

# In Defence of Normic De Minimis Expected Utility Theory

Martin Smith

In a recent paper, Björn Lundgren and H. Orri Stefánsson (forthcoming) present three objections to normic de minimis expected utility theory (NDEUT) – a decision theoretic framework defended in Smith (2024). In this paper, I respond to these objections and outline some possible ways in which NDEUT might be modified or further developed. Like any de minimis framework, NDEUT employs a risk threshold to sort possibilities into those that must be considered when making a decision, and those that can be legitimately ignored – the so-called ‘de minimis risks’. While this threshold would usually take the form of a probability value, in NDEUT the de minimis risks are identified instead on the basis of their *abnormality*. In the first section I will set up some of the formalism used in NDEUT and consider Lundgren and Stefánsson’s first objection – which is essentially formal in nature. In the second section I will delve further into the interpretation of NDEUT and consider the second and third objections, which turn on particular examples.

## 1 Dominance

The motivation behind NDEUT is to introduce, into expected utility theory, the idea of a de minimis risk – a risk that it is so small that it can be legitimately ignored for the purpose of decision making. The formal machinery of NDEUT is in many ways similar to that of standard expected utility theory, but with a few additions which I’ll highlight. Let propositions be modelled as subsets of a set of possible worlds  $W$ . Let  $\Omega$  be a Boolean algebra of propositions – a set of propositions which includes  $W$  and is closed under negation and disjunction. Amongst the members of  $\Omega$  we distinguish two important classes of propositions: (i) *Actions* are propositions that an agent can choose to make true in a given decision problem. (ii) *States* are propositions that lie beyond the agent’s control and are relevant to determining the outcomes of actions. The set of actions  $\mathbf{A}$  and the set of states  $\mathbf{S}$  will both be partitions of  $W$  – which is to say that exactly one action will be performed and exactly one state will obtain at each possible world.

We then introduce three functions: (i) A function  $u$  taking pairs of actions and states (the members of  $\mathbf{A} \times \mathbf{S}$ ) into the set of real numbers. (ii) A function  $\text{Pr}$  taking the propositions in  $\Omega$  into the set of real numbers in the unit interval in a way that conforms to the probability axioms. (iii) A function  $\text{ab}$  taking the propositions in  $\Omega$  into the set of nonnegative integers, plus infinity, in a way that conforms to the negative ranking axioms (Smith, 2024, sections 4 and 5).  $u(A, S)$  represents the utility yielded by action  $A$  in state  $S$ .  $\text{Pr}(X)$  and  $\text{ab}(X)$  represent, respectively, the probability of proposition  $X$  and the abnormality of proposition  $X$ , given the agent’s evidence. The function  $\text{ab}$  is a new addition to the formalism of expected utility theory, its inclusion inspired by the idea that probability and abnormality offer different ways of conceptualising the *risk* of a given state (Ebert, Smith and Durbach, 2020). We assume that that the probability and abnormality of a state are both independent of which action is chosen – so  $\text{Pr}(S | A) = \text{Pr}(S \wedge A)/\text{Pr}(A) = \text{Pr}(S)$  and  $\text{ab}(S | A) = \text{ab}(S \wedge A) - \text{ab}(A) = \text{ab}(S)$  for each  $A \in \mathbf{A}$  and  $S \in \mathbf{S}$ .

We then set an abnormality threshold  $t$  such that any state with an abnormality greater than  $t$  is classified as a de minimis risk. This is another new ingredient; standard expected utility theory has no use for such a threshold, and even previous de minimis frameworks would tend work with a probabilistic threshold, characterising de minimis risks as those states that are highly *unlikely* (see for

instance the discussion in Shrader-Frechette, 1985, and Lundgren and Stefánsson, 2020). In NDEUT the de minimis risks are characterised instead as those states that would be highly *abnormal*. Let  $\psi$  be the conjunction of the negations of all de minimis risks:  $\psi = \wedge\{\sim S \mid S \in \mathbf{S} \wedge \text{ab}(S) > t\}$ . For any action  $A \in \mathbf{A}$  we can now calculate its expected utility ( $EU(A)$ ), for which every state is considered, and its normic de minimis expected utility ( $NDEU(A)$ ), for which the de minimis risks (if any) are discounted:

$$EU(A) = \sum_{S \in \mathbf{S}} Pr(S) \cdot u(A, S)$$

$$NDEU(A) = \sum_{S \in \mathbf{S}} Pr_{\psi}(S) \cdot u(A, S) = \sum_{S \in \mathbf{S}} Pr(S \mid \psi) \cdot u(A, S)$$

