



Decision Theory and De Minimis Risk

Martin Smith¹

Accepted: 4 October 2022
© The Author(s) 2022

Abstract

A *de minimis risk* is defined as a risk that is so small that it may be legitimately ignored when making a decision. While ignoring small risks is common in our day-to-day decision making, attempts to introduce the notion of a *de minimis risk* into the framework of decision theory have run up against a series of well-known difficulties. In this paper, I will develop an enriched decision theoretic framework that is capable of overcoming two major obstacles to the modelling of *de minimis risk*. The key move is to introduce, into decision theory, a non-probabilistic conception of risk known as *normic risk*.

1 Ignoring Small Risks

While cooking one evening I notice that I'm missing one of the ingredients I need, and have to decide whether to drive to the local shops to try and get it. In making this decision there are several things that I might consider; the time and effort involved in a trip to the shops, whether I can make do with what I have, the possibility that the shops are closed or don't have what I'm looking for etc. One thing I probably wouldn't consider—at least not seriously—is whether I might cause a fatal car accident on this short, familiar drive. It's not that I regard this as impossible—it's just that I would take it to be such a small risk as to be not worth factoring into the decision. In short, I would treat the risk of a fatal car accident as being *de minimis*. This term is taken from a well-known practice in risk management¹—a practice that employs a risk threshold to sort potential negative outcomes into those that need to be taken seriously and those that can be legitimately ignored. *De minimis risk management* is controversial, and is something I will return to—but, as the present example illustrates, the basic idea of ignoring small risks is one that is common in day-to-day decision making.

¹ Which, in turn, borrowed the term from the legal principle *de minimis non curat lex* or 'the law does not concern itself with trifles' (see Peterson, 2002, p. 47).

✉ Martin Smith
Martin.Smith@ed.ac.uk

¹ University of Edinburgh, Edinburgh, UK

One place that we don't find this idea however is in *expected utility theory*—the best known *decision theory*, or formal framework for navigating decisions under uncertainty. To formally model a decision problem we require, first, a pairwise exclusive and jointly exhaustive set of available *actions*. Second, for each of these actions, we require a pairwise exclusive and jointly exhaustive set of possible *outcomes*. Next we need a *risk* function which assigns some measure of risk to each outcome, given each action, and a *value* function which assigns some measure of value to each outcome. The final element is a *decision procedure* which selects an action or actions based on the risks and values. All of the decision theoretic frameworks that I consider here will appeal to these basic components.

Within standard expected utility theory (EUT), the risk function assigns *probabilities* to the outcomes that could result from an action—positive real numbers that sum to 1—while the value function assigns *utilities* to these outcomes—positive or negative numbers. The decision procedure selects the action or actions that have the highest *expected utility*, where the expected utility of an action is equal to the average of the utilities of its possible outcomes, weighted by the probabilities of those outcomes. Utilities, when negative, might be referred to as *disutilities* and, if it is potential negative outcomes that bulk large in a particular decision problem, we might say that the decision procedure selects the action or actions that have the lowest expected disutility. If A is the set of available actions, O is the set of possible outcomes, Pr is the probability function, and u is the utility function then, as is familiar, the expected utility of $A \in A$ can be written like this:

$$EU(A) = \sum_{O \in O} \text{Pr}(O | A) \cdot u(O) \text{ } ^2$$

Though it will play only a small role in what is to come, it may be worth delving a little deeper into the formal foundations of this framework. Following Jeffrey (1965), the A s and O s can be regarded as *propositions*: An action proposition is, in effect, a proposition that an agent can choose to make true in a given decision problem, while the truth or falsity of the outcome propositions lies beyond the agent's control, once an action has been chosen. Let propositions be modelled as subsets of a set of possible worlds W and let Ω be a Boolean σ -algebra of propositions—a set of propositions that contains W and is closed under negation and disjunction (and under countable disjunction in case Ω is infinite). The set of actions A and the set of outcomes O will both be partitions of W that are subsets of Ω . We assume here that outcomes are individuated in such a way as to include the action that gives rise

² The formula used here is characteristic of *evidential* decision theory. In *causal* decision theory the probability assigned to an outcome O , given an action A , is understood not as a conditional probability, but as a kind of 'causal' probability which, roughly speaking, reflects only A 's *causal influence* upon O , and brackets any purely *evidential* connection between A and O (see Gibbard and Harper, 1978, Lewis, 1981, Joyce, 1999, Egan, 2007, Buchak, 2016, Sect. 2). The dispute between evidential and causal decision theory (on which I mean to take no position here) is orthogonal to the issues under discussion—in the sense that none of the examples I consider would prompt a different treatment from the two approaches. There are substantial questions about how one might develop 'causal' versions of the various decision theoretic frameworks I will consider—but these will have to be postponed to another occasion.

to them, in which case \mathcal{O} will be a refinement or fine-graining of \mathcal{A} (every element of \mathcal{O} is a subset of an element of \mathcal{A}).³ \Pr can be defined as a function taking the members of Ω to real numbers in the unit interval and conforming to the probability axioms.⁴ The conditional probabilities required for calculating expected utilities can be derived according to the standard ratio formula: $\Pr(X \mid Y) = \Pr(X \wedge Y) / \Pr(Y)$, and is undefined in case $\Pr(Y) = 0$.

According to EUT, the action or actions that maximise expected utility will be rationally permissible for an agent, while the rest are rationally prohibited. In a case where only a single action maximises expected utility, this action will be rationally obliged. The ‘rationality’ here is, of course, relativised to the probability and utility functions, which can themselves be interpreted in different ways. Sometimes the interpretation is purely subjective— u reflects the agent’s desires while \Pr reflects the agent’s own probability judgments which, aside from general formal constraints, can be made in any way that the agent sees fit. For present purposes, I will adopt a more objective interpretation on which \Pr represents the probabilities imposed by the agent’s evidence, or the probability judgments that the agent’s evidence supports.

One thing we can immediately observe is that EUT has no mechanism whereby a low risk outcome can be ignored. On the contrary, any outcome that has some risk of resulting from an action, no matter how small, will feature in the expected utility calculation for that action and, depending on the utility or disutility it is assigned, could have a decisive effect upon the result. One might, at first, take the view that this is wholly appropriate. While ignoring low risk outcomes may be a convenient way of simplifying or expediting certain decisions, one might think that it has no place in a theory of idealised rational decision making. The idea, pursued a bit further, is that the ignoring of low risk outcomes might serve as a helpful *heuristic*—a way of reliably hitting upon a rational decision even when it would be too complex

³ A model of a decision problem will sometimes include, in addition, a set of *states* of the world, which are relevant to determining the outcomes of the available actions. Formally, the set of states \mathcal{S} will be another partition of W that is a subset of Ω , such that every outcome $O \in \mathcal{O}$ is associated with an action, state pair. The probability of an outcome O , given an action A will be equal to the sum of the probabilities, given A , of those states that would, in combination with A , lead to O . States, in effect, allow us to factor an outcome into a contribution that is due to the agent and a contribution that is due to the world. This additional structure allows us to define richer relations between actions—such as relations of dominance. I don’t include states in the formalism that I use in the main text, but it is easy enough to add them in, and I will mention them again in n9 and n19.

⁴ For $X, Y \in \Omega$:

P1	$\Pr(W) = 1$
P2	$\Pr(X) \geq 0$
P3	If X and Y are inconsistent then $\Pr(X \vee Y) = \Pr(X) + \Pr(Y)$

If we are dealing with an infinite stock of propositions we might strengthen P3 to:
 If $\Pi \subseteq \Omega$ is a set of pairwise inconsistent propositions then $\Pr(\vee \Pi) = \sum_{X \in \Pi} \Pr(X)$.

or laborious to try and consider everything—but when it comes to our theory of what *constitutes* a rational decision, we want every possible outcome to count.⁵

This may be a natural thought to have—and yet, the fact that every possible outcome ‘counts’ in EUT is implicated in at least two well-known problems that arise for the framework. The first problem stems from an attempt to model, within EUT, the relation of *lexical priority*. The second problem concerns the way in which the predictions of EUT can, in some cases, be manipulated by the provision of highly suspect information.

Say that one moral requirement is lexically prior to another just in case we are always morally obliged to uphold the first requirement at the expense of the second—no matter how many times the second must be violated thereby. Suppose we are offered a choice between taking the life of an innocent person and inflicting a short-lived, mild headache upon each of a number of people. Some would claim that the preservation of life takes lexical priority over the avoidance of headaches and, as a result, we should always choose the latter option, no matter how many headaches this involves. While this claim is disputable (Norcross, 1997, 1998), some regard it as obviously true (Thomson, 1990, p. 163) and it is sometimes described as a part of ‘common sense morality’ (see, for instance, Brennan, 2006, p. 251, Dorsey, 2009, pp. 36–37, Kirkpatrick, 2018, p. 107). But if an agent assigns a constant disutility to each headache, as seems reasonable, then the only way to capture this preference in EUT is to allow the agent to assign an *infinite disutility* to the death of an innocent (Jackson & Smith, 2006, Huemer, 2010, Colyvan et al., 2010, Sect. 3, Hansson, 2013).⁶

⁵ Some have explicitly endorsed the view that the ignoring of low risk outcomes could only benefit boundedly rational agents, and is something that an ideal reasoner should dispense with (see for instance Adler, 2007). If the points made here are on the right track, then this is not the right picture.

⁶ If we are prepared to give up the idea that every headache should add the same disutility, and assign the headaches a diminishing marginal disutility that tends towards 0, then it will be possible to capture this preference without the use of infinite disutilities. In this case, the disutility of n headaches would tend to a finite limit as n approached infinity and the death of an innocent could be assigned a disutility that exceeds this limit (Lazar and Lee-Stronach, 2019). Suppose, for instance, an agent assigns the first headache a disutility of 1, the second headache a disutility of $\frac{1}{2}$, the third headache a disutility of $\frac{1}{4}$ and so on. Provided the agent assigns the death of an innocent a disutility of 2 or more, this will exceed the disutility of any number of headaches. There are significant drawbacks, however, to a utility assignment like this. Most obviously, it requires the agent to place different value on the wellbeing of different people—to prefer relieving the first headache to relieving the tenth, even though they are equally painful etc. I won’t discuss this further here—but for more on the strategy of capturing lexical priority by placing finite bounds on the utility and disutility of certain benefits and costs see Black (2020).

Another option for modelling lexical priority is to allow the value of an outcome to be represented by a *multidimensional* utility—a sequence of positive or negative numbers. Multidimensional utilities are then subject to a lexicographic ordering—ordered with respect to the first dimension, then with respect to the second dimension in the case of any ties, then the third dimension in case of further ties and so on (much as words would be ordered alphabetically). In order for the death of an innocent to be reckoned worse than any number of headaches, it would be enough for it to be assigned disutility relative to an earlier dimension. I think that this approach to the modelling of lexical priority is, in many ways, more promising than the use of infinite disutilities, but it won’t, on its own, solve the problem identified in the main text. I discuss this further in Smith (ms)—see also Hájek (2003, Sect. 4.2) and Lee-Stronach (2018, Sect. 3).

The trouble with an infinite disutility, though, is that it *remains infinite* when multiplied by any positive probability whatsoever. If the death of an innocent is assigned an infinite disutility, then any action that has some risk of resulting in the death of an innocent will have negative infinity as one of the terms in its expected utility calculation. As a result, the other terms in the calculation (if finite) no longer matter—the action will have an infinite *expected* disutility. So if an agent would choose any number of headaches over the death of an innocent then, within EUT, the agent would also be committed to choosing any number of headaches over an action that merely *risks* the death of an innocent. But this commitment is certainly not a part of common sense morality—particularly if the risk involved is very low. If a friend is feeling a headache coming on then many of us would be perfectly willing to take a short drive to the pharmacy to buy some painkillers even though we are aware, at some level, that this action might result in an accident in which an innocent person loses their life.

One could of course reject the claim that the death of an innocent should be assigned an infinite disutility, and deny that the preservation of life takes lexical priority over the avoidance of headaches. But this would not altogether resolve the problem. What we have here are two intuitive preferences: I may be unwilling to sacrifice a human life in order to avoid any number of headaches, but I might also choose to drive to the pharmacy to help a friend avoid a single headache. Given the above assumptions, these preferences simply cannot be recovered within the standard EUT framework. If we think that these are the morally right preferences to have then the problem for EUT is obvious. But even if we think that these preferences are not morally right, the fact that they cannot even be *accommodated* within a given decision theoretic framework could still be seen as a serious limitation, and reason enough to seek greater flexibility.⁷

The second problem, in any case, doesn't involve devices such as infinite disutilities, and can be generated by a utility function that anyone should find acceptable. Suppose I am approached by a mysterious stranger who demands that I hand over my wallet. When I refuse he tells me that he has the authority to order a nuclear strike and, if I continue to refuse, he will do so. The resulting nuclear exchange, he informs me, would be so devastating as to cost billions of human lives. While I don't for a moment believe what the stranger is claiming, I can't conclusively rule it out. There may be some people on Earth who have the capability to trigger a nuclear exchange and, while it seems incredibly farfetched that the stranger is one of them and is sincere in his threat, it's not impossible. Even if the probability that the stranger is speaking the truth is only one-in-a-billion, the extreme disutility of the scenario he describes means that this could still have a significant effect on the expected utility of keeping my wallet. If there is a one-in-a-billion chance that

⁷ There are some who spin these observations in a rather different way. For some, the difficulties involved in modelling lexical priority within the standard EUT framework give us reason to reject the idea of lexical priority and the kinds of preferences that it demands (see for instance Jackson and Smith, 2006, 2016, Huemer, 2010). The contrary idea that these difficulties expose the limitations of our decision theoretic modelling is pursued by Lee-Stronach (2018). I discuss this issue in more detail in Smith (ms).

keeping my wallet will result in several billion deaths, and every death is assigned a constant disutility, then this option will have an expected disutility equal to the disutility of several deaths, which would surely exceed the expected disutility of losing my wallet. Given these utilities and probabilities, EUT predicts that I am rationally obliged to hand my wallet over.⁸

One might complain that the probability of the stranger's claim shouldn't be as high as one-in-a-billion—perhaps one-in-ten-billion or one-in-a-hundred-billion would be more realistic. But even if that's right, any reduction in the probability of the stranger's claim could be compensated by an increase in the disutility assigned to the prospect of global nuclear war. The disutility assumed above only takes account of the loss of human life and doesn't, for instance, factor in the loss of animal and plant life, the suffering that would be experienced by the survivors, the countless future humans that would be prevented from coming into existence and so on. Whatever the rationale, I could legitimately assign this scenario a disutility that is trillions of times greater than that of losing my wallet—but surely this wouldn't commit me to handing my wallet to the stranger. Surely it is rationally permissible for me to hold on to my wallet *irrespective* of the disutility I assign to the scenario that the stranger describes.

Both of these problems revolve around an expected utility calculation that is, in effect, hijacked by a very low risk outcome—an outcome of the kind that we would ordinarily ignore. When it comes to the decisions described, considering every possible outcome is not just difficult or cumbersome—it seems to force us into the *wrong* decision or, at any rate, a decision that isn't rationally mandated. One way to make progress with these problems is to introduce, into decision theory, the notion of a *de minimis* risk—which might be defined as an outcome that is at such low risk of resulting from an action that we may ignore it for the purposes of decision making. In this paper I will outline two proposals as to how this might be done. The first proposal has been put forward, in slightly different forms, by a number of authors. This proposal offers a straightforward solution to the lexical priority and suspect information problems, but runs into two problems of its own—the disjunction problem and the lottery problem. The second proposal is more radical—in that it involves introducing a completely new component into the decision theoretic framework—but is able to avoid all four of these problems.

Some would argue that all of this fuss seems premature; to reject EUT on the strength of the lexical priority and suspect information problems is an overreaction. I don't altogether disagree. It's true that the problems are far from decisive, and there are a number of things that defenders of EUT could potentially say in response. But, on the other side, we shouldn't feel beholden to any particular formal framework,

⁸ This example is based on the problem of 'Pascal's mugging' described by Bostrom (2009) (see also <http://www.lesswrong.com/posts/a5JAiTdytjou3Jg749/pascal-s-mugging-tiny-probabilities-of-vast-utilities>). In Bostrom's original case the mugger makes extravagant promises in the event that the wallet is handed over rather than extravagant threats in the event that it isn't—but the effect is the same. The case described here is closer to the version of Pascal's mugging given by Balfour (2021). The idea that these cases present a problem for EUT is (I think) implicit in Bostrom's paper, but is developed in detail by Monton (2019).

particularly if there are alternatives available. The middle ground, I think, is to treat these problems as an incentive to develop alternatives to the standard EUT framework and see what, if anything, they have to offer. The two alternatives that I develop here are very much presented in this spirit.

2 De Minimis Expected Utility Theory

On one level, introducing de minimis risk into decision theory is straightforward—after all, even if an outcome has some positive probability of resulting from an action, it could still, in principle, be ignored by a decision procedure. Suppose we specify a probability value t greater than, but close to, 0 to serve as the de minimis threshold. One natural option is to consider a decision procedure in which expected utilities are calculated while *excluding* any de minimis risks—any outcomes with a probability below t . Rather than using Pr to assign the weights in an expected utility calculation for an action A , we now use an amended function, which results from conditionalising Pr on the conjunction of the negations of all de minimis risks—of all outcomes which, given A , have a probability below t . That is, if $\varphi = \wedge\{\sim O \mid O \in \mathcal{O} \wedge \text{Pr}(O \mid A) < t\}$ we have it that:

$$DEU(A) = \sum_{O \in \mathcal{O}} \text{Pr}_{\varphi}(O \mid A) \cdot u(O) = \sum_{O \in \mathcal{O}} \text{Pr}(O \mid A \wedge \varphi) \cdot u(O)^9$$

The decision procedure selects the action or actions that have the highest *de minimis expected utility*, as given by this formula. We might call this framework *de minimis expected utility theory* (DEUT).

Less formally, we might think of the calculation of a de minimis expected utility as a two-step process: First, we list all of the outcomes that could possibly result from an action and eliminate any outcomes that have a probability below t . Second, we update the probabilities of those outcomes that remain and use these probabilities to calculate an expected utility in the usual way. According to DEUT, the action

⁹ Since φ entails A , the formula could be written more simply:

$$DEU(A) = \sum_{O \in \mathcal{O}} \text{Pr}_{\varphi}(O) \cdot u(O) = \sum_{O \in \mathcal{O}} \text{Pr}(O \mid \varphi) \cdot u(O)$$

The way the formula is presented in the main text is more perspicuous, however, and has a closer correspondence with the familiar formula for calculating expected utility.

Pr_{φ} could also, in principle, be defined in a way that is not relative to a given action. Recall that, in the present framework, every outcome $O \in \mathcal{O}$ includes a unique action. If we let α be a function mapping each outcome to its associated action, φ could then be defined as $\wedge\{\sim O \mid O \in \mathcal{O} \wedge \text{Pr}(O \mid \alpha(O)) < t\}$. With φ defined in this way, we could not use the above simplification but, as can be easily checked, the outputs of the original formula will be unchanged for any action $A \in \mathcal{A}$. Another option is to define φ in terms of *states* (if we include them in the model)— $\varphi = \wedge\{\sim S \mid S \in \mathcal{S} \wedge \text{Pr}(S) < t\}$. This will lead to the same results as the definition in the main text provided we make the following assumptions: (i) Outcomes are associated with a unique state as well as a unique action—that is, \mathcal{O} is a more fine-grained partition than both \mathcal{A} and \mathcal{S} . (ii) States and actions are probabilistically independent—for any $S \in \mathcal{S}$ and $A \in \mathcal{A}$, $\text{Pr}(S \mid A) = \text{Pr}(S)$.

or actions that maximise de minimis expected utility, calculated in this way, will be rationally permissible for an agent. I won't assume here that these are the *only* actions that are rationally permissible. A de minimis risk was defined as an outcome that an agent *may* ignore when making a decision—but that's not to say that it *has to be* ignored. As a result, I leave it open whether the actions which maximise standard expected utility (if different) might also be regarded as rationally permissible.

Even if I assign an infinite disutility to the loss of a life and a finite disutility to a mild headache, within DEUT I may still be rationally permitted to drive to the pharmacy to buy painkillers for my friend. Although there is some probability that my driving to the pharmacy will result in a death, provided this is below the de minimis threshold, the action can still have a positive de minimis expected utility—and one that exceeds the de minimis expected utility of letting my friend suffer a headache.¹⁰ Even if I assign the prospect of global nuclear war a disutility that is trillions of times greater than that of losing my wallet, within DEUT I may still be rationally permitted to hold on to my wallet, in spite of the stranger's claim. Provided the probability of the stranger's claim is below the de minimis threshold, it will have no effect on the de minimis expected utility of holding on to my wallet which, all else equal, will exceed the de minimis expected utility of handing it over.¹¹ If we want the stronger result that I am rationally *obliged* to hold on to my wallet, we will need the stronger assumption that the *only* actions that are rationally permissible for an agent are those that maximise de minimis expected utility—and, as noted above, I don't commit to this here.¹²

One question that immediately arises for DEUT is that of where the de minimis threshold should be set. One place we might look for inspiration is the literature on de minimis risk management, where the problem of setting a de minimis threshold

¹⁰ This is close to a suggestion made by Lee-Stronach in the context of an attempt to capture lexical priority within decision theory (Lee-Stronach, 2018, Sect. 5). Rather than amending the decision procedure, however, Lee-Stronach suggests that any possible outcomes with a probability lower than the de minimis threshold could simply be left out of a model of a decision problem. While this may, in one sense, lead to the same results, I am inclined to think that there are some theoretical advantages to having the discounting of possible outcomes represented in the decision procedure itself, rather than being something that is concealed within the construction of a decision model. On Lee-Stronach's approach, the probability function Pr can no longer be interpreted as representing an agent's evidential probabilities, since there may be outcomes with a positive evidential probability that are nevertheless missing from the model. Pr should presumably be interpreted along the lines of Pr_{ϕ} above—an evidential probability function conditionalised upon the negations of each outcome that represents a de minimis risk. And yet, since neither the de minimis threshold, nor the discounted outcomes, are represented anywhere in the model, the precise meaning of Pr is still left opaque. Relatedly, a decision model will not, on this approach, contain sufficient information to construct a new model in the event that the de minimis threshold is adjusted or new evidence is acquired. I discuss this further in Smith (ms).

¹¹ This is close to a suggestion made by Monton as a way of dealing with Pascal's mugging scenarios (Monton, 2019, Sect. 4). Monton does not develop the suggestion formally—rather, he sees himself as providing an 'impetus' for a better decision theory (Monton, 2019, Sect. 8). His views are I think consistent with a development along the lines of DEUT, but could also be developed in alternative ways. Others who have suggested the discounting of low probability outcomes include Smith (2014) and Buchak (2014, pp. 73–74) and Kosonen (2021)—though their motivations are different from those that I focus on here. The idea is also present in much earlier sources, such as Cournot (1843).

¹² The stronger assumption is endorsed by Monton (see Monton, 2019, p. 10).

has received a good deal of attention. De minimis risk management, as mentioned in the introduction, involves sorting the potential adverse consequences of a policy decision into two categories. Those that have a risk below the de minimis threshold are dismissed, while those with a risk above the threshold are subjected to a full risk analysis to determine if they can be further controlled or mitigated.

Some advocates of de minimis risk management propose a precise numerical threshold that might be used in any context (see for instance Comar, 1979; Mumpower, 1986)—an approach that immediately attracts charges of arbitrariness. Others propose methods for deriving a threshold from the parameters of a given decision. One prominent method is to let the threshold be guided by the level of equivalent background risk to which we are all exposed in our daily lives (see for instance Weinberg, 1985, section VII). Suppose we are considering a policy that will affect a given population, such as the licensing of a new food additive. If we consider the worst outcome that could befall individuals in the population as a result of exposure to the additive then, on the present approach, the de minimis threshold should be determined in relation to the background risk of comparable outcomes—the risk to which individuals are already subject as a result of ordinary everyday activities. Some have complained that this approach is still unacceptably arbitrary, however, as background levels of risk depend upon a range of contingent factors, can vary considerably from one time to another and one place to another, and are not always ignored (Peterson, 2002, pp. 52–53, Lundgren and Stefánsson, 2020, Sect. 2).

Concerns about fixing the de minimis threshold have led some to conclude that the very practice of de minimis risk management is wrongheaded (Lundgren and Stefánsson, 2020). But it would be hasty, I think, to draw the same conclusion about a framework like DEUT. Unlike a set of guidelines or instructions for managing risks, DEUT doesn't have to specify a de minimis threshold, or even supply determinate criteria for finding one. Rather, the de minimis threshold can simply be seen as another input into the decision procedure—another variable that gives the framework greater flexibility when it comes to capturing preferences. Indeed, the preferences involved in the lexical priority and suspect information problems already place constraints upon where the de minimis threshold might be set—and this suggests a methodology on which our judgments about the rationality of certain preferences can serve as a guide to establishing the threshold (see Monton, 2019, Sect. 6.1).¹³ In any case, I won't pursue this further here, as I think that DEUT is subject to at least two further problems—problems which show, in my view, that the framework is not viable as it stands.

¹³ It is often observed that de minimis risk management cannot be squared with standard EUT (see for instance Peterson, 2002, p. 49, Adler, 2007). While this practice would undoubtedly fit better with a framework like DEUT, that's not to say that accepting DEUT would commit us to adopting the practice. As explained in the main text, DEUT may offer the flexibility to capture preferences and practices that deviate from the recommendations of EUT, but it doesn't, in and of itself, force us to endorse them.

3 The Disjunction and Lottery Problems

A disjunction is, in general, more probable than its disjuncts. If we set a de minimis threshold t , the probability that an outcome O_1 will result from action A and the probability that an outcome O_2 will result from action A may both be below t , even though the probability that either O_1 or O_2 will result from A may be above t . In this case it looks as though O_1 and O_2 can be legitimately ignored when deciding about A —as these are both de minimis risks. And yet, if we ignore O_1 and we ignore O_2 then this is tantamount to ignoring $O_1 \vee O_2$ which is *not* a de minimis risk and, thus, is the sort of thing that we are supposed to *consider* when making a decision. As a result, the very idea of a de minimis risk seems to lead to inconsistent recommendations as to what we are allowed to ignore (Ebert et al., 2020, pp. 438–440, for related discussion see Lundgren & Stefánsson, 2020, Sect. 4).

In a way, adopting a particular decision model resolves this dilemma, in that it forces us to specify the possible outcomes in one particular way. But we now face the prospect that the rational permissibility of an action may hinge upon how finely the outcomes are individuated in the model. Suppose again that the probability that O_1 will result from A and the probability that O_2 will result from A are both below t , while the probability that either O_1 or O_2 will result from A is above t . If our model includes O_1 and O_2 as separate outcomes then, according to DEUT, they can be legitimately ignored. If, on the other hand, we combine O_1 and O_2 into a single outcome— $O_{1 \vee 2}$ —then, according to DEUT, this outcome must be considered. This could make all the difference as to whether or not A is predicted to be permissible. Following Lee-Stronach, we might call this the *disjunction problem* (Lee-Stronach, 2018, Sect. 5, see also Monton, 2019, n26, Kosonen, ms, Sect. 1).

Suppose I'm considering whether to drill into the wall to hang a painting, and I'm concerned about the possibility that the wall contains asbestos. Suppose I find out that, of all the houses built in the same area and around the same time as mine, one in 10,000 has asbestos in the walls, and that half of these contain blue asbestos and half contain brown asbestos. If I drill into the wall then, given my evidence, there is a one in 20,000 chance that this will release blue asbestos dust (O_1) and a one in 20,000 chance that this will release brown asbestos dust (O_2)—both of which are equally harmful to human health. Suppose we set the de minimis threshold at one in 15,000. If, in modelling this decision, we decide to distinguish these two outcomes then, since they both have a probability below the de minimis threshold, they will be discounted for the purposes of calculating a de minimis expected utility. If these are the only negative outcomes that could result from my drilling into the wall then DEUT will predict that it is rationally permissible for me to proceed. If, on the other hand, we roll these together into a single outcome—in which asbestos dust of *some kind* is released ($O_{1 \vee 2}$)—then this will have a probability above the de minimis threshold and will be included in the calculation of de minimis expected utility. If the disutility that I assign this outcome is sufficiently high, and eclipses the utility of hanging the painting, then DEUT will predict that it is not rationally permissible for me to proceed.

