

IGNORING THE IMPROBABLE

ABSTRACT

Many problems in ethics and decision theory are solved by ignoring small probabilities. This paper argues that ignoring possibilities—even extremely remote ones—comes with a hefty cost for one’s ability to make rational decisions. First, it makes rationality depend on how the world is described. Second, agents become insensitive to differences in the probabilities and values contained in excluded states. Third, ignoring a state affects the expected value of other acts in the decision problem, making agents inaccurately assess the choiceworthiness of those acts. (For instance, agents will refuse arbitrarily favourable bets on the excluded state.) And unfortunately, none of the most intuitive patches to the view can avoid these issues.

1. INTRODUCTION

A number of proposals in ethics and decision theory instruct agents to ignore sufficiently improbable possibilities—a strategy I will call “probability discounting”. Probability discounting solves many stubborn problems. Absolutist moral theories are able to give sufficient guidance in contexts of risk if we ignore sufficiently small chances of violating our duty (see especially Lee-Stronach 2018; also Kagan 1989: 89-92; Aboodi, Borer, and Enoch 2008; Hawley 2008; Haque 2012; Bjorndahl, London, and Zollman 2017; Lazar 2017; Tarsney 2018). Pascal’s Mugging is solved by ignoring the possibility that the mugger is telling the truth (Jordan 1994: 218-19; Chalmers 2017; Schwitzgebel 2017: 273; Monton 2019). The St. Petersburg Paradox (and its variants) are solved by ignoring the possibility that the game goes on a sufficient number of coin

flips (Bernoulli 1738; d'Alembert 1761; Buffon 1777; Condorcet 1785; Borel 1962; Jordan 1994: 217-18; Buchak 2013: 73-74; Smith 2014, 2016; Robert 2018; Monton 2019). (I will limit my discussion to probability discounting's use as a strategy for solving these three decision-theoretic problems, ultimately focusing on the problem from moral absolutism, but it is worth noting its use in myriad other contexts for a variety of reasons.¹) This paper argues that ignoring possibilities comes with a price.

1.1 MORAL ABSOLUTISM AND RISK

Absolutist moral theories claim that consequentialist considerations can never justify violating a moral duty. For instance, if agents have an absolutist moral duty to avoid killing innocent people, no number of averted headaches could justify killing an innocent. Absolutism famously has trouble telling agents how much *risk* of violating their duties agents may permissibly take.

¹ The suggestion that “*de minimis*” (negligible) risk should be ignored entirely in decision making is ubiquitous in public health, medical, and risk-assessment scholarship. (Peterson (2002) provides a good introduction to *de minimis* principles and their use in these fields.)

It has also been used by philosophers in the literature on collective action. For instance, Miller and Sartorius (1979) argue that free-riding on the contribution of others is often rational and moral, because the probability of any one individual's contribution (e.g., voting) determining whether the (perceived) mutually advantageous outcome is produced (e.g., the best candidate being elected) is so small as to be rationally negligible. In a similar vein, Buchanan (1979) criticises the Marxist claim that the proletariat will revolt due to their collective interests being served by the revolution. Buchanan contends that while each proletarian might be greatly benefitted by the revolution, revolting is not in any individual's self-interest because the probability that *their* revolting will make the revolution successful is rationally negligible. More recently, Kosonen (2021) has endorsed probability discounting to attack *ex ante* pareto principles: those holding that if some prospect is better for everyone, then that prospect is simply better. Kosonen argues (in response to Nebel (2019: 320)), that a prospect's risk to each individual might be small enough to ignore, while the risk to the group is too great to ignore. So, the prospect might be better for each individual but worse from an impartial perspective.

At first pass, we might think that the absolutist views killings as *infinitely* more disvaluable than headaches.² Otherwise, some number of headaches would outweigh the killing. The most plausible way of modelling this view is with a multi-dimensional value function in a lexicographical ordering (e.g., Lee-Stronach 2018: 796-99): agents maximise primary value (that of upholding their duties), only looking to secondary value (that of upholding consequentialist considerations) to break ties.³ But on this approach, agents must *always* perform the act with the smallest risk of violating their duty. They “have to stay huddled in the corner, not daring to move, trying not to breathe too loudly” out of fear of raising their probability of killing someone (Kagan 1989: 89).⁴ This result is unacceptable because probabilities matter. Absolutists must disregard sufficiently small risks while nevertheless allowing substantial risks to outweigh *any amount* of secondary value.⁵

1.2 PASCAL’S MUGGING

Pascal is approached by Mugger, who asks for his wallet, promising to pay him back double its value. Being committed to maximising expected value, Pascal declines on the grounds that the probability of Mugger telling the truth is less than 0.5, so the deal has negative expected value.

² This modelling assumption is made by: Jackson and Smith (2006, 2016); Colyvan, Cox, and Steele (2010); Huemer (2010); Hayenhjelm and Wolff (2012); Hansson (2013); Bjorndahl, London, and Zollman (2017).

³ There are some alternative ways of modelling absolutist theories. Notably, Lazar and Lee-Stronach (2019) make the value of upholding the secondary consideration bounded in such a way that it can never agglomerate to outweigh the value of upholding one’s duty; However, the infinite-value approach is by far the most common, and the lexicographic model is the most plausible formalization of this approach.

⁴ Problems of this form are offered by: McKerlie (1986); Ashford (2003: 298); Jackson and Smith (2006, 2016); Colyvan, Cox, and Steele (2010); Huemer (2010); Fried (2012); Sobel (2012); Isaacs (2014); Holm (2016); Tenenbaum (2017); Alexander (2018).

⁵ Note that the puzzle here does not *only* arise for absolutist theories. For instance, consequentialist theories that endorse *value superiority* face the same task of answering the question for what probabilities is a gamble on the superior good and otherwise nothing better than a large quantity of the inferior good for sure?

Mugger then asks what probability Pascal assigns to the possibility of him telling the truth. Whatever (non-zero) number Pascal responds with, Mugger makes an offer valuable enough to counteract the risk, giving the deal positive expected value. The price they settle on is 1,000 quadrillion happy days of life. If Pascal maximises expected value, he will keep giving his wallet away to any crook who makes a large enough promise. Maximising expected value here appears to be irrational (Bostrom 2009).

1.3 THE ST. PETERSBURG PARADOX

How much would you be willing to pay to play the following game? A fair coin is tossed until it comes up heads. You will then be paid $\$2^n$, where n is the number of times the coin was flipped. So, if the coin comes up heads on the first toss, you win $\$2$; on the second toss, $\$4$; on the third toss, $\$8$. Intuitively, you should not pay very much to play the game. While you *could* end up with a substantial amount of money, the probability of this is vanishingly small. But if you maximise expected value, you will be willing to pay *any* finite amount for a single game because its expected value appears to be infinite:⁶

$$\frac{1}{2} \cdot 2 + \frac{1}{4} \cdot 4 + \frac{1}{8} \cdot 8 \cdots = 1 + 1 + 1 \cdots = \sum_{n=1}^{\infty} \left(\frac{1}{2}\right) \cdot 2^n = \infty$$

⁶ See Monton (2019: 2-3) for a finite version of the game, where bounding the agent's utility function does not solve the problem (a solution defended by Arrow 1971; Bassett 1987; Samuelson 1977; McClennen 1994).

1.4 PROBABILITY DISCOUNTING

A solution to all three problems is to posit a probabilistic threshold, t , governing which possibilities you must attend to, and which are so improbable that you should ignore them. If a state's probability is greater than or equal to t , you attend to it as usual. But if its probability is strictly less than t , you simply ignore it, treating it as though it were zero. (Presumably, after sufficiently improbable states have been excluded, you redistribute your credence into the remaining states by conditionalisation.)

Probability discounting allows absolutists to uphold their duties under risk by ignoring small chances of violating their duties while allowing substantial risks to outweigh any other considerations. It enables agents to hold onto their wallets in the face of philosophically literate muggers since they will not attend to the possibility of them telling the truth. And it caps the amount agents will be willing to pay to play the St. Petersburg game since the remote possibilities in which they win large sums of money will not contribute to the game's expected value. In what follows, I will focus on the problem for moral absolutism, but each problem can be extended to all views that endorse probability discounting.

2. PARTITION SENSITIVITY

Probability discounting generates partition-sensitive prescriptions; when states are described coarsely, the theory prescribes one act, but when they are described finely, it prescribes another. This feature results in inconsistent prescriptions when two relevantly identical decision problems should be described at different levels of generality:

Boris and Doris are two military commanders with absolutist moral duties to avoid killing civilians. Both would like to test an identical missile, which has a 0.04 probability of malfunctioning. If Boris's malfunctions, it will strike a building containing 100 civilians, while if Doris's malfunctions, it will hit one of two buildings: one containing 99 civilians and one containing 101.

Suppose $t = 0.03$. Since Boris's case involves a 0.04 probability of killing 100 people, he cannot ignore the possibility of killing civilians, and so is prohibited from testing the missile. (Let $-\infty$ correspond to the value of a duty violation, which is infinitely weightier than consequentialist value.)

	Safe (0.96)	\neg Safe (0.04)
Fire	Success	100 Deaths
\neg Fire	0	0

However, since Doris's missile could strike one of two buildings, she must split the 0.04 probability of malfunction into two states:

	Safe (0.96)	\neg Safe ₉₉ (0.02)	\neg Safe ₁₀₁ (0.02)
Fire	Success	99 Deaths	101 Deaths
\neg Fire	0	0	0

Since the probabilities of \neg Safe₉₉ and \neg Safe₁₀₁ are each 0.02, and $t = 0.03$, she ignores every possibility in which she would kill innocents and is permitted to test the missile. Probability discounting gives the bizarre result that testing the missile is forbidden for Boris but permitted for Doris.

This feature results in another problem: probability discounters may ignore arbitrarily large probabilities: no matter how risky an act is, there is some number of states n large enough that $1/n < t$; if the probability of violating the duty is split equally among the n states, each state's

probability will be below the threshold. As a result, the (arbitrarily probable) duty violation will be ignored and the (arbitrarily risky) act permitted.

2.1 PARTITION-SENSITIVE THRESHOLDS

One solution to the problem is to relativise the threshold to the partition: when states are partitioned coarsely, the threshold is correspondingly higher than when they are partitioned finely. For instance, perhaps Doris's threshold—where the probability of violating the duty is split among two states—should be half that of Boris's. As a result, both Boris and Doris would be forbidden from firing the missile.

But this solution is inadequate. Consider Morris, a third military commander whose decision is the same as Boris's, but whose missile is safer: the probability of Morris's missile malfunctioning is slightly less than t ($t - \epsilon$).

	Safe ($0.97 + \epsilon$)	\neg Safe ($0.03 - \epsilon$)
Fire	Success	100 Deaths
\neg Fire	0	0

Morris excludes \neg Safe from the decision problem and is permitted to test the missile. So far so good.

	Safe ₁ ($((0.97 + \epsilon)/n)$)	Safe ₂ ($((0.97 + \epsilon)/n)$)	...	Safe _n ($((0.97 + \epsilon)/n)$)	\neg Safe ($0.03 - \epsilon$)
Fire	Success + Prize ₁	Success + Prize ₂	...	Success + Prize _n	100 Deaths
\neg Fire	0	0	...	0	0

Now, pick some n large enough that $1/n < t$, and suppose that if the missile does *not* malfunction, Morris will be rewarded with one of n (differently valued) prizes. The partition will

now contain n equally probable states ($\text{Safe}_1 \dots \text{Safe}_n$), corresponding to the possible outcomes of the lottery.

Since the probability of each remaining state is less than $1/n$ —and by stipulation, $1/n$ is below the threshold—the probability discounter will ignore all states in the decision problem.

The discounter's solution is to lower t to below $1/n$, allowing states $\text{Safe}_1 \dots \text{Safe}_n$ to enter the equation. However, since $\neg\text{Safe}$ is lurking just below t , lowering the threshold will cause $\neg\text{Safe}$ to enter the equation, making the act forbidden. The result is that Fire is permitted *unless* it comes with the lottery—even though the lottery merely sweetens the deal! Since there is no way to have t below states ($\text{Safe}_1 \dots \text{Safe}_n$) but above $\neg\text{Safe}$, there is no way to avoid this result.

Partition problems appear to be a relatively deep feature of probability discounting theories. Orthodox decision theory avoids them because states' probabilities will have the same sum regardless of how they are partitioned. But when we discount improbable states, if some probability is split among sufficiently improbable states, they will count for nothing. Nevertheless, even if we think partition sensitivity is merely a framing problem (and so might be solved by standardising the way we describe decision problems), trickier problems await.

3. PROBABILITY INSENSITIVITY

The probabilistic threshold is too blunt an instrument. In excluding a problematic state from the decision problem, important features of that state go with it. In the first place, excluding a state prevents agents from deliberating on its probability:

You are a clinician who is preparing a vaccine for your next patient. A small portion of the vaccines are defective, and injecting someone with a defective vaccine will kill them. Despite your absolutist duty to avoid killing your patients, you can ignore risks less than one in 500,000. And luckily, the probability of each vaccine dose being defective is only one in 100 million. However, you recently learned that mixing a certain chemical into a vaccine dose has two effects. First, it turns the clear liquid bright pink. And second, it makes the dose one hundred times more deadly. While waiting for your patient, you try mixing the chemical into a dose to watch it change colour. And as your patient enters the room, you wonder if it would be wrong to inject her with it.

Obviously, you should not inject your patient with the pink vaccine: you have a perfectly good alternative that is one hundred times less likely to kill her. However, even though the probability of Pink being defective is one hundred times greater than that of the original vaccine, both possibilities are under the 1/500k threshold. So, you ignore these states and reason that your patient will be safely vaccinated, whichever vaccine you pick: you are indifferent between the two.⁷

	Neither Defective ($1 - 1/1m - 1/100m$)	Pink Defective (1/1m)	Original Defective (1/100m)
Pink	Vaccinated	Death	Vaccinated
Original	Vaccinated	Vaccinated	Death
None	0	0	0

But this result is wrong. There is no reason to impose the additional risk upon your patient, and no theory should treat this difference as entirely irrelevant to your decision.

⁷ For simplicity, the formalization of the decision problem ignores the state in which both are defective.

3.1 THRESHOLD MANIPULATION

Suppose the probability discounter objects to the fact that *we* were in charge of setting the threshold. “You tell *me* the relevant probabilities”, she insists, “and *I’ll* tell you the value of *t*”. Her strategy becomes apparent: wherever we set the probabilities, she will set *t* between the two, forcing us to attend to one possibility while ignoring the other and avoiding indifference. So, in a choice between a $1/1m$ and a $1/100m$ risk, perhaps $t = 1/50m$.

	Neither Defective ($1 - 1/1m - 1/100m$)	Pink Defective ($1/1m$)	Original Defective ($1/100m$)
Pink	Vaccinated	Death	Vaccinated
Original	Vaccinated	Vaccinated	Death
None	0	0	0

As the state where Pink is defective is now above the threshold, the theory gives the correct prescription: you ought to administer Original. If we lower the probability of Pink Defective to $1/60m$, then the probability discounter simply shifts *t* down to $1/70m$ —once again avoiding indifference. As long as there is *any* gap between the two probabilities, she will be able to set the threshold between them.

This solution—waiting to see the available acts before placing the threshold between the best and second-best acts—in addition to raising arbitrariness concerns, also violates

Irrelevant Alternatives One’s preference ordering between acts does not change if new acts are added to the decision problem.⁸

The principle gets at the following idea: if you prefer an apple to an orange, adding the option of a banana should not make you suddenly prefer an orange to an apple.

⁸ This articulation is from Peterson (2017: 59).

Consider the previous case, where the value of t is $1/50m$, so you ignore Original Defective but attend to Pink Defective. As a result, you express the following preference ordering: Original \succ None \succ Pink. However, now suppose you are told that new vaccines have just arrived that have been transported in ultra-low-temperature freezers, which reduces the rate of defectiveness tenfold—and makes them take on a bluish hue.

Naturally, the blue vaccine is now your best option. To extract this prescription, the probability discounter moves t from $1/50m$ to somewhere between $1/100m$ and $1/1b$, so you will ignore Blue Defective but not ignore Original Defective.

	Neither Defective ($1 - 1/1m - 1/100m - 1/1b$)	Pink Defective ($1/1m$)	Original Defective ($1/100m$)	Blue Defective ($1/1b$)
Pink	Vaccinated	Death	Vaccinated	Vaccinated
Original	Vaccinated	Vaccinated	Death	Vaccinated
Blue	Vaccinated	Vaccinated	Vaccinated	Death
None	0	0	0	0

However, your preference ordering is now: Blue \succ None \succ Original \sim Pink. While you preferred Original to Pink in the original problem, you are now indifferent between them. And while you preferred Original to None in the first case, the addition of Blue has reversed this ordering: you now prefer None to Original. Unfortunately for the probability discounter, there is no way to manipulate the threshold without opening the door for violations of Irrelevant Alternatives. But even if we are not concerned by the violation of Irrelevant Alternatives, a related problem cannot be solved by threshold manipulation.