Finally, we have one decision procedure that selects the action or actions with the highest expected utility and one decision procedure that selects the action or actions with the highest normic de minimis expected utility. The selected action or actions will be rationally permissible for the agent, while the remainder will be rationally prohibited. The use of these two decision procedures is, again, somewhat nonstandard, but reflects the idea that it is legitimate for an agent to ignore de minimis risks (and evaluate actions in terms of their normic de minimis expected utilities) and also legitimate for an agent to consider them (and evaluate actions in terms of their standard expected utilities).

The formalism just described is, in fact, a little different from that used in Smith (2024) which dispenses with states and works instead with actions and outcomes (the foregoing is, in effect, the formalism gestured at in n3, n9 and n19). States provide a way of factoring an outcome into a contribution due to the agent and a contribution due to the world, and allow us to define more complex relations between actions, including relations of *statewise dominance* which are central to Lundgren and Stefánsson's first objection. Say that an action  $A_1$  statewise dominates an action  $A_2$  just in case  $A_1$  yields no less utility than  $A_2$  in every state and more utility than  $A_2$  in some state. Some maintain that an agent should never choose an action that is dominated by another – and yet, NDEUT predicts that this can be rationally permitted under certain circumstances.

Lundgren and Stefánsson give an example involving two actions  $A_1$  and  $A_2$  and three states  $S_1$ ,  $S_2$  and  $S_3$  such that  $A_1$  yields 1 utility in  $S_1$ ,  $S_2$  and  $S_3$ , while  $A_2$  yields 1 utility in  $S_1$  and  $S_2$  and -1 utility in  $S_3$ :

	$S_1$	$S_2$	$S_3$
$A_1$	1	1	1
$A_2$	1	1	-1

In this case  $A_1$  dominates  $A_2$ . And yet, if  $S_3$  were a de minimis risk – if  $\text{ab}(S_3) > t$  – then  $A_1$  and  $A_2$  would have the same normic de minimis expected utility (i.e. 1) and would both be rationally permitted. The first thing to note about this objection is that it is not in any way specific to NDEUT – rather, it applies to *any* de minimis framework (as Lundgren and Stefánsson note). That is, the distinctive way in which de minimis risks are identified in NDEUT plays no role in the objection (indeed we are yet to say anything substantial about it)<sup>1</sup>.

---

<sup>1</sup> Lundgren and Stefánsson (2020, section 5) level the same objection at a probabilistic construal of de minimis risks. More generally, Lundgren and Stefánsson's criticisms of NDEUT are not intended to clear the way for some alternative account of de minimis risks – rather, they would appear to reflect a general scepticism about the very idea of a de minimis risk.

One response to this objection is to bite the bullet. If we are genuinely ignoring  $S_3$  then we would regard  $A_1$  and  $A_2$  as equivalent actions. And if it is *legitimate* to ignore  $S_3$  then it is *legitimate* to regard  $A_1$  and  $A_2$  as equivalent actions. The present issue can only arise in a situation in which two actions produce exactly the same utilities in all sufficiently normal states and, according to NDEUT, these are the only states that I'm required to consider when making a decision. On the bullet-biting response, an action that is dominated, with respect to the totality of states, could still be rationally chosen, provided it is not dominated with respect to the sufficiently normal states.

