵ More recently, Smith and Monton have argued for the same idea.³⁶ Monton argues that in order to avoid the fate of the expected utility maximizers, we need to either limit the high utility numbers or discount the small probability numbers in cases that involve very small probabilities of huge payoffs.³⁷ However, he argues that introducing a utility cap would be ethically problematic because one can always add more agents into the utility calculation and that the utilities of those individuals matter regardless of how many agents already exist. Thus, bounding utility is not viable³⁸, which leaves the only other option: discounting very small probabilities. Monton then argues that very small probabilities need to be discounted down to zero instead

³³Many decision theorists reject an unbounded utility function. However, even if utility is bounded, the expected utility of the St. Petersburg game can still be very high if the upper bound of utility is very high.

³⁴Huemer (2016, pp. 34–35) and Russell and Isaacs (2021).

³⁵Pulskamp (2013, p. 2); the German original von Spieß (1975).

³⁶Smith (2014) considers it permissible to discount very small probabilities down to zero, while Monton (2019) argues that one is rationally required to do so. Smith argues that Probability Discounting is a way of getting a reasonable expected utility for the Pasadena game. See Nover and Hájek (2004).

³⁷Monton (2019, p. 5).

³⁸Standard axiomatizations of expected utility maximization, such as the von Neumann-Morgenstern utility theorem, require utility to be bounded—or else the continuity axiom is violated. See Kreps (1988, p. 63).

of merely reducing those probabilities because one can always increase the payoffs of the games by a sufficient amount to compensate for those reduced probabilities.³⁹

To summarize, orthodox decision theory gives highly counter-intuitive recommendations in cases that involve very small probabilities of huge payoffs. In response to such cases, some have argued that we ought to discount very small probabilities down to zero.⁴⁰ Next, I will argue that Probability Discounting can also solve the Intrapersonal Addition Paradox.

3 A response to Nebel's inductive argument

In this section, I will argue that the inductive step of Nebel's inductive argument for Weak Pareto for Equal Risk is unjustified due to the cumulative nature of probabilities if one engages in Probability Discounting.

Recall that Sally's parents need to make a choice between \mathcal{A} , $\mathcal{A}+$ and \mathcal{Z} . For the purposes of this chapter, let's grant the Probable Addition Principle. Consequently, $\mathcal{A}+$ is better than \mathcal{A} for Sally. Also, I accept Minimal Prudence; \mathcal{Z} is better than $\mathcal{A}+$ for Sally when the probability of state 1 is very small because very small probabilities should be discounted down to zero. Sally's parents should ignore the small possibility of Sally getting a life at welfare level $a+$ and compare the options for their remaining outcomes: a life at welfare level

³⁹Monton (2019, p. 5).

⁴⁰Although there may also be some more fundamental justification for Probability Discounting. According to Monton, maximizing expected utility is a mistake because "you only live once", and the prescription to maximize expected utility does not take seriously the importance of how one's life actually goes. In contrast, Smith's argument is that decision theory tells us to ignore outcomes with zero probability, and because decision-making is a practical activity, infinite precision cannot be required. Smith also argues that Probability Discounting is a way of getting a unique expected value for the Pasadena game. See Nover and Hájek (2004) on the Pasadena game. For a discussion of other possible justifications, see Monton (2019). For further discussion of discounting small probabilities down to zero, see Smith (2014, 2016), Hájek (2014) and Isaacs (2016).

$z-$ in $\mathcal{A}+$ and a life at welfare level z in \mathcal{Z} . Because z is greater than $z-$, \mathcal{Z} is better than $\mathcal{A}+$ for Sally. Consequently, the argument for the Intrapersonal Repugnant Conclusion goes through.⁴¹ However, I will argue that this does not lead to the Repugnant Conclusion.