4. VALUE INSENSITIVITY

Ignoring a possibility also prevents us from deliberating on its value:

You are a surgeon tasked with implanting a neurostimulator in a patient to treat their depression. You must choose which product to use: NS1 or NS2. The effectiveness of each device largely depends on the patient's genes, which you are unable to determine: Both devices have a 0.49 probability of successfully curing the patient's illness, a 0.49 probability of being unsuccessful, and a 0.02 chance of malfunctioning and paralyzing the patient. Although you have a duty to avoid harming your patient, you decide that the risk of paralysis is small enough to ignore. The only difference between the devices is that in the event of a malfunction, NS2 would not only paralyse her but *also* cure her depression.

	NS1 Works (0.499)	NS2 Works (0.499)	Malfunction (0.002)
NS1	Cured	0	Paralysis
NS2	0	Cured	Paralysis + Cured
None	0	0	0

Since you are only looking at the first two states, and in these states, each act has a 0.5 probability of curing the patient with no side effects and a 0.5 probability of being ineffective, NS1 and NS2 have the same expected value. However, you should not be indifferent between the two options. NS2 guarantees that, even if the patient is paralysed, her depression will be cured, while NS1 does not. Clearly, you should use NS2.

4.1 DOMINANCE PRINCIPLES

Some probability discounters deal with problems like this by adding a dominance principle to their theories, stipulating that agents always prefer dominating acts (Hájek 2014: 557; Smith 2014; Monton 2019: 20).⁹ But this would not help in the case presented because neither act is dominating

⁹ Isaacs (2016) criticizes Smith's theory on the grounds that accepting dominance (as both parties believe we should) results in further problems: Agents will risk (low probabilities of) arbitrarily disvaluable outcomes for arbitrarily small reward.

the other: NS1 is better in S_1 and NS2 is better in S_2 . A more careful solution would be to instead add a *stochastic dominance* principle, stipulating that agents prefer stochastically dominating acts. But this view would have problems too: Agents will accept conditional bets but refuse to take the same bets once those conditions are met—as they will discount the unconditional probability but not the conditional probability of losing the bet.

5. PERIPHERAL ACTS

States can matter differently to different acts, and so the effects of ignoring a state will not be isolated to a single act.

You are watching TV late at night with your friend Jimmy when he suggests playing trumpet instead. You would like to, but you promised not to play the trumpet while your housemate, Leonard, is home, and you have a duty to keep that promise. The problem is you are unsure if he is home. Leonard is the fire department chief, so he only sleeps at home on one (randomly determined) night per calendar year. And since today is January 1st, each of the next 365 nights has a $1/365$ chance of being his night off.

Suppose you are permitted to ignore probabilities smaller than $1/350$, so you can ignore the state where Leonard is home and enjoy your trumpeting.

	Home ($1/365$)	\neg Home ($364/365$)
Trumpet	Broken Promise	Fun
None	0	0

But before you can give Jimmy an answer, he offers you a wager: if Leonard is home, he will pay you \$1 billion; if he is not home, you owe him a penny. (Jimmy is a hundred-billionaire with a gambling addiction, so you know he is good for the money.) Obviously, the bet is extremely

valuable: a ticket to a 365-ticket lottery that pays out \$1 billion is well worth the price of 1¢. However, probability discounting theories do not get this result.

	Home (1/365)	¬Home (364/365)
Trumpet	Broken Promise	Fun
Bet	\$1b	-1¢
None	0	0

Since you have discounted the probability of Leonard being home to zero, you see the bet as offering a 100% chance of losing 1¢ with a 0% chance of winning \$1 billion. According to probability discounting theories, the bet is tantamount to throwing your money away.

5.1 ACT-RELATIVE THRESHOLDS

Once a bet against the excluded state has been added to the decision problem, we *want* to shift t below that state, so we are sensitive to the value of the bet. Unfortunately, as we saw in §3.1, shifting t when an option is added leads to violations of Irrelevant Alternatives. So perhaps the probability discounter would employ a different strategy: relativising the threshold to acts. On this view, the value of Trumpet's outcome in Home would not contribute to its expected value, while the value of *Bet's* outcome in Home *would* contribute to its expected value. But unfortunately, there are acts that this approach cannot accommodate.