What we have is a fine-grained model, which includes O_1 and O_2 as outcomes, and which predicts that it is permissible for me to drill into the wall, and a coarse-grained model, which includes $O_1 \vee_2$ as an outcome, and which predicts that it is impermissible for me to drill into the wall. EUT doesn't face this problem. In the standard EUT framework, if the release of asbestos dust is assigned a sufficiently high disutility, both the fine-grained and the coarse-grained models will predict that drilling into the wall is impermissible. More generally, the predictions of EUT are guaranteed to be *invariant* under fine-graining.¹⁴

One way to approach the disjunction problem is to argue that certain ways of modelling a decision are to be strictly preferred over others. A fine-grained model will, in effect, be more detailed and will include more information than a coarse-grained one. A model in which the outcomes are differentiated more finely comes closer to Savage's notion of a 'grand world' decision model in which every potentially relevant detail is included (see Savage, 1954, chap. 5). Thus, it's plausible that the predictions of a fine-grained model ought to take precedence over those of a coarse-grained model, in case those predictions conflict (Joyce, 1999, Sect. 2.6, Thoma, 2019, section V). If this thought is on the right track then, in the above example, we should favour the first model, in which drilling into the wall is predicted to be permissible. Whatever one thinks of this prediction, we would at least appear to have a definitive verdict about the case.

On closer inspection, though, matters are not so clear. Aside from asbestos, there are any number of other materials that could be released if I drill into the wall—plaster, wood, concrete etc. I may be indifferent to these possibilities—but they could still be represented in a decision model. If we took note of the precise quantity and chemical composition of the materials that could be released then we could, in principle, distinguish millions of possible outcomes that could result from my drilling into the wall—each of which has a probability below the de minimis threshold. If we do this, however, then we will be *unable* to calculate a de minimis expected utility for the action. That is, if O_1, O_2, \dots, O_n are all the possible outcomes which could result from an action A , such that $\Pr(O_1 | A) \approx \Pr(O_2 | A) \approx \dots \approx \Pr(O_n | A) < t$, then $\varphi = \sim O_1 \wedge \sim O_2 \wedge \dots \wedge \sim O_n = \sim A$ and $\sum_{O \in \mathcal{O}} \Pr_\varphi(O | A) \cdot u(O) = \sum_{O \in \mathcal{O}} \Pr(O | A \wedge \sim A) \cdot u(O)$, which is *undefined*. Rather than guiding us towards a definitive verdict about an action, the policy of differentiating outcomes as finely as possible will, in a case like this, cause the DEUT framework to collapse altogether.¹⁵

¹⁴ That is, substituting a coarse-grained outcome $O_{1 \vee \dots \vee n}$ for a series of more fine-grained outcomes O_1, \dots, O_n will make no difference to the standard expected utility of any action, provided the utility of $O_{1 \vee \dots \vee n}$ in the coarse grained model is equal to the probability weighted average of the utilities of O_1, \dots, O_n in the fine-grained model— $u(O_{1 \vee \dots \vee n}) = \sum_{1 \leq x \leq n} \Pr(O_x | O_1 \vee \dots \vee O_n) \cdot u(O_x)$, where u is the utility function for the coarse-grained model and u' is the utility function for the fine-grained model. The proof is provided in the appendix. As the above example illustrates, even if we observe this constraint, substituting $O_{1 \vee \dots \vee n}$ and O_1, \dots, O_n can still make a difference to the de minimis expected utility of an action.

¹⁵ One way to avoid this result is to use what Lee-Stronach calls an 'odds-based' threshold, rather than an absolute probability threshold for the identification of de minimis risks (Lee-Stronach, 2018, p. 801). On this approach, an outcome is deemed a de minimis risk if and only if its probability is low *relative* to some other outcome in the model—that is, lower than some set fraction of the probability of another outcome. Moving to an odds-based threshold makes it impossible for all outcomes to be discounted, and

While the disjunction problem draws attention to a kind of structural defect within DEUT, the second problem involves a straightforward counterexample: Consider a random process that could result in the death of an innocent. Suppose a single marked ticket is spun in a barrel along with a number of blank tickets, before a ticket is randomly drawn. If the result is a blank ticket then no action is taken, but if the marked ticket is drawn then an innocent person is put to death. Or suppose, more vividly, that a single bullet is placed in a single chamber of one out of a set of revolvers before a revolver is selected at random, the cylinder is spun, it is pointed at an innocent person and the trigger is pulled. Wherever we set the *de minimis* threshold, if the number of tickets or the number of revolvers is large enough, DEUT will predict that the outcome in which the person dies counts as a *de minimis* risk, which can be legitimately ignored. But even if we accept the general idea that there are such things as *de minimis* risks, there is something deeply objectionable about the suggestion that an innocent dying in a ‘death lottery’ of these kinds could ever count as one. If an agent was seriously considering subjecting a person to one of these set-ups, they would be *obliged* to consider the possibility that the person will die, irrespective of how low its *probability* might be.

We might term this the ‘lottery problem’—in death lottery cases DEUT seems to make incorrect predictions about which outcomes count as *de minimis* risks. And this will lead to incorrect predictions about our rational obligations. As noted in Sect. 2, in most circumstances I would be perfectly willing to drive to the pharmacy to buy painkillers to prevent a friend’s headache, even though there is some low risk that this action could result in the death of an innocent. But I would *not* be willing to subject an innocent person to a death lottery in order to buy painkillers, no matter the number of tickets or revolvers involved.¹⁶ If the number is sufficiently large, however, then the probability of an innocent dying in the lottery may be lower than that incurred by a drive to the pharmacy. In this case, if the death of an innocent counts as a *de minimis* risk relative to the pharmacy option, it must also count as a *de minimis* risk relative to the lottery option, and the two would have the same *de minimis* expected utility. Further, it would not take much to tip the balance in favour of the death lottery. If we added a small amount of additional utility to this option—say the painkillers come cheaper than they would at the pharmacy—then DEUT will predict that I am rationally obliged to choose this over the pharmacy option, as it

Footnote 15 (continued)

ensures that the *de minimis* expected utility of an action is always defined. The use of an odds-based threshold introduces a new problem, however—while a policy of finely individuating outcomes will no longer cause DEUT to break down, there is a danger that it will render the *de minimis* aspect of the framework *inert*. If O_1, O_2, \dots, O_n are the possible outcomes which could result from an action A, such that $\Pr(O_1 | A) \approx \Pr(O_2 | A) \approx \dots \approx \Pr(O_n | A)$ then, using an odds-based threshold, *no* outcome will count as a *de minimis* risk and the *de minimis* expected utility and the standard expected utility of A will coincide. I won’t pursue this further here.

¹⁶ There is a strong intuition to the effect that the former action is morally permissible while the latter is morally prohibited—a point made clearly by Thomson (1983, Sect. 4) and anticipated by Nozick (1974, pp. 73–84). For more recent discussion of this contrast see Hayenhjelm and Wolff (2012, Sect. 7), Holm (2016, Sect. 3).

would have both a higher de minimis expected utility and a higher standard expected utility.

4 Normic Risk

There is nothing in either EUT or DEUT that can make sense of the preferences just described; I would be willing to drive to the pharmacy to buy painkillers to prevent a friend's headache, but I would not be willing to subject an innocent person to a death lottery for the same result, irrespective of the number of tickets involved. That's not to say, however, that such preferences are baseless. The set-up of a death lottery invites the thought that *some* ticket has to be drawn, and it could just as easily be the marked ticket as any other—or the bullet has to be in *some* chamber, and it might just as well be the chamber that rotates into alignment with the barrel when the trigger is pulled. When it comes to driving to the pharmacy, this kind of reasoning doesn't seem to apply. Indeed, when evaluating this option, it would be natural to have the diametrically opposed thought that the only way in which an innocent could die is if something goes drastically *wrong*—the brakes fail, or I lose consciousness, or someone runs right in front of the car etc. Under *normal* circumstances, a drive to the pharmacy is not going to result in a death. But this seems not to be true of a death lottery, in which all of the outcomes would be equally normal. These factors are not represented in a standard decision model—and yet they could well have some effect on our decision making.

According to the conventional understanding of risk, the risk of a possible outcome is determined by its *probability*. More precisely, the risk that a particular outcome would result from an action, given the agent's evidence, depends upon how probable it is that the outcome would result from the action, given the agent's evidence. This account of risk has been taken for granted across a range of areas, including decision theory. Recently, however, several authors have challenged the hegemony of this account, and put forward alternatives, such as the *modal* account (Pritchard, 2015), the *relevant alternatives* account (Gardiner, 2021) and the *normic* account (Ebert et al., 2020)—which will be my focus here. According to the normic account, the risk of a possible outcome is determined by its *abnormality*. More precisely, the risk that a particular outcome would result from an action, given the agent's evidence, depends upon how abnormal it would be for the outcome to result from the action, given the agent's evidence.

But what is the notion of abnormality that is being invoked here? Sometimes when we describe an event as 'abnormal' we mean only that it's *infrequent*. Suppose you are playing poker, holding a three-of-a-kind and, looking at your opponent, I whisper 'she wouldn't normally have a better hand than that'. This may just be another way of saying that it would be relatively *rare* for your opponent to be dealt a better hand. On other occasions, however, when we describe a possibility as abnormal, we are not just making a claim about frequencies. Suppose you're trying to decide whether to take the bus home and I remark 'the bus ride wouldn't normally take more than 20 minutes'. Part of what I'm saying here is that circumstances would have to *conspire against you* in some way in order for the ride to take

more than 20 minutes—it would have to be that the bus breaks down, or runs out of petrol, or gets stuck in traffic, or is diverted by roadworks etc. but, absent any of these interfering factors, the trip would take 20 minutes or shorter. Put differently, if you get on the bus, and the trip ends up taking longer than 20 minutes, there would have to be some *special explanation* as to how this happened. In contrast, if your opponent happened to be dealt a hand that beats your three-of-a-kind then, while you may think yourself unlucky, no special explanation is needed for this. When we say that a given outcome would be abnormal, what we are sometimes claiming is that there would have to be some special explanation if it were to result from the action in question (Smith, 2010, 2016, chap. 2). It is this notion of abnormality that is appealed to in the normic account of risk.

While it may be highly unlikely for the one marked ticket to be drawn, this is not something that would require a special explanation of any kind—as noted above, the marked ticket could be drawn just as easily as any other ticket. If, on the other hand, an innocent person were to die as a result of my driving to the shops, there would have to be some explanation for this. Possible explanations have already been suggested above—perhaps the brakes failed, or I lost consciousness, or someone ran right in front of the car etc. Whatever the case, this could only happen as a result of some serious disruption to the normal course of events. While a death is a possible outcome of a drive to the shops, it is a highly *abnormal* outcome. But no disruption to the normal course of events is required for the marked ticket to be drawn. Although a death lottery could, in principle, present a lower probabilistic risk of death than a drive to the shops, it will present a higher *normic risk*, and will do so irrespective of the number of tickets involved.

It's clear that the normic risk of various possible outcomes is something that could potentially be factored into a decision. In order for this notion to play a role in a formal theory of decision making, however, it will need to be made more formally tractable. Suppose that propositions can be placed in an abnormality ordering, reflecting how much explanation their truth would require, given an agent's evidence (see Smith, 2022, Sect. 5). Suppose that any two propositions can be compared for their normalcy—that is, suppose that, for any two propositions, either one is more normal than the other or both are equally normal. Given these assumptions, propositions may be assigned numerical abnormality degrees—the maximally normal propositions will be assigned an abnormality degree of 0, the next most normal propositions will be assigned an abnormality degree of 1 and so on.¹⁷

No proposition can be more normal than the universal proposition *W* which should always be assigned degree 0. Similarly, no proposition can be less normal

¹⁷ If we are dealing with an infinite stock of propositions, then one further assumption is needed; that every (potentially infinite) set of propositions has maximally normal members (the existing assumptions already guarantee this for finite sets).

than the empty proposition \emptyset which should always be assigned an infinite degree of abnormality. The final constraint to be placed upon an abnormality ordering concerns disjunctions: The only way in which $X \vee Y$ can be true is if either X is true or Y is true. To explain the truth of $X \vee Y$ is to either explain the truth of X or the truth of Y . As a result, the amount of explanation demanded by the truth of $X \vee Y$ will be equal to either the amount demanded by X or by Y —whichever is less. And the degree of abnormality of $X \vee Y$ will be equal to the degree of abnormality of X or of Y —whichever is lower (see Smith, 2022, Sect. 5).

If, as above, Ω is a Boolean σ -algebra of propositions built from the elements of W , we can define an evidential abnormality function ab taking propositions in Ω into the set of nonnegative integers plus infinity— $\{0, 1, 2, 3, 4 \dots \infty\}$. Given the above constraints, ab will satisfy the axioms for a negative ranking function (see, for instance, Huber, 2009, Sect. 4, Spohn, 2012, chap. 5, Smith, 2016, p. 169). For $X, Y \in \Omega$:

R1	$ab(W) = 0$
R2	$ab(\emptyset) = \infty$
R3	$ab(X \vee Y) = \min\{ab(X), ab(Y)\}$ ¹⁸

Conditional degrees of abnormality can be calculated according to the following formula: $ab(X \mid Y) = ab(X \wedge Y) - ab(Y)$ (where $\infty - \infty = 0$)—which corresponds to the standard definition of conditional negative ranks (Huber, 2009, p. 19, Spohn, 2012, Sect. 5.3). In this case, the abnormality of X given Y is equal to the abnormality that X adds to the existing abnormality of Y —the amount of additional explanation that $X \wedge Y$ requires over and above that required by Y . This is equivalent to saying that the abnormality of Y and the abnormality of X given Y will together add up to the abnormality of $X \wedge Y$. On the normic account of risk, the risk that an outcome O will result from an action A can be gauged by the conditional abnormality of O given A —by $ab(O \mid A)$. The lower the value of $ab(O \mid A)$, the greater the risk that O will result from A , with the risk being maximal when $ab(O \mid A) = 0$ —indicating that O represents one of A 's normal outcomes.

Before proceeding further, it is important to emphasise that I am not proposing the normic account of risk as a *competitor* to the probabilistic account. Rather, I take the view that both of these accounts capture prominent patterns in ordinary risk judgments, and that each may be regarded as a legitimate way of refining our ordinary risk concept (Ebert et al., 2020, Sect. 6). When comparing the pharmacy option and the lottery option, on the present 'pluralist' approach there is no final fact of the matter as to which of these presents a higher risk that an innocent will die; the pharmacy option presents a higher probabilistic risk, while the lottery option presents a

¹⁸ If, once again, we dealing with an infinite stock of propositions, we might strengthen R3 to: $ab(V \mid \Pi) = \min\{ab(X) \mid X \in \Pi\}$ for any $\Pi \subseteq \Omega$. According to this principle, the degree of abnormality of the disjunction of a (potentially infinite) set of propositions will be equal to the degree of abnormality of its most normal members.

higher normic risk. If forced to choose between these options, it might seem that we are forced to choose also between these two kinds of risk—to choose whether probabilistic or normic risk will serve as our guide. In the next section, though, I will outline a picture on which *both* kinds of risk can play a role in determining which actions are rationally permissible.

5 Normic De Minimis Expected Utility Theory

There would be little purpose in introducing normic risk into EUT. The notion of risk, in effect, plays only a single role in this framework—namely, determining the weights that are assigned to different outcomes for the purpose of calculating an expected utility—and probabilistic risk may be uniquely suited to this. In DEUT, however, probabilistic risk is called upon to play two *distinct* roles; to determine which outcomes may be excluded from consideration and to determine the weights of those outcomes that remain. The fundamental idea behind the final decision theoretic framework that I will describe is that these two roles should be *decoupled*—while probabilistic risk continues to determine how each of the outcomes in an expected utility calculation is weighted, it is normic risk that determines which outcomes have to be included in the calculation in the first place.

Suppose we specify some abnormality rank t to serve as the de minimis threshold, and consider a decision procedure in which expected utilities are calculated while excluding those outcomes that have an abnormality greater than t . Rather than using Pr to assign the weights in an expected utility calculation for an action A , we use an amended function, which results from conditionalising Pr on the conjunction of the negations of all outcomes which, given A , have an abnormality greater than t . That is, if $\psi = \wedge \{ \sim O \mid O \in \mathbf{O} \wedge \text{ab}(O \mid A) > t \}$ we have it that:

$$\text{NDEU}(A) = \sum_{O \in \mathbf{O}} \text{Pr}_{\psi}(O \mid A) \cdot u(O) = \sum_{O \in \mathbf{O}} \text{Pr}(O \mid A \wedge \psi) \cdot u(O) \quad 19$$

The decision procedure selects the action or actions that have the highest *normic de minimis expected utility*, as given by this formula. We might call this framework *normic de minimis expected utility theory* (NDEUT).

NDEUT continues to embed the idea that if an outcome has a sufficiently low risk of resulting from a given action, then that outcome can be legitimately ignored—but the ‘risk’ is now interpreted as *normic*, rather than probabilistic. According to NDEUT, the action or actions that maximise normic de minimis expected utility, calculated in this way, will be rationally permissible for an agent. As with DEUT, I

¹⁹ An alternative non-action-relative definition of ψ could be given by $\wedge \{ \sim O \mid O \in \mathbf{O} \wedge \text{ab}(O \mid \alpha(O)) > t \}$, where α , as introduced in n9, is a function mapping each outcome to its unique associated action. As with DEUT, we could, once again, apply the de minimis threshold directly to states, defining ψ as $\wedge \{ \sim S \mid S \in \mathbf{S} \wedge \text{ab}(S) > t \}$. This will lead to the same results on the assumption that every outcome is associated with a unique state, and states and actions are *normically* independent—for any $S \in \mathbf{S}$ and $A \in \mathbf{A}$, $\text{ab}(S \mid A) = \text{ab}(S)$.

leave it open whether actions that maximise standard expected utility might also be regarded as rationally permissible.

NDEUT, like DEUT, offers a straightforward solution to the problems of lexical priority and suspect information. Not only would it be highly improbable for my drive to the shops to cause a fatal car accident or for the stranger to follow through on his threat, these outcomes would also be highly abnormal. That is, these outcomes should be regarded as de minimis risks on either a probabilistic or a normic construal. Given, then, an appropriate choice of threshold, driving to the shops and holding on to my wallet are options that could both have a positive normic de minimis expected utility—and a normic de minimis expected utility which exceeds that of handing my wallet over or of letting my friend suffer a headache.