That's one way to go – but a more concessive response is also available, using resources already present within NDEUT. As noted, NDEUT is equipped with *two* decision procedures – one which selects actions based on their expected utilities and one which selects actions based on their normic de minimis expected utilities. One option is to allow considerations of expected utility to enter into the second procedure as well, as a way of breaking ties. The proposal, more precisely, is this: If there is only one action that maximises normic de minimis expected utility then the procedure will select this action. If two or more actions are tied for maximum normic de minimis expected utility then the procedure will select, from amongst these, those actions that have the highest expected utility<sup>2</sup>. Even though  $A_2$ , in the above example, has the same normic de minimis expected utility as  $A_1$ , with this modification in place, it won't be selected by the second decision procedure, as the tie will be broken in favour of  $A_1$ . More generally, this modification guarantees that an action which is dominated by another will never be selected and will never be deemed rationally permissible within NDEUT<sup>3</sup>.

It's true that, under this modification, the abnormality of the different states would be irrelevant in any decisions that involve statewise dominance (or, more precisely, any decisions in which a single action dominates all the others). But even if we accept a framework like NDEUT, we shouldn't expect that abnormality will prove to be an important factor in *all* decisions. We might just as well say that *probability* is irrelevant in decisions involving statewise dominance, since we can simply choose the dominant action without worrying about the probabilities of the different states.

One might be concerned that this modification makes the second decision procedure more cumbersome and less efficient, as it now requires us to calculate *both* normic de minimis expected utilities *and* standard expected utilities, at least in cases where the former turn out to be tied. But such cases are unusual, and there's no reason to think that this procedure should be less efficient, across the board, than the calculation of standard expected utilities. And even if it were, it's not clear to what extent this would undermine NDEUT. While the notion of a de minimis risk is often linked with a desire to make decisions easier and more streamlined (see, for instance, Mumpower, 1986,

---

<sup>2</sup> If we make this modification, we may need to qualify the informal definition of a de minimis risk – changing it to something like 'A de minimis risk is a risk that can be legitimately ignored for the purposes of decision making, unless considering it makes a difference between two otherwise optimal actions'.

<sup>3</sup> For symmetry, we might also allow the *second* decision procedure to break any ties that arise with respect to the *first*. On this proposal, if two or more actions are tied for maximum expected utility, the first procedure will select, from amongst them, those actions that have the highest normic de minimis expected utility. We might think of decisions in NDEUT as having a *multi-criterial* character, with expected utility and normic de minimis expected utility providing two incommensurable criteria of choice-worthiness. In the original unmodified theory, one is permitted to choose according to either criterion. But the two proposed modifications amount to an additional requirement that one not choose actions which are *dominated*, where this is now meant in the multi-criterial decision making sense of an option that performs no better than another according to any of the criteria and worse according to some. Looking at things in this way can make the two proposed modifications seem quite natural.

Adler, 2007) one of the central themes of Smith (2024) is that this is *not* the primary motivation for adopting a de minimis framework. On the contrary, I think it can be rational to ignore low-risk possibilities even when this brings *no* gains in decision making efficiency. I will briefly return to this topic in the next section.

## 2 Cost-effectiveness and severity

Lundgren and Stefánsson's second objection to NDEUT is that it allows us to ignore risks that could be managed or mitigated in a cost-effective way. They illustrate this with a medical example. According to Lundgren and Stefánsson, conditions such as diabetes, Addison's disease and hypothyroidism all have readily available and effective treatments but, being physiologically abnormal, may be possibilities that it is legitimate to ignore, according to NDEUT. One thing it is important to note right away is that the abnormality of a possibility is something that should always be assessed relative to a body of *evidence* (Smith, 2024, pp2181-2182). If we are thinking about a case of medical *diagnosis* then this would always proceed on the basis of specific symptoms and/or risk factors. Although the possibility that someone has a given medical condition may count as abnormal in the absence of any relevant evidence, it need not be abnormal relative to a particular profile of symptoms and/or risk factors<sup>4</sup>.

Lundgren and Stefánsson anticipate this point (p6<sup>5</sup>) and note that some conditions – such as Lyme disease – may have no identifiable symptoms in the early stages. Surely this, then, is a possibility that NDEUT would allow us to ignore? Let's think through a concrete example. Suppose I am considering whether to book a trip overseas in a couple of weeks time and am struck the thought that I might be in the early stages of Lyme disease, in which case I may end up too ill to travel. I don't have any symptoms however and neither am I aware of any risk factors – that is, I'm not suffering from fever or headaches or a rash, as far as I know I've not been bitten by a tick etc. It is still *possible* that I have the condition of course – I might not notice a tick bite, Lyme disease may not be associated with these symptoms in the early stages etc. – but, under these circumstances, isn't this precisely the kind of possibility that I would be entitled to ignore? That's not to say that I *have to* ignore this possibility – as noted above, NDEUT would allow me to make my decision on the basis of standard expected utility and, thus, to factor this possibility in if I see fit to do so. But in allowing me to ignore this it's not clear that NDEUT is giving the wrong result.