I will use the following principle in my argument:

Risky Non-Repugnance: q chance (or greater) of obtaining at least one life at a high welfare level a is better than certainly obtaining any number of lives at a low welfare level z , where q is the smallest probability that should not be discounted down to zero.^{42,43}

Risky Non-Repugnance states that a non-negligible probability of at least one very good life is better than certainty of any number of low positive lives. It can be supported with Perfectionism and some other types of Value Superiority. Value Superiority avoids the Repugnant Conclusion because, according to it, no quantity of low positive lives could ever be as good as some number of very good lives—very good lives are lexically superior to low positive lives. Then, one can argue that whatever the smallest probability that should not be discounted is, one should choose that probability of obtaining at least one very good life

⁴¹Actually, it might not: If one discounts the probability of state 1 down to zero, then the only outcome left in \mathcal{A} is non-existence. Thus, one cannot use the Probable Addition Principle when comparing \mathcal{A} and $\mathcal{A}+$ because one is comparing existence with non-existence.

⁴²Alternatively, one could discount anything up to and including some small probability q .

⁴³A related principle could be called *Reckless Risky Non-Repugnance*: any non-zero probability of obtaining one life at a high welfare level a is better than certainly obtaining any number of lives at a low welfare level z . Risky Non-Repugnance is also similar to what one might call *Intrapersonal Risky Non-Repugnance*: q chance of life at a high welfare level a is better for some person S than certainty of life at a low welfare level z . Another related principle could be called *Intrapersonal Risky Welfare-Level Superiority*: q chance of a life of some length t at a high momentary well-being level a is better for a person S than certainty of any length of life at a low momentary well-being level z . While Risky Non-Repugnance (and its reckless version) are about the contributive value of lives to the value of a population, the latter two principles are about what is good for individuals. One can also make reckless versions of them by replacing ‘ q chance’ with ‘any non-zero probability’.

instead of certainty of any number of low positive lives.⁴⁴

However, one might think Risky Non-Repugnance is implausible because the smallest probability that should not be discounted down to zero might be very small. Say it is one-in-a-trillion. Then, according to Risky Non-Repugnance, a one-in-a-trillion chance of one very good life is better than certainty of any number of low positive lives. A few things can be said in favor of Risky Non-Repugnance. First, there are cases in which Risky Non-Repugnance does not seem counter-intuitive: q chance of one very good life is better than certainty of any number of 10-second-lives consisting of a barely positive emotion. No quantity of 10-second-lives could ever be as good as one very good life, and q chance of a very good life is still better than those 10-second-lives. Secondly, q might actually be higher than one-in-a-trillion. However, by definition, q is a probability that should not be discounted down to zero—it is a probability that we should pay attention to and consider non-negligible.

Lastly, finding an intuitively acceptable population axiology is notoriously difficult, and this task gets even harder when we take risk into account. Finding Risky Non-Repugnance

⁴⁴Some versions of Value Superiority might accept the following principle instead of Risky Non-Repugnance:

Weak Risky Non-Repugnance: There is some probability p (less than 1) and some number n of very good lives such that p chance of n very good lives is better than certainty of any number of low positive lives.

On these versions of Value Superiority, the value of additional positive lives at some welfare level w diminishes the more such lives there already are, and their total contributive value approaches some upper bound. As this bound is higher for very good lives than for low positive lives, some number of very good lives is better than any number of low positive lives. However, the probability of obtaining very good lives might have to be high for that prospect to be better than certainty of any number of low positive lives. Lazar and Lee-Stronach (2019) defend this kind of approach in the context of limited aggregation and risk. They argue against an *infinitist* approach (such as Risky Non-Repugnance), which posits an infinite value difference between higher and lower considerations.

counter-intuitive is not a decisive reason for rejecting it if one must bite the bullet anyway and the alternatives are even worse. Consider, for example, this implication of Expected Total Utilitarianism:

Risky Very Repugnant Conclusion: For any number of very good lives that could be obtained for certain, there is a prospect that certainly gives many very bad lives together with a small probability that in addition there will exist some very large number of barely good lives, which is better, provided that the quantity of barely good lives is sufficient.

The quantity of barely good lives would, of course, have to be enormous for this to be true, as the goodness of those mediocre lives must be enough to outweigh the badness of the very bad lives and the risk introduced. Nevertheless, the world would almost certainly be arbitrarily bad. And, in the best-case scenario, it would only contain very bad and barely positive lives—yet Expected Total Utilitarianism would still recommend that option. In comparison, Risky Non-Repugnance does not seem counter-intuitive.⁴⁵

Now, recall that Weak Pareto for Equal Risk states that we ought to prefer prospects that are better for everyone. Next, I will argue that—contrary to the inductive step—the addition of only a single possible person can make the difference between the true and false instances of Weak Pareto for Equal Risk. I will illustrate my argument using the choice Sally’s parents have to make, and for simplicity, I will compare \mathcal{A} and \mathcal{Z} (instead of $\mathcal{A}+$ and \mathcal{Z}). Consider the following situation:

⁴⁵However, views that imply Risky Non-Repugnance might have other—even more counter-intuitive—implications than Risky Non-Repugnance. Ultimately, we must compare complete theories against one another.

Iterated Sally's Parents' Choice: A great number of couples face the choice between \mathcal{A} and \mathcal{Z} , and the probability of obtaining a very good life is independent every time if \mathcal{A} is chosen repeatedly.

If \mathcal{A} is chosen repeatedly, the probability of obtaining at least one very good life accumulates. Thus, with some k number of choices of \mathcal{A} , the probability of obtaining at least one very good life is less than the threshold for Probability Discounting (and thus should be discounted down to zero). But with $k + 1$ number of choices, that probability is above or equal to the threshold (and thus should *not* be discounted down to zero). If \mathcal{A} is chosen enough times, the probability that at least one person gets a very good life is greater than or equal to q —a prospect that is better than certainty of any number of low positive lives, according to Risky Non-Repugnance.

The inductive step of the inductive argument for Weak Pareto for Equal Risk is unsound. Certainty of k individuals obtaining low positive lives is better than the prospect of them all having p chance of getting a very good life because the cumulative probability of obtaining at least one very good life is still rationally negligible. The former prospect is also better for each individual who might exist because the probability of them obtaining a very good life is also rationally negligible. However, certainty of $k + 1$ individuals obtaining low positive lives is worse than them all having a p chance of getting a very good life. This is because the cumulative probability of obtaining at least one very good life is no longer rationally negligible, and Risky Non-Repugnance judges that prospect to be better than $k + 1$ individuals obtaining low positive lives for certain.⁴⁶ So, Weak Pareto for Equal Risk is true for

⁴⁶A weaker principle than RNR would be sufficient here, as the principle only needs to state that a q chance of a very good life is better than certainty of $k + 1$ low positive lives.

k individuals, but it is not true for $k + 1$ individuals. The accumulation of probabilities is the relational fact that renders the inductive step false.⁴⁷ Consequently, the value of a series of choices is not just an aggregation of the value of each individual choice.^{48,49}

It is worth pointing out that those who accept a probability-discounting version of Expected Total Utilitarianism must also reject the inductive step. Consider the case of Sally again. Say that the very good a -life is sufficiently good such that the expected value of obtaining one such life with probability q is higher than the value of certainly obtaining $k + 1$ lives at welfare level z . As before, with up to k number of choices, \mathcal{A} is judged overall worse than \mathcal{Z} because the possibility of getting at least one very good life is rationally negligible. However, with $k + 1$ number of choices, \mathcal{A} is judged overall better than \mathcal{Z} because q chance of at least one very good a -life has a higher expected value than certainty of $k + 1$ lives at welfare level z . Thus, the addition of a single possible person can make the difference between the true and false instances of Weak Pareto for Equal Risk if one accepts a probability-discounting version of Expected Total Utilitarianism.

Lastly, the probabilities (of getting a very good life) are not independent in the second stage of Nebel's argument because either Bob, Cat or Dan would have a very good life for

⁴⁷One could object that it is implausible that a very small difference in probabilities (from just below q to just above it) can make all the difference. One possible response is that the threshold q might be vague.

⁴⁸There is also the question of what decision-makers should do when they know they face a series of choices involving a very small probability of a huge payoff. Should they refrain from discounting in the last choice of the series, even if they would discount in a similar one-off choice? To deal with iterated choices, probability discounters could accept *Resolute Choice*. A resolute agent chooses according to any plan they have adopted earlier as long as nothing unexpected has happened since then. Probability discounters can then form a plan to not discount in any of the choices in the series, even if they would discount in a similar one-off choice. See McClennen (1990) on Resolute Choice. However, Probability Discounting in combination with Resolute Choice leads to untenable results. See §3.2 in Chapter 5 of this thesis.