Suppose you go to play trumpet with Jimmy, having determined that you can ignore the possibility of Leonard being home. He offers you a choice of trumpet—Gold or Silver—over which you have no preference. However, he explains that if you take Gold, it will be understood that you have accepted the bet, and if you take Silver, it will be understood that you have declined. Obviously, you should prefer Gold to Silver: Both acts are identical except Gold comes with a favourable bet. But there is no possible value of t that generates this prescription: If t is above

1/365, you will ignore the possibility of winning the bet (and will be indifferent between the trumpets); if t is below (or equal to) t , you will attend to the possibility of winning the \$1 billion but *also* to the possibility of waking Leonard—and no amount of money could justify breaking your promise. Wherever we set the threshold, the probability discounter cannot prefer to take the bet.

	Home (1/365)	¬Home (364/365)
Gold	Broken Promise + \$1b	Fun - 1¢
Silver	Broken Promise	Fun
None	0	0

As a result, trumpet-playing is permitted, but trumpet-playing *with a favourable bet* is forbidden.

6. CONCLUSION

The challenges for probability discounting views appear to be serious. Partition sensitivity seems to be a deep feature of the view because arbitrarily probable states may always be split into enough states that each state's probability is below t . As a result, agents will receive different prescriptions in relevantly similar cases and be willing to take arbitrarily great risks. Insensitivity to both the probability and the value of risks also appears inescapable since excluded states cannot contribute to the expected value of acts. Finally, since a state can be relevant to multiple acts, excluding states from a decision problem has unintended effects on the choiceworthiness of peripheral acts.

References

- Aboodi, Ron, Adi Borer, and David Enoch. 2008. 'Deontology, Individualism, and Uncertainty: A Reply to Jackson and Smith', *The Journal of Philosophy*, 105: 259-72.
- Alexander, Larry. 2018. 'The Need to Attend to Probabilities—For Purposes of Self-Defense and Other Preemptive Actions', *San Diego Law Review*, 55: 223-29.
- Arrow, Kenneth J. 1971. *Essays in the theory of risk-bearing* (North-Holland Pub. Co.: Amsterdam).
- Ashford, Elizabeth. 2003. 'The Demandingness of Scanlon's Contractualism', *Ethics*, 113: 273-302.
- Bassett, Gilbert W., Jr. 1987. 'The St. Petersburg Paradox and Bounded Utility', *History of Political Economy*, 19: 517-23.
- Bernoulli, Daniel. 1738. 'Exposition of a New Theory on the Measurement of Risk', *Econometrica*, 22: 23-36.
- Bjorndahl, Adam, Alex John London, and Kevin J. S. Zollman. 2017. 'Kantian Decision Making Under Uncertainty: Dignity, Price, and Consistency', *Philosophers' Imprint*, 17.
- Borel, Emile. 1962. *Probabilities and life* (Dover: New York).
- Bostrom, Nick. 2009. 'Pascal's mugging', *Analysis*, 69: 443-45.
- Buchak, Lara. 2013. *Risk and Rationality* (Oxford University Press).
- Buchanan, Allen. 1979. 'Revolutionary Motivation and Rationality', *Philosophy & Public Affairs*, 9: 59-82.
- Buffon, Georges-Louis Leclerc de. 1777. 'Essai d'Arithmétique Morale.' in, *Supplement a l'Histoire Naturelle* (L'Imprimerie: Paris).
- Chalmers, Adam. 2017. 'An offer you can't (rationally) refuse: systematically exploiting utility-maximisers with malicious gambles', University of Sydney.
- Colyvan, Mark, Damian Cox, and Katie Siobhan Steele. 2010. 'Modelling the Moral Dimension of Decisions', *Noûs*, 44: 503-29.
- Condorcet, Jean-Antoine-Nicolas de Caritat. 1785. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix* (Imprimerie royale: Paris).
- d'Alembert, Jean Le Rond. 1761. *Opuscules mathématiques, ou Mémoires sur différens sujets de géométrie, de mécanique, d'optique, d'astronomie &c* (David: Paris).
- Fried, Barbara H. 2012. 'What Does Matter? The Case for Killing the Trolley Problem (Or Letting It Die)', *The Philosophical Quarterly*, 62: 505-29.
- Hájek, Alan. 2014. 'Unexpected Expectations', *Mind*, 123: 533-67.
- Hansson, Sven Ove. 2013. *The ethics of risk: ethical analysis in an uncertain world* (Palgrave Macmillan).

- Haque, Adil Ahmed. 2012. 'Killing in the Fog of War', *Southern California Law Review*, 86: 63-116.
- Hawley, Patrick. 2008. 'Moral Absolutism Defended', *Journal of Philosophy*, 105: 273-75.
- Hayenhjelm, Madeleine, and Jonathan Wolff. 2012. 'The Moral Problem of Risk Impositions: A Survey of the Literature', *European Journal of Philosophy*, 20: e26-e51.
- Holm, Sune. 2016. 'A Right against Risk-Imposition and the Problem of Paralysis', *Ethical Theory and Moral Practice*, 19: 917-30.
- Huemer, Michael. 2010. 'Lexical Priority and the Problem of Risk', *Pacific Philosophical Quarterly*, 91: 332-51.
- Isaacs, Yoav. 2014. 'Duty and Knowledge', *Philosophical Perspectives*, 28: 95-110.
- . 2016. 'Probabilities Cannot Be Rationally Neglected', *Mind*, 125: 759-62.
- Jackson, Frank, and Michael Smith. 2006. 'Absolutist Moral Theories and Uncertainty', *Journal of Philosophy*, 103: 267-83.
- . 2016. 'The Implementation Problem for Deontology.' in, *Weighing Reasons* (Oxford University Press: New York).
- Jordan, Jeff. 1994. 'The St. Petersburg paradox and Pascal's Wager', *Philosophia*, 23: 207-22.
- Kagan, Shelly. 1989. *The Limits of Morality* (Oxford University Press).
- Kosonen, Petra. 2021. 'Discounting Small Probabilities Solves the Intrapersonal Addition Paradox', *Ethics*, 132: 204-17.
- Lazar, Seth. 2017. 'Risky Killing', *Journal of Moral Philosophy*, 16: 1-26.
- Lazar, Seth, and Chad Lee-Stronach. 2019. 'Axiological Absolutism and Risk', *Noûs*, 53: 97-113.
- Lee-Stronach, Chad. 2018. 'Moral priorities under risk', *Canadian Journal of Philosophy*, 48: 793-811.
- McClennen, Edward F. 1994. 'Pascal's Wager and Finite Decision Theory.' in Jeff Jordan (ed.), *Gambling on God: Essays on Pascal's Wager* (Rowman & Littlefield: Boston).
- McKerlie, Dennis. 1986. 'Rights and Risk', *Canadian Journal of Philosophy*, 16: 239-51.
- Miller, Frank, and Rolf Sartorius. 1979. 'Population policy and public goods', *Philosophy and Public Affairs*, 8: 148-74.
- Monton, Bradley. 2019. 'How to Avoid Maximizing Expected Utility', *Philosophers' Imprint*, 19: 1-25.
- Nebel, Jacob M. 2019. 'An Intrapersonal Addition Paradox', *Ethics*, 129: 309-43.
- Peterson, Martin. 2002. 'What Is a de Minimis Risk?', *Risk Management*, 4: 47-55.
- . 2017. *An Introduction to Decision Theory* (Cambridge University Press: Cambridge).

- Robert, David. 2018. 'Expected Comparative Utility Theory: A New Theory of Rational Choice', *Philosophical Forum*, 49: 19-37.
- Samuelson, Paul A. 1977. 'St. Petersburg Paradoxes: Defanged, Dissected, and Historically Described', *Journal of Economic Literature*, 15: 24-55.
- Schwitzgebel, Eric. 2017. '1% Skepticism', *Noûs*, 51: 271-90.
- Smith, Nicholas J. J. 2014. 'Is Evaluative Compositionality a Requirement of Rationality?', *Mind*, 123: 457-502.
- . 2016. 'Infinite Decisions and Rationally Negligible Probabilities', *Mind*, 125: 1199-212.
- Sobel, David. 2012. 'Backing Away from Libertarian Self-Ownership', *Ethics*, 123: 32-60.
- Tarsney, Christian. 2018. 'Moral Uncertainty for Deontologists', *Ethical Theory and Moral Practice*, 21: 505-20.
- Tenenbaum, Sergio. 2017. 'Action, Deontology, and Risk: Against the Multiplicative Model', *Ethics*, 127: 674-707.