The notion of abnormality that figures in NDEUT is tied to the need for *special explanation* (Smith, 2010, 2016, chap. 2, 2024, section 4). So  $ab(X)$  measures, in effect, the extent to which a possibility  $X$  would require more explanation than  $\sim X$ , given the agent's evidence. If the amount of explanation demanded by  $X$  is significantly greater than that demanded by  $\sim X$  then, depending on where we set the threshold,  $X$  may count as a de minimis risk. Here is not the place to go into this conception of abnormality in detail – but it may be instructive to briefly consider how it could apply to the present example. In the absence of symptoms or risk factors the possibility that I am in the early stages of Lyme disease would require significant explanation and could count as a de minimis risk within NDEUT. If, say, I found evidence of a tick-bite then that would reduce the need for

---

<sup>4</sup> Smith (forthcoming, section 7) outlines a more complicated formalism in which the evidence is treated as a variable, rather than being built into the interpretations of  $Pr$  and  $ab$ . As well as serving to emphasise the evidence-relativity of abnormality (and probability), this framework is also capable of modelling the effects of evidence *change*. I won't explore this further here.

<sup>5</sup> Pagination refers to the online first version of Lundgren and Stefánsson (forthcoming).

explanation in the event that I have the condition. If, on the other hand, I was displaying certain symptoms associated with Lyme disease, such as fever, headaches, or a rash, then this would *increase* the need for explanation in the event that I *don't* have the condition, as the symptoms would need to be otherwise explained. In either case, the explanation gap between my having the condition and my not having the condition would be narrowed and – depending on where we set the threshold – this could mean that the former no longer qualifies as a *de minimis* risk, and would have to be considered in any decisions that I make.

Lundgren and Stefánsson's final objection to NDEUT is that it allows us to ignore possibilities that are very *severe* in the sense that they are associated with a high disutility. They ask us to consider the risks posed by *radon*. Radon is a radioactive gas that is part of the decay chain of uranium and thorium, and can be released by building materials that naturally contain these elements – though usually only in minute quantities. According to Lundgren and Stefánsson, it would be very abnormal to find high levels of radon in building materials, but this is nevertheless a possibility that should not be ignored, given its serious consequences for human health. As noted, though, NDEUT will never classify something as a *de minimis* risk *tout court*. It is only with respect to a particular decision problem – and a particular body of evidence – that a possibility can count as a *de minimis* risk. If the relevant evidence is just that radon is sometimes released by building materials but, on average, in very low amounts then this would not, in itself, generate a need for explanation in the event that high amounts of radon are present in a given sample of materials. If, on the other hand, the materials had been thoroughly tested for radioactivity and given the all clear then *that* would generate a need for explanation if they were nevertheless releasing high levels of radon, and might lead to this possibility being classified as a *de minimis* risk within NDEUT. But information about average radon levels will not have this effect<sup>6</sup>.

Contrary to what Lundgren and Stefánsson claim, NDEUT would not classify the high radon possibility as a *de minimis* risk – or not, at any rate, given a natural way of filling in the relevant evidential details. But Lundgren and Stefánsson use this example to illustrate a more general point which is quite correct; there is nothing to prevent NDEUT from classifying *some* extremely severe possibilities as *de minimis* risks. And this runs counter to Lundgren and Stefánsson's apparent view that extreme possibilities should always be considered, regardless of how low their risk is. As they write 'But clearly, the most extreme risks should not be treated as *de minimis*' (Lundgren and Stefánsson, forthcoming, p7). I don't think that this claim is so clear however.