⁴⁹As noted by an anonymous reviewer, other views on which the value of conjunction of acts is different from the sum of the value of the individual acts can also block Nebel's argument. For a discussion of this possibility in the context of limited aggregation and risk, see Tadros (2019).

certain if A^* were chosen. However, my argument provides a justification for rejecting the inductive step of Nebel's inductive argument. Next, I will argue that one should reject Ex Ante Pareto principles such as Weak Pareto for Equal Risk if one discounts very small probabilities down to zero.

4 An argument against Ex Ante Pareto from discounting small probabilities

Contrary to Ex Ante Pareto, a prospect can be impersonally better, even if it is worse for everyone. This happens when the probability of an individual receiving some good (or harm) is discounted down to zero, but—collectively—the probability that at least one person receives that good (or harm) is large enough to be taken into account. Z is better than A for Sally, but it would be impersonally worse if the choice of Z over A was repeated a great number of times. Also, Z is better than A^* for all Bob, Cat and Dan, but A^* is impersonally better than Z . In both cases, the probability of obtaining a very good life is rationally negligible for the individual, but the probability of someone obtaining such a life is non-negligible.

Accepting Ex Ante Pareto and engaging in Probability Discounting gets one in trouble. Consider the following case:

Celebratory Gunfire: Someone shoots into the air in an area full of people during a celebration, which causes people to feel excitement for a few seconds. The probability of any particular individual being hit by the bullet when it falls is negligibly small, but there is a high probability that someone is hit by it.

In this case, the prospect of shooting into the air is *ex ante* better than not shooting for

everyone; each individual feels excitement for a few seconds, and the probability of any particular individual being hit by the bullet is rationally negligible. However, the goodness of everyone feeling excitement is not enough to outweigh the badness of the likely injury. Consequently, shooting into the air is *ex ante* impersonally worse than not shooting—which contradicts Ex Ante Pareto. If one accepts Probability Discounting, one should also hold that impersonally better prospects are possible, or one would permit the infliction of arbitrarily severe harms for little or no benefits.⁵⁰

I have argued that sometimes we should not choose prospects that are better for everyone. However, it also seems that sometimes it is permissible to only care about what is good for particular individuals instead of what would be impersonally best. This seems appropriate when one attempts to benefit one's family, friends or oneself. So, it would be permissible for Sally's parents to choose \mathcal{Z} over \mathcal{A} , even if a great number of parents faced the same choice, because they are concerned with the welfare of their future child instead of attempting to make the world better overall. This discrepancy between what is *ex ante* good for individuals and what is impersonally *ex ante* good is a price to pay for Probability Discounting.

5 Conclusion

I have argued that one can solve the Intrapersonal Addition Paradox if one discounts very small probabilities down to zero. The Repugnant Conclusion does not follow from its intrapersonal analogue because we should reject one of the premises of the argument, namely, Weak Pareto for Equal Risk. First, this principle is not supported by the inductive argument

⁵⁰Pareto principles have also been challenged before. See for example Sen (1970), Mongin (1995) and Temkin (2000).

if one engages in Probability Discounting. This is because the inductive step is unjustified due to the cumulative nature of probabilities. Secondly, this principle licenses the infliction of arbitrarily severe harms for little or no benefits if combined with Probability Discounting. This happens when the probability of harming each individual is small, but there is a high probability that someone is harmed.

Furthermore, we have independent reasons for engaging in Probability Discounting because it enables one to get intuitively right responses in decision-theoretic problems that involve very small probabilities of huge payoffs. To conclude, in order to avoid ending up like the expected utility maximizers in these situations, we must discount very small probabilities down to zero. But then, we must give up Ex Ante Pareto. And then, we can solve the Intrapersonal Addition Paradox. This argument has implications for other ethical debates as well, as this solution is somewhat similar to the solution posed to the problem of aggregation and risk.⁵¹ This chapter adds to the idea that this solution has a principled foundation in a more general claim in decision theory: very small probabilities have no prudential or moral significance.⁵²

References

Arrhenius, G. (2000), 'An impossibility theorem for welfarist axiologies', *Economics and Philosophy* **16**(2), 247–266.