NDEUT also offers a solution to the lottery problem. Unlike DEUT, NDEUT predicts that the outcome of a random lottery can never qualify as a de minimis risk, no matter how many tickets are involved. As I argued in the previous section, all of the outcomes in a random lottery will be at maximal normic risk, as none would require special explanation. Ultimately, though, we don't need to rely here on this judgment—or any particular judgment about what does and doesn't require explanation. As long as it is accepted that all lottery outcomes are equally normal, and one of them is guaranteed to eventuate, then the *formal* properties of normic risk will ensure that the normic risk of each outcome is maximal. Suppose O_1, \dots, O_n are the possible outcomes of the lottery—one for each ticket in the barrel. Since the lottery must have some outcome, we have it that $ab(O_1 \vee \dots \vee O_n) = 0$. If the outcomes are all equally normal then $ab(O_1) = \dots = ab(O_n)$. It now follows, using R3, that all outcomes are maximally normal— $ab(O_1) = \dots = ab(O_n) = 0$. Within NDEUT, no lottery outcome—including one that would result in the death of an innocent—could ever be regarded as a de minimis risk.

As a result, NDEUT can recover the preferences described at the end of Sect. 3. Even if the probability of an innocent dying as a result of my drive to the shops is greater than the probability of an innocent dying in the death lottery, within NDEUT the former may count as a de minimis risk while the latter does not. In this case, the normic de minimis expected utility of driving to the shops will exceed that of subjecting an innocent person to a death lottery, and it can exceed it by any amount, depending upon the disutility that one assigns to the death of an innocent.

Finally, NDEUT offers a solution to the disjunction problem. While a disjunction can be more *probable* than each of its disjuncts, it can never be more *normal* than each of its disjuncts. The only way in which $O_1 \vee O_2$ can result from an action A is if O_1 results from A or if O_2 results from A. Therefore, $O_1 \vee O_2$ cannot represent a more normal possibility than both O_1 and O_2 . Rather, the abnormality of $O_1 \vee O_2$ will be equal to either the abnormality of O_1 or the abnormality of O_2 —whichever is

less. Given R3, and the definition of conditional abnormality as detailed in the last section, it is straightforward to prove that, for any outcomes O_1 and O_2 that could result from an action A , $\text{ab}(O_1 \vee O_2 \mid A) = \min\{\text{ab}(O_1 \mid A), \text{ab}(O_2 \mid A)\}$.²⁰

On the normic account, unlike the probabilistic account, it is not possible for an outcome which is *not* a de minimis risk to be divided, without remainder, into outcomes which *are* de minimis risks. If $O_1 \vee O_2$ is not a de minimis risk then, on the normic account, the same will be true of O_1 or O_2 (or both). Within NDEUT, a policy of finely individuating outcomes will never generate a decision model in which *every* outcome that could result from an action represents a de minimis risk. As a result, the normic de minimis expected utility of an action, unlike the de minimis expected utility, will always be defined, no matter how finely the outcomes in the model are individuated.

While dividing outcomes more finely will never cause NDEUT to break down, it is important to note that it can still make a *difference* to the normic de minimis expected utility of an action, in which case the predictions of NDEUT, unlike those of EUT, are not completely invariant under fine-graining. Consider a coarse-grained model which includes an outcome $O_{1 \vee 2}$ and a fine-grained model in which $O_{1 \vee 2}$ is divided into the outcomes O_1 and O_2 . In NDEUT, if $O_{1 \vee 2}$ is not a de minimis risk then, as we have seen, O_1 and O_2 cannot both be de minimis risks—but one of them may be. Suppose O_2 is a de minimis risk. In this case, dividing $O_{1 \vee 2}$ will lead to an adjustment in what we are required to consider; rather than having to consider $O_{1 \vee 2}$ per se, we need only consider one of the ways in which $O_{1 \vee 2}$ could be realised—the O_1 way—and can legitimately ignore the O_2 way. This, in turn, could make a difference to the normic de minimis expected utility of the action which has $O_{1 \vee 2}$ as one of its possible outcomes. Importantly, though, this process of adjustment cannot continue indefinitely. It can be shown that there is a point in the course of fine-graining at which the predictions of NDEUT will effectively *stabilise*, and will not change as a result of further fine-graining. More precisely, it can be shown that, for any NDEUT decision model, there is a more fine-grained model for which the normic de minimis expected utilities of all possible actions will be unaffected by any further fine-graining. The proof of this result is provided in the appendix.

Having set out some of the advantages of NDEUT, I will conclude by briefly noting one potential problem. While the death lottery case revolves around a negative outcome which has a low probabilistic risk, but a high normic risk, cases of the *reverse* kind are also possible. Suppose a building suffers structural damage during an earthquake and 20 supports are put in place to hold up the roof. Suppose these

20

$$\begin{aligned}
 \text{ab}(O_1 \vee O_2 \mid A) &= \text{ab}((O_1 \vee O_2) \wedge A) - \text{ab}(A) && [\text{Def}^{\text{II}}] \\
 &= \text{ab}((O_1 \wedge A) \vee (O_2 \wedge A)) - \text{ab}(A) \\
 &= \min\{\text{ab}(O_1 \wedge A), \text{ab}(O_2 \wedge A)\} - \text{ab}(A) && [\text{R3}] \\
 &= \min\{\text{ab}(O_1 \wedge A) - \text{ab}(A), \text{ab}(O_2 \wedge A) - \text{ab}(A)\} \\
 &= \min\{\text{ab}(O_1 \mid A), \text{ab}(O_2 \mid A)\} && [\text{Def}^{\text{III}}]
 \end{aligned}$$

supports are only just able to bear the required weight—provided they all hold, the roof will remain in place, but if any one of them fails then the roof will collapse. Suppose finally that each support has been carefully checked and tested and the risk of it failing is low, on both a probabilistic and a normic interpretation. Probabilistic risk, as is well known, can *aggregate*. That is, even if the probabilistic risk of any particular support failing is low, the probabilistic risk of *at least one* support failing—and the roof collapsing—may be high. Normic risk, on the other hand, does not aggregate in this way. If each support has a low normic risk of failing then, by R3, the normic risk of at least one of the supports failing will also be low—in which case, according to NDEUT, this could potentially count as a de minimis risk. But if I was considering, say, taking up residence in the building, could it really be legitimate for me to ignore the outcome in which the roof collapses?

One immediate comeback is to point out that NDEUT will always *allow* us to set a threshold in such a way that this outcome will not count as a de minimis risk. But, unless there is some independent reason as to why the threshold should be especially demanding in this kind of case, the response is liable to seem ad hoc. I leave this problem open here—but I think it is important to correct at least one potential misconception. One might think that cases of this kind would serve to favour *DEUT*—but this is not so. As the discussion of the disjunction problem makes clear, the outcome in which the roof collapses could easily be divided into a series of fine-grained outcomes with probabilities below any de minimis threshold that we set, in which case *DEUT* would make the same problematic prediction. This also serves to illustrate why the problem posed by this case could not be resolved by moving to a mixed normic/probabilistic conception of de minimis risk, on which an outcome can only be legitimately ignored if it has a low normic risk *and* a low probabilistic risk. If we adopt a policy of finely individuating outcomes, then the second condition will effectively become redundant, and the mixed approach will collapse into NDEUT.

If the gist of this paper is right, then the ignoring of small risks is not just a way of cutting corners when decision problems become too complex or unwieldy—it goes deeper than this, and may after all have a place in a formal theory of idealised decision making. I have experimented with two attempts to import the idea of a de minimis risk into decision theory. The first attempt, based upon a probabilistic conception of a small risk, runs into two significant difficulties—it mishandles lottery cases, and threatens to collapse if the outcomes in a decision model are differentiated too finely. The second attempt, based on a normic conception of a small risk, fares better, and is able to overcome these obstacles. Whether this approach is ultimately acceptable will depend of course on the overall balance of its costs and benefits—and this is not something I have been able to comprehensively assess here.

Appendix 1

In n14 I claimed that dividing a coarse-grained outcome into a series of fine-grained outcomes will make no difference to the predictions of EUT—that is, will make no difference to the expected utility of any action—provided the utility assigned to the coarse-grained outcome is equal to the probability weighted average of the utilities

assigned to the fine-grained outcomes. In [Appendix 1](#), I make this claim, and the assumptions on which it rests, more precise, and outline a proof. The theorem proved is, in effect, a variant upon a well-known result (see for instance [Joyce, 1999](#), theorem 4.1, p. 121).

An EUT decision model has the form $\langle W, \Omega, A, O, Pr, u \rangle$. W is a set of possible worlds, with propositions modelled as subsets of W . $\Omega \subseteq \wp(W)$ is a Boolean σ -algebra of propositions. A and $O \subseteq \Omega$ are partitions of W —sets of pairwise exclusive and jointly exhaustive nonempty subsets of W —which represent the set of actions and of outcomes respectively. If $\Pi \subseteq \Omega$ is a set of propositions, let $cl(\Pi)$ be the closure of Π under the operation of disjunction. Since O is a refinement of A we have it that $cl(A) \subseteq cl(O)$ (with $cl(A) = cl(O)$ in the case of a decision under certainty). Finally, Pr is a probability function taking members of Ω into the set of real numbers in the unit interval and u is a utility function taking the members of O into the set of positive and negative real numbers (plus $-\infty$ if desired). Say that a model $\langle W, \Omega, A, O', Pr, u' \rangle$ is a *fine-graining* of a model $\langle W, \Omega, A, O, Pr, u \rangle$ just in case $cl(O) \subseteq cl(O')$ and u and u' agree for any outcomes common to O and O' .