One of the primary motivations for NDEUT (Smith, 2024, section 1) concerns cases like the following (taken from Balfour, 2021): Suppose I am deciding between two shops at which to buy my weekly groceries. The route to the first shop runs closer to some important diplomatic buildings and, as a result, there is a slightly higher chance, on this route, that I will collide with a car carrying an important envoy, preventing them from playing a pivotal role in urgent international negotiations and indirectly triggering global nuclear war. If severity is just a function of disutility, then this would have to be counted as an extremely severe risk. And even if severity were measured by the product of probability and disutility, this outcome could still count as severe. While the probability of this outcome may be minute, given its vast disutility, the result of multiplying the two could still be significant. And yet – for this very reason – attempting to factor this possibility into my decision-

---

<sup>6</sup> If we were thinking about *de minimis* risks in *probabilistic* terms then information about average radon levels, along with information about the associated distribution, could well lead to the high level possibility being classified as a *de minimis* risk. This illustrates one significant difference between probabilistic and normic criteria for identifying *de minimis* risks, concerning the efficacy of 'purely statistical' evidence (Smith, 2024, section 4).

making would threaten to derail the decision altogether. That is, if we include this possibility in an expected utility calculation, it could easily overwhelm considerations such as the length of the two routes, the average prices at the two shops etc. – that is, the comparatively mundane factors on which the decision *should*, intuitively, rest<sup>7</sup>.

In this paper I have considered Lundgren and Stefánsson's three objections against NDEUT. I have argued that the second and third objections can be disarmed by emphasising the evidence-relativity of abnormality. Depending on how we specify the evidence, NDEUT may not allow us to ignore the possibilities that Lundgren and Stefánsson highlight – and, in those cases where it does, the prediction seems quite appropriate. While NDEUT may, in some decision problems, classify extreme possibilities as de minimis risks, I have argued that this is also the right prediction – and is in fact a prediction that plays a crucial role in one of the primary motivations for NDEUT. When it comes to Lundgren and Stefánsson's first objection, I have outlined a way of modifying NDEUT in order to avoid the result that dominated actions can be rationally permitted.

## References

- Adler, M. (2007) 'Why de minimis?' *Faculty Scholarship Paper* 158. Retrieved from [http://scholarship.law.upenn.edu/faculty\\_scholarship/158](http://scholarship.law.upenn.edu/faculty_scholarship/158)
- Balfour, D. (2021) 'Pascal's mugger strikes again' *Utilitas* v33(1), pp118-124
- Ebert, P., Smith, M. and Durbach, I. (2020) 'Varieties of risk' *Philosophy and Phenomenological Research* v101(2), pp432-455
- Lundgren, B. and Stefánsson, H. O. (forthcoming) 'Can the normic de minimis expected utility theory save the de minimis principle?' *Erkenntnis*
- Lundgren, B. and Stefánsson, H. O. (2020) 'Against the de minimis principle' *Risk Analysis* v40(5), pp908-914
- Mumpower, J. (1986) 'An analysis of the de minimis strategy for risk management' *Risk Analysis* v6(4), pp437-446
- Shrader-Frechette, K. (1985) 'Technological risks and small probabilities' *Journal of Business Ethics* v4(6), pp431-445
- Smith, M. (2010) 'What else justification could be' *Noûs* v44(1), pp10-36
- Smith, M. (2016) *Between Probability and Certainty: What Justifies Belief* (Oxford: Oxford University Press)
- Smith, M. (2024) 'Decision theory and de minimis risk' *Erkenntnis* v89(6), pp2169-2192

---

<sup>7</sup> While ignoring the nuclear war possibility would undoubtedly make for a simpler and more efficient decision, notice that this is *not* the primary reason to ignore it. Rather, this is a case in which attempting to consider every possibility, no matter how low the risk, seems to steer us in the *wrong* direction, and could lead to a decision that is intuitively irrational or, at any rate, not rationally required – such as deciding to go to the second shop, even though the first is closer, has cheaper average prices etc. Notice that NDEUT would clearly predict that the nuclear war possibility counts as a de minimis risk; relative to the background evidence, significant further explanation would be required in the event that a drive to the shops results in nuclear war.

Smith, M. (forthcoming) 'How to model lexical priority' *Ergo*