Barrington, M. (forthcoming), 'Superiority Discounting implies the Preposterous Conclusion', *Utilitas* .

⁵¹Thanks to an anonymous reviewer of *Ethics* for raising this point. For discussions of aggregation and risk, see Lazar (2018), Lazar and Lee-Stronach (2019), Tadros (2019), and Horton (2020).

⁵²See Barrington (forthcoming) for criticism of combining Probability Discounting with a lexical view.

- 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.
- Fleurbaey, M. (2010), 'Assessing risky social situations', *Journal of Political Economy* **118**(4), 649–680.
- Greaves, H. (2017), 'Population axiology', *Philosophy Compass* **12**(11), e12442.
- Hájek, A. (2014), 'Unexpected expectations', *Mind* **123**(490), 533–567.
- Horton, J. (2020), 'Aggregation, risk, and reductio', *Ethics* **130**(4), 514–529.
- Huemer, M. (2008), 'In defence of repugnance', *Mind* **117**(468), 899–933.
- Huemer, M. (2016), *Approaching infinity*, Palgrave Macmillan, New York.
- Isaacs, Y. (2016), 'Probabilities cannot be rationally neglected', *Mind* **125**(499), 759–762.
- Kosonen, P. (2021), 'Discounting small probabilities solves the Intrapersonal Addition Paradox', *Ethics* **132**(1), 204–217.
- Kreps, D. M. (1988), *Notes on the Theory of Choice*, Westview Press, Boulder.
- Lazar, S. (2018), 'Limited aggregation and risk', *Philosophy & Public Affairs* **46**(2), 117–159.
- Lazar, S. and Lee-Stronach, C. (2019), 'Axiological absolutism and risk', *Noûs* **53**(1), 97–113.
- McClellenn, E. F. (1990), *Rationality and Dynamic Choice: Foundational Explorations*, Cambridge University Press, Cambridge.

- McTaggart, J. M. E. (1927), *The Nature of Existence*, Vol. 2, Cambridge University Press, Cambridge.
- 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.
- Mongin, P. (1995), 'Consistent bayesian aggregation', *Journal of Economic Theory* **66**(2), 313–351.
- Monton, B. (2019), 'How to avoid maximizing expected utility', *Philosophers' Imprint* **19**(18), 1–24.
- Nebel, J. M. (2019), 'An intrapersonal addition paradox', *Ethics* **129**(1), 309–343.
- Nover, H. and Hájek, A. (2004), 'Vexing expectations', *Mind* **113**(450), 237–249.
- Parfit, D. (1984), *Reasons and Persons*, Clarendon Press, Oxford.
- Parfit, D. (2004), Overpopulation and the quality of life, *in* T. Tännsjö and J. Ryberg, eds, 'The Repugnant Conclusion: Essays on Population Ethics', Springer Netherlands, Dordrecht, pp. 7–22.
- Pulskamp, R. J. (2013), '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_petersburg_game.pdf

- Russell, J. S. and Isaacs, Y. (2021), 'Infinite prospects', *Philosophy and Phenomenological 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.
- Sen, A. (1970), 'The impossibility of a paretian liberal', *Journal of Political Economy* **78**(1), 152–157.
- 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.
- Tadros, V. (2019), Localized restricted aggregation, in D. Sobel, P. Vallentyne and S. Wall, eds, 'Oxford Studies in Political Philosophy Volume 5', Oxford University Press, Oxford, pp. 171–203.
- Temkin, L. S. (2000), Equality, priority, and the levelling down objection, in M. Clayton and A. Williams, eds, 'The Ideal of Equality', Macmillan, London, pp. 126–161.
- Temkin, L. S. (2008), 'Is living longer living better?', *Journal of Applied Philosophy* **25**(3), 193–210.
- Temkin, L. S. (2012), *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*, Oxford University Press, New York.

von Spieß, O. (1975), Zur vorgeschichte des Petersburger problems, *in* B. van der Wæerden, ed., 'Die Werke von Jakob Bernoulli Band 3: Wahrscheinlichkeitsrechnung', Birkhäuser, Basel, pp. 557–567.