Consider a model $\langle W, \Omega, A, O, Pr, u \rangle$ which includes $O_1 \vee \dots \vee O_n$ as one of the outcomes in O , and a fine-graining $\langle W, \Omega, A, O', Pr, u' \rangle$ in which O_1, \dots, O_n are each included in O' . Assume that O and O' are otherwise identical. Assume that the utility of $O_1 \vee \dots \vee O_n$ in the coarse-grained model is equal to the probability weighted average of the utilities of O_1, \dots, O_n in the fine-grained model— $u(O_1 \vee \dots \vee O_n) = \sum_{1 \leq x \leq n} Pr(O_x | O_1 \vee \dots \vee O_n) \cdot u'(O_x)$. For any action $X \in A$, $EU(X) = EU'(X)$. *Proof* Since $cl(A) \subseteq cl(O)$, $O_1 \vee \dots \vee O_n$ must entail a particular action $A \in A$, in which case $O_1 \vee \dots \vee O_n$ is logically equivalent to $(O_1 \vee \dots \vee O_n) \wedge A$. Since $u(O_1 \vee \dots \vee O_n) = \sum_{1 \leq x \leq n} Pr(O_x | O_1 \vee \dots \vee O_n) \cdot u'(O_x) = \sum_{1 \leq x \leq n} Pr(O_x | (O_1 \vee \dots \vee O_n) \wedge A) \cdot u'(O_x)$ we have it that:

$$\begin{aligned} Pr(O_1 \vee \dots \vee O_n | A) \cdot u(O_1 \vee \dots \vee O_n) &= Pr(O_1 \vee \dots \vee O_n | A) \cdot \sum_{1 \leq x \leq n} Pr(O_x | (O_1 \vee \dots \vee O_n) \wedge A) \cdot u'(O_x) \\ &= \sum_{1 \leq x \leq n} Pr(O_1 \vee \dots \vee O_n | A) \cdot Pr(O_x | (O_1 \vee \dots \vee O_n) \wedge A) \cdot u'(O_x) \\ &= \sum_{1 \leq x \leq n} 1/Pr(A) \cdot Pr(O_x \wedge A) \cdot u'(O_x) \\ &= \sum_{1 \leq x \leq n} Pr(O_x | A) \cdot u'(O_x) \end{aligned}$$

In this case, substituting these two quantities in an expected utility calculation for A will make no difference to the result. Since O and O' only differ with respect to $O_1 \vee \dots \vee O_n$ and O_1, \dots, O_n , and u and u' agree for all outcomes common to O and O' , any other terms in the expected utility calculation for A will be the same in the coarse-grained and fine-grained models, in which case $EU(A) = EU'(A)$. Since O and O' only differ with respect to $O_1 \vee \dots \vee O_n$ and O_1, \dots, O_n , and u and u' agree for all outcomes common to O and O' , the expected utility calculation for any other action in A will be the same in the coarse-grained and fine-grained models in which case for any $X \in A$, $EU(X) = EU'(X)$. □

Appendix 2

Within NDEUT we have enriched decision models of the form $\langle W, \Omega, A, O, \text{Pr}, \text{ab}, u \rangle$, where ab is an abnormality function taking the propositions in Ω into the set of nonnegative integers plus ∞ and meeting the constraints R1, R2 and R3 as detailed in Sect. 4. In Sect. 5, I claimed that, for any NDEUT decision model, there is a more fine-grained model for which the normic de minimis expected utility of every action is stable and unaffected by further fine-graining. In Appendix 2, I make this claim, and the assumptions on which it rests, more precise, and outline a proof.

As before, say that $\langle W, \Omega, A, O', \text{Pr}, \text{ab}, u' \rangle$ is a *fine-graining* of a model $\langle W, \Omega, A, O, \text{Pr}, \text{ab}, u \rangle$ just in case $\text{cl}(O) \subseteq \text{cl}(O')$ and u and u' agree for all outcomes in $O \cap O'$. Say that $\langle W, \Omega, A, O, \text{Pr}, \text{ab}, u \rangle$ is *normically calibrated* just in case for every outcome $O \in O$ if X is a nonempty proposition in Ω such that $X \subseteq O$ then $\text{ab}(X) = \text{ab}(O)$. Less formally, $\langle W, \Omega, A, O, \text{Pr}, \text{ab}, u \rangle$ is normically calibrated just in case none of the outcomes in O ‘cuts across’ normalcy ranks—no outcome can be divided into more fine-grained outcomes that differ in terms of their normalcy.

Any NDEUT decision model $\langle W, \Omega, A, O, \text{Pr}, \text{ab}, u \rangle$ has a fine-graining that is normically calibrated. *Proof* For each normalcy rank i , let $N_{i \geq} = \vee \{X \in \Omega \mid \text{ab}(X) \geq i\}$. $N_{i \geq}$ is the disjunction of all propositions that have an abnormality rank of at least i , and may be thought of as the set of worlds that have an abnormality of at least i . Since Ω is a σ -algebra (and closed under countable disjunction), $N_{i \geq}$ will be included in Ω , for each normalcy rank i . For each i , let $N_i = N_{i \geq} \wedge \sim N_{i+1 \geq}$ (or $N_i = N_{i \geq} \wedge \sim N_{i >}$). N_i might be thought of as the set of worlds that have an abnormality of *exactly* i ²¹. If there are no propositions in Ω that are assigned an abnormality rank of i then $N_i = \emptyset$. Let N be the set of all nonempty propositions so defined. Given a set of outcomes O , let $O+ = \{O \wedge N_i \mid O \in O, N_i \in N, O \wedge N_i \neq \emptyset\}$. Less formally, if an outcome O in O crosses several normalcy ranks then, in $O+$, O is divided according to these ranks. If, say, some of the worlds at which $O \in O$ is true have abnormality 1, some have abnormality 2 and some have abnormality 3 then, instead of O , $O+$ will contain three outcomes: O -in-the-abnormality-1-way, O -in-the-abnormality-2-way and O -in-the-abnormality-3-way.

From the definition of $O+$ it follows that $\text{cl}(O) \subseteq \text{cl}(O+)$ in which case, if u is a utility function defined on O and $u+$ is a utility function defined on $O+$ such that u and $u+$ agree for any outcomes in $O \cap O+$ then $\langle W, \Omega, A, O+, \text{Pr}, \text{ab}, u+ \rangle$ will meet the conditions for a fine-graining of $\langle W, \Omega, A, O, \text{Pr}, \text{ab}, u \rangle$. For any outcome $O+$ in $O+$, $O+$ will be equal to $O \wedge N_i$ for some $O \in O$ and $N_i \in N$. Consider a nonempty proposition $X \subseteq O \wedge N_i$. It follows that $X \subseteq N_i \subseteq \sim N_{i+1 \geq}$ in which case $\text{ab}(X) \leq i$ and $X \subseteq N_i \subseteq N_{i \geq}$ in which case $\text{ab}(X) \geq i$. Therefore $\text{ab}(X) = i$. It follows that $\langle W, \Omega, A, O+, \text{Pr}, \text{ab}, u+ \rangle$ is normically calibrated. \square

²¹ If Ω contains the singleton of each world in W —that is, if $\{w\} \in \Omega$ for each $w \in W$ —then it will be the case that $N_{i \geq} = \{w \in W \mid \text{ab}(\{w\}) \geq i\}$ and $N_i = \{w \in W \mid \text{ab}(\{w\}) = i\}$. If Ω does not contain the singleton of each world in W —and this constraint is not demanded here—then the possible world descriptions of $N_{i \geq}$ and N_i given in the main text won’t be reflected in the formalism (but can still serve as informal heuristics).

Suppose that we have a model $\langle W, \Omega, A, \mathcal{O}, Pr, ab, u \rangle$ which includes $O_1 \vee \dots \vee O_n$ as one of the outcomes in \mathcal{O} and a fine-graining $\langle W, \Omega, A, \mathcal{O}', Pr, ab, u' \rangle$ in which O_1, \dots, O_n are each included in \mathcal{O}' . Assume that \mathcal{O} and \mathcal{O}' are otherwise identical. Assume that $u(O_1 \vee \dots \vee O_n) = \sum_{1 \leq x \leq n} Pr(O_x | O_1 \vee \dots \vee O_n) \cdot u(O_x)$. Assume, finally, that $\langle W, \Omega, A, \mathcal{O}, Pr, ab, u \rangle$ is normically calibrated. For any action $X \in A$, $NDEU(X) = NDEU'(X)$. *Proof* As above, $O_1 \vee \dots \vee O_n$ must entail a given action A in which case $O_1 \vee \dots \vee O_n$ and $(O_1 \vee \dots \vee O_n) \wedge A$ are logically equivalent. Let $\psi = \wedge \{ \sim O \mid O \in \mathcal{O}, ab(O \mid A) > t \}$ and $\psi' = \wedge \{ \sim O \mid O \in \mathcal{O}', ab(O \mid A) > t \}$. Suppose $ab(O_1 \vee \dots \vee O_n \mid A) > t$. It follows, by R3, that $ab(O_1 \mid A) > t, \dots, ab(O_n \mid A) > t$ in which case $Pr_\psi(O_1 \vee \dots \vee O_n \mid A) = 0$ and $Pr_{\psi'}(O_1 \mid A) = \dots = Pr_{\psi'}(O_n \mid A) = 0$ and $u(O_1 \vee \dots \vee O_n)$ and $u'(O_1) \dots u'(O_n)$ will make no difference, respectively, to $NDEU(A)$ or $NDEU'(A)$. Since \mathcal{O} and \mathcal{O}' only differ with respect to $O_1 \vee \dots \vee O_n$ and O_1, \dots, O_n and u and u' agree for all outcomes common to \mathcal{O} and \mathcal{O}' , it follows immediately that for any $X \in A$, $NDEU(X) = NDEU'(X)$.

Assume instead that $ab(O_1 \vee \dots \vee O_n \mid A) \leq t$. Given normic calibration, it follows that $ab(O_1 \vee \dots \vee O_n \mid A) = ab(O_1 \mid A) = \dots = ab(O_n \mid A) \leq t$ and that $\psi = \psi'$. It follows further that $O_1 \vee \dots \vee O_n$ entails ψ in which case, for any proposition $X \in \Omega$, $Pr(X \mid O_1 \vee \dots \vee O_n) = Pr_\psi(X \mid O_1 \vee \dots \vee O_n)$ and $Pr(X \mid (O_1 \vee \dots \vee O_n) \wedge A) = Pr_\psi(X \mid (O_1 \vee \dots \vee O_n) \wedge A)$. Since $u(O_1 \vee \dots \vee O_n) = \sum_{1 \leq x \leq n} Pr(O_x \mid (O_1 \vee \dots \vee O_n) \wedge A) \cdot u'(O_x) = \sum_{1 \leq x \leq n} Pr_\psi(O_x \mid (O_1 \vee \dots \vee O_n) \wedge A) \cdot u'(O_x)$ we have it that:

$$\begin{aligned} Pr_\psi(O_1 \vee \dots \vee O_n \mid A) \cdot u(O_1 \vee \dots \vee O_n) &= Pr_\psi(O_1 \vee \dots \vee O_n \mid A) \cdot \sum_{1 \leq x \leq n} Pr_\psi(O_x \mid (O_1 \vee \dots \vee O_n) \wedge A) \cdot u'(O_x) \\ &= \sum_{1 \leq x \leq n} Pr_\psi(O_1 \vee \dots \vee O_n \mid A) \cdot Pr_\psi(O_x \mid (O_1 \vee \dots \vee O_n) \wedge A) \cdot u'(O_x) \\ &= \sum_{1 \leq x \leq n} 1/Pr_\psi(A) \cdot Pr_\psi(O_x \wedge A) \cdot u'(O_x) \\ &= \sum_{1 \leq x \leq n} Pr_\psi(O_x \mid A) \cdot u'(O_x) \end{aligned}$$

In this case, substituting these two quantities in a normic de minimis expected utility calculation for A will make no difference to the result. Since \mathcal{O} and \mathcal{O}' only differ with respect to $O_1 \vee \dots \vee O_n$ and O_1, \dots, O_n , and u and u' agree for all outcomes common to \mathcal{O} and \mathcal{O}' , any other terms in the normic de minimis expected utility calculation for A will be the same in both, in which case $NDEU(A) = NDEU'(A)$. Since \mathcal{O} and \mathcal{O}' only differ with respect to $O_1 \vee \dots \vee O_n$ and O_1, \dots, O_n , and u and u' agree for all outcomes common to \mathcal{O} and \mathcal{O}' , the normic de minimis expected utility calculation for any other action in A will be the same in both in which case for any $X \in A$, $NDEU(X) = NDEU'(X)$. \square

Acknowledgements This paper was presented at the Luck, Risk and Competence workshop at the University of Seville in February 2020, the Centre for Ethics, Philosophy and Public Affairs seminar series at the University of St Andrews in October 2020, the Varieties of Risk seminar series in March 2022 and the British Society for the Theory of Knowledge Conference at the University of Glasgow in September 2022. Thanks to the audiences on these occasions, and to four anonymous referees for this journal, for comments and criticisms that have significantly improved the paper. Work on this paper was supported by the Arts and Humanities Research Council (Grant No. AH/T002638/1).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adler, M. (2007). 'Why de minimis?' *Faculty Scholarship Paper* 158. Retrieved from http://scholarship.law.upenn.edu/faculty_scholarship/158
- Balfour, D. (2021). Pascal's mugger strikes again. *Utilitas*, 33(1), 118–124.
- Black, D. (2020). Absolute prohibitions under risk. *Philosophers Imprint*, 20(20), 1–26.
- Bostrom, N. (2009). Pascal's mugging. *Analysis*, 69(3), 443–445.
- Brennan, S. (2006). Moral lumps. *Ethical Theory and Moral Practice*, 9(3), 249–263.
- Buchak, L. (2014). *Risk and rationality*. Oxford University Press.
- Buchak, L. (2016). Decision theory. In A. Hájek & C. Hitchcock (Eds.), *Oxford handbook of probability and philosophy*. Oxford University Press.
- Colyvan, M., Cox, D., & Steele, K. (2010). Modelling the moral dimension of decisions. *Noûs*, 44(3), 503–529.
- Comar, C. (1979). Risk: A pragmatic de minimis approach. *Science*, 203(4378), 319.
- Cournot, A. A. (1843). *Exposition de la théorie des chances et des probabilités*. Hachette.
- Dorsey, D. (2009). Headaches, lives and value. *Utilitas*, 21(1), 36–58.
- Ebert, P., Smith, M., & Durbach, I. (2020). Varieties of risk. *Philosophy and Phenomenological Research*, 101(2), 432–455.
- Egan, A. (2007). Some counterexamples to causal decision theory. *Philosophical Review*, 116(1), 93–114.
- Gardiner, G. (2021). Relevance and risk: How the relevant alternatives framework models the epistemology of risk. *Synthese*, 199(1–2), 481–511.
- Gibbard, A., & Harper, W. (1978). Counterfactuals and two kinds of expected utility. *University of Western Ontario Series in Philosophy of Science*, 15, 153–190.
- Hájek, A. (2003). Waging war on Pascal's wager. *Philosophical Review*, 112(1), 27–56.
- Hansson, S. O. (2013). *The ethics of risk: Ethical analysis in an uncertain world*. Basingstoke: Palgrave Macmillan.
- Hayenbjelm, M., & Wolff, J. (2012). The moral problem of risk impositions: A survey. *European Journal of Philosophy*, 20(S1), e26–e51.
- Holm, S. (2016). A right against risk imposition and the problem of paralysis. *Ethical Theory and Moral Practice*, 19(4), 917–930.
- Huber, F. (2009). Belief and degrees of belief. In F. Huber & C. Schmidt-Petri (Eds.), *Degrees of belief*. Springer.
- Huemer, M. (2010). Lexical priority and the problem of risk. *Pacific Philosophical Quarterly*, 91(3), 332–351.
- Jackson, F., & Smith, M. (2006). Absolutist moral theories and uncertainty. *Journal of Philosophy*, 103(6), 267–283.
- Jackson, F., & Smith, M. (2016). The implementation problem for deontology. In E. Lord & B. Maguire (Eds.), *Weighing reasons*. Oxford University Press.
- Jeffrey, R. (1965). *The logic of decision*. McGraw-Hill.
- Joyce, J. (1999). *The foundations of causal decision theory*. Cambridge University Press.
- Kirkpatrick, J. (2018). Permissibility and the aggregation of risks. *Utilitas*, 30(1), 107–119.
- Kosonen, P. (2021). Discounting small probabilities solves the intrapersonal addition paradox. *Ethics*, 132(1), 204–217.
- Kosonen, P. (ms). How to discount small probabilities.
- Lazar, S., & Lee-Stronach, C. (2019). Axiological absolutism and risk. *Noûs*, 53(1), 97–113.

- Lee-Stronach, C. (2018). Moral priorities under risk. *Canadian Journal of Philosophy*, 48(6), 793–811.
- Lewis, D. (1981). Causal decision theory. *Australasian Journal of Philosophy*, 59(1), 5–30.
- Lundgren, B., & Stefánsson, O. (2020). Against the de minimis principle. *Risk Analysis*, 40(5), 908–914.
- Monton, B. (2019). How to avoid maximizing expected utility. *Philosophers Imprint*, 19(18), 1–25.
- Mumpower, J. (1986). An analysis of the de minimis strategy for risk management. *Risk Analysis*, 6(4), 437–446.
- Norcross, A. (1997). Comparing harms: Headaches and human lives. *Philosophy and Public Affairs*, 26(2), 135–167.
- Norcross, A. (1998). Great harms from small benefits grow: How death can be outweighed by headaches. *Analysis*, 58(2), 152–158.
- Nozick, R. (1974). *Anarchy, state and utopia*. Blackwell.
- Peterson, M. (2002). What is a de minimis risk? *Risk Analysis*, 4(2), 47–55.
- Pritchard, D. (2015). Risk. *Metaphilosophy*, 46(3), 436–461.
- Savage, L. (1954). *The foundations of statistics*. Wiley.
- Smith, M. (2010). What else justification could be. *Noûs*, 44(1), 10–36.
- Smith, M. (2016). *Between probability and certainty: What justifies belief*. Oxford University Press.
- Smith, M. (2022). The hardest paradox for closure. *Erkenntnis*, 87(4), 2003–2028.
- Smith, M. (ms) Modelling lexical priority
- Smith, N. (2014). Is evaluative compositionality a requirement of rationality? *Mind*, 123(490), 457–502.
- Spohn, W. (2012). *The laws of belief*. Oxford University Press.
- Thoma, J. (2019). Risk aversion and the long run. *Ethics*, 129(2), 230–253.
- Thomson, J. J. (1983). Some questions about government regulation of behavior. In T. Machan & M. Johnson (Eds.), *Rights and regulation: Ethical, political and economic issues*. Ballinger.
- Thomson, J. J. (1990). *The realm of rights*. Harvard University Press.
- Weinberg, A. (1985). Science and its limits: The regulator's dilemma. *Issues in Science and Technology*, 2(1), 59–72.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.