UNCERTAINTY, CHANGE, AND THREE ACCOUNTS OF RISK

Abstract

Standardly, risk is analyzed in terms of probability times severity. Recently, two alternative accounts of risk have emerged that replace the probability component in the standard account with modal closeness and normalcy, respectively. Here, these competing accounts of risk are evaluated to see if they are fit to be used in a normative theory on decision-making. Based on evaluating them in light of two elements that are part of most decisions under risk—uncertainty and change—it is argued that the modal and normic accounts of risk are unacceptable in any normative theory of decision-making.

1. Introduction

A risk is an unwanted possible event or outcome. According to the standard account of risk, risk is understood in terms of probability times severity (i.e., statistical expectation value) of the unwanted event or outcome (Hansson 2023).[1] Thus, on the standard account of risk—that is, what is sometimes called "the probabilistic account of risk"—the higher the probability, the higher the risk is, *ceteris paribus*.

Recently, the probability element in the standard analysis has been questioned, by two non-standard accounts of risk. First, according to Pritchard (2015; 2016; 2022), we ought to replace the probability element in the standard analysis with the "modal account of risk." On the modal account, it follows that the closer the nearest possible world in which the risk it obtains is, the

[1] Sometimes the notion of "risk" is used to denote only the probability component—and sometimes, at least colloquially, only the severity notion.

higher the risk is. Otherwise put, the risk of an event is higher the more similar the most similar possible world in which the risk obtains is to the actual world; the risk of an event is lower, the more dissimilar the most similar possible in which the risk obtains is to actual world, *ceteris paribus*. More recently, Ebert, Smith, and Durbach (2020) presented a similar account of risk—the so-called "normic account of risk"—according to which the risk is determined by the outcome's normalcy/abnormality. The risk of an event is higher, the more normal it is for the event to occur; and, the risk of an event is lower, the more abnormal it is for the event to occur, *ceteris paribus*.[2]

In this article, I look at two basic elements that are central to any decision-making involving risk—uncertainty and change—to compare the two non-standard accounts of risk against the standard account. The aim is to determine which account of risk is best suited—as part of a normative theory of decision-making—to manage the choices involving these basic elements.[3]

The article is structured as follows. First, I aim to further clarify the three alternate accounts of risk (in section 2). In doing so, I also raise some basic conceptual challenges for the non-standard accounts of risk. Second, I argue that the non-standard accounts of risk are less suitable than the

[2] In all fairness, no one has yet (as far as I know), defended the idea of promoting the normic account as the sole account of risk. In fact, Ebert et al. (2020, p. 446) recognizes "that the normic account is incapable of providing everything that we require from a notion of risk." Instead, they think that risk must be understood pluralistically. More recently, Smith (2022) uses normic risk to define a *de minimis* principle part of his decisional framework, while other risks are analyzed using the probabilistic account of risk (see Lundgren and Stefánsson, 2023, for a critique). Nevertheless, it is worthwhile to evaluate the normic account a serious contender (if only to add to the reasons to set it aside).

[3] Sometimes one distinguishes between *decision under risk* and *decision under uncertainty* as decision involving known and unknown probabilities, respectively. However, in most policy decision-making situations, one is necessarily engaging with risks involving some degree of uncertainty—even if it often can be quantified (e.g., as a range of probabilities).

standard analysis of dealing with problems of uncertainty (in section 3) and change (in section 4). I end the article with some brief concluding comments on future research (in section 5).

2. Explicating the non-standard accounts of risk and some basic conceptual problems

In this section, I aim to explain the non-standard accounts of risk and present some basic conceptual problems for their analyses, starting with the modal account.

As stated in the introduction, the risk of an event—on the modal account of risk—is determined solely by how similar the closest possible world in which the risk obtains is to the actual world. To see this more clearly consider the following quote from Pritchard (2015, p. 443):

> we naturally order possible worlds, and thus the possible events that obtain in those worlds, in terms of their similarity to the actual world, where similarity is determined by how much needs to change in the actual world in order to get to this possible world where the target event occurs. A close possible world is thus one that does not require much change in the actual world. A far-off possible world, in contrast, is one that does require a great deal of change in the actual world.

The normic account of risk also depends on a modal comparison, but while the modal account determines risk based on similarity to the actual world, the normic account considers how normal the most normal world in which the risk obtains is. Given that the actual world is not maximally normal, the evaluations of risk on the modal and normic accounts come apart.

Before turning to discuss how they come apart, let us first consider how the modal and normic account of risk differs from the standard—probabilistic—account of risk. Suppose we have a fair lottery such that the probability of a win is equal to 1/n (where *n* is the number of lottery tickets). In such cases, the probability of winning depends on the size of n (i.e., if n is high, then the probability is low, and if n is low, then the probability is high). However, each possibility of

winning is equally close to the actual world and equally normal, which means that modal closeness and normalcy/abnormality are independent of the size of n.

While the above example neatly illustrates the difference between the non-standard and standard account(s) of risk, most events cannot be modeled as fair lotteries; instead, most events depend on a more complex set of circumstances. As stated above, modal closeness and normalcy come apart in various situations because similarity to the actual world does not necessarily imply normalcy. Although this is clear in theory, it may not be obvious what is normal and what is abnormal unless we understand these notions relative to some set of norms. Smith attempts to explain this with the following example:

> 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. (Smith, 2022)

Thus, the idea is that (the degree of) abnormality can be understood by the need for a special explanation; the more special the needed explanation is, the more abnormal the outcome is. However, we can question whether this call for a *special explanation* really can serve the role of distinguishing what is normal from what is abnormal (and serve to help quantify that notion). Although there is a clear technical understanding of the notion:

> 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. (Smith 2022)

That does not help us unless we know how to rank propositions relative to their normalcy/abnormality. If we are supposed to use the idea of special explanations, then there are at least two problems. First, we are merely pushing the question of a normalcy/abnormality ranking to a ranking of how *special* an explanation is. Second, and more problematically, it seems as if the need for a special explanation is not tracking the normalcy/abnormality ranking as it should. Indeed, there seems to be no need for a special explanation for *abnormal events* in *abnormal situations*. On the contrary, it seems as if normal events sometimes would require special explanations in abnormal situations. Thus, the notion of special explanation does not seem to properly track the normalcy/abnormality distinctions, nor help us with its ranking.[4]

We can therefore question whether the notion of normalcy/abnormality is sufficiently well-developed to be used in practice (cf. Lundgren and Stefánsson 2023). However, for the sake of the upcoming arguments, we should aim to set these issues aside. Nevertheless, it is worthwhile to note that similar worries have been raised against the modal account of risk. Recently, Smith (2023), discussed a situation in which we are evaluating the risk that there is asbestos in a wall. Smith argues:

> On the modal account, what it means for there to be a 'low risk' that the wall contains asbestos is for the wall to not contain asbestos in any close worlds. So when you say 'As long as there's a low risk that the wall contains asbestos you should go ahead and drill' this could be paraphrased as 'As long as the wall does not contain asbestos in any close worlds you should go ahead and drill'. But surely that would invite the rebuke 'That's no help! I don't even know if there is asbestos in the wall in the actual world, let alone the other close worlds'. If I can't reasonably judge that the wall does not contain asbestos in the actual world then I can't reasonably judge that the wall does not contain asbestos in any close worlds. And if I *can* reasonably judge that the wall does not contain asbestos in the actual world then why would I care about the other close worlds? (Smith, 2023, p. 3)

[4] A potential solution might be to provide some normalcy/abnormality norms, but this is not supplied by the proponents of the normic account of risk.

I will not dwell on these issues, but it is relevant to keep them in the back of our heads as we turn to discuss the two types of concerns that I wish to raise for the non-standard accounts of risk.

3. Uncertainty

In this section, I argue that the non-standard accounts of risk are not well-suited for dealing with uncertainty, a standard element in most forms of decision-making (cf. fn. 3).

Roughly speaking we can distinguish between modal, empirical, and normative uncertainty (cf. Bradley and Drechsler, 2014). While we may disagree about whether these distinctions hold under closer scrutiny (e.g., is modal uncertainty simply a deeper form of empirical uncertainty?) or make sense as distinct metaphysical categories, that is not of any relevance here since all we need is a set of categories that can be used to evaluate the alternate accounts of risk.

Given that I am evaluating accounts of risk to see whether they are suitable for a normative theory of decision-making, I will set aside normative uncertainty since it goes to the question of what to do in cases when we are uncertain about what we ought to do (or which normative theory that is correct).

Hence, let us turn to modal uncertainty. Broadly speaking, modal uncertainty is uncertainty about possibilities and necessities. In the case of decision-making under risk and uncertainty, it is most relevant to focus on modal uncertainty as applying to the possible outcomes of an action under consideration. Modal uncertainty thus understood is problematic for any account of risk. However, the situation is strictly more problematic for non-standard accounts of risk as they require more modal knowledge than the standard account. Take the modal account of risk. To judge how modally close a specific possible outcome is, we not only need to know details about that possible outcome, but we also need to be able to rank it against other possibilities so that we

can judge how close or distant it is. The same problem extends to the normic account of risk, even if we are comparing possibilities relative to their normalcy/abnormality. Simply put, the non-standard analyses of risk require a broader set of knowledge of possibilities beyond the possible outcomes; specifically, what is needed is information that allows for the ranking of the possible outcomes, relative to other possibilities and the actual world, to determine their closeness and normalcy/abnormality, respectively. For the standard account of risk, no such information is needed. All we need to be able to do is evaluate the probability (or a range thereof) of the possible outcome, which does not require any detailed knowledge of other possibilities. This holds even if we think of probabilities as a modal concept—for example, if we think of probabilities as being a function of distributions of possible worlds, all we need is to know the distributions, we do not need to possess any further information about their state of affairs to determine modal closeness or normalcy. Generally, modal uncertainty does not imply an inability to make a probability judgment—unless we are fully uncertain about the possible outcome of an action, but in such cases, the standard and non-standard analyses of risk all fare equally badly.

What about empirical uncertainty? Roughly speaking, empirical uncertainty can be understood as uncertainty about the actual state of affairs. On any account of risk, complete empirical uncertainty would make us completely helpless. However, empirical uncertainty usually comes in degrees, or it is reducible to a set of facts. Nevertheless, even limited empirical uncertainty affects all accounts of risk when it comes to the normative evaluation of whether an action is motivated (i.e., even if we know what consequences will result from a given action, we may lack sufficient information about the actual state of affairs to know whether that outcome is beneficial or not). In such cases, the negative effects of empirical uncertainty affect all three accounts of risk. However, it is easy to see how the non-standard accounts of risk would fare worse. Indeed, while the evaluation—as just described—would affect all accounts equally, the quantification of the

uncertainty element would affect the non-standard accounts more. The reason for this is simple, just as with modal uncertainty, empirical uncertainty affects the ability to evaluate the modal closeness and normalcy/abnormality, respectively. Whether something is modally close to the actual world requires certainty information, none of which would affect any probability estimates. Similarly, ranking of normalcy/abnormality would also require certainty informational input. However, the situation is strictly worse for the modal account of risk since risk is determined relative to the actual world. While for the normic account of risk, the actual world is just one possible world in the modal space according to which the normalcy/abnormality of an outcome is judged.

4. Change

In this section, I argue that the non-standard accounts of risk fare worse in their ability to evaluate change (a necessary component in any risk). To make the example clearer I focus on transformative change or impact.

We know that certain actions can have a transformative impact on the world. That is, it can change the world radically. For example, so-called "socially disruptive technologies" often have a transformative impact on the world (see, e.g., Hopster, 2021). Potential examples of socially disruptive technologies include telephones, cars, electricity, flight, computers, smartphones, social media, AI, and CRISPR/Cas9.

What these technologies have in common is that they either have a transformative impact on the world or they have the potential to have a transformative impact on the world. (Whether this is good or bad, I am here setting aside, since I am analyzing the probability, closeness, and normalcy/abnormality element in the risk accounts.)

Let us start with the modal account for risk. While transformative impacts can vary in how radically they change the world, the change is radical, which means that the outcome is, per definition, *modally distant*. Hence, on the modal account of risk, the risk of transformative impact and radical change is, per definition, low. However, as the history of technological development shows, radical changes are sometimes to be expected. Indeed, those who worry about AI safety would hardly accept the argument that scenarios of how AI can cause existential risk are low-risk events because the outcome would be so different from the actual world. Indeed, it seems odd to think that radical change should—per definition—be a low-risk event. Momentarily, I will present an example to illustrate this problem more clearly, but before doing so I will first mention the normic account of risk so that I can present a case direct against both.

For the normic account of risk, the issue is a bit more complicated since a transformative outcome can make the world more normal. However, that does not seem to help since it implies that the quantification of the risk—in terms of normalcy and abnormality—is arbitrary. That is, for some subset of such risks, the risk of transformative impact will be considered low. For some other subsets, the risk of transformative impact will be considered high. But none of these subsets seems to track the actual risk properly since that is independent of its normalcy/abnormality. Thus, the normic account of risk will suffer both from arbitrariness and from incorrect risk quantification.

To see this more clearly, let us turn to an example of the problem of incorrect risk quantification: Suppose we have two possible risks, A and B, which are equal relative to their severity and probability. However, B differs from A in being substantially more modally distant *and* less normal outcome. On the standard account of risk, A and B are equivalent, but on the non-standard accounts, they are not. On the non-standard accounts of risk, B is, in virtue of being more distant and less normal, a smaller risk than A.

In itself, this result is not a problem since, after all, on the standard account of risk, the risks are equivalent. Hence, under most normative theories of decision-making giving priority to one over the other would—as it should—be permissible.

However, if we accept that B is a smaller risk than A, then that would have unacceptable consequences in situations when A and B differ in severity (or, more specifically, goodness and badness). No matter how the severity element is combined with the non-standard accounts of risk, it would follow that there is some risk B* such that B* differs from B only in so far that the severity of B* is strictly worse than B and B* is a smaller risk than A. But this is an unacceptable outcome since it would imply that if we used the modal or normic account in a decision-making framework we would, over time, make the world worse.

When it comes to the standard account of risk, the probabilistic account of risk is not dependent upon (dis)similarity or normalcy/abnormality, so it will not be sensitive to these examples. Simply put, the risk of an impact will depend on its likelihood (and severity), irrespective of how radical the change of that impact is.

5. Concluding comments

In this paper, I have presented two challenges for using the modal and normic accounts of risk as part of a normative theory of decision-making, arguing that they fail in cases involving uncertainty and change, two basic elements in most standard decisions. Moreover, based on a brief analysis of some of the basic conceptual problems of the non-standard accounts of risk, I indicated that they may need further development before they can be applied in practice.

Therefore, I conclude that we should retain the standard—probabilistic—account of risk when it comes to normative theories of decision-making.

However, as I noted earlier, Ebert et al. argue that risk should be analyzed pluralistically (see fn. 2). This raises the question of whether there is room for a pluralistic account of risk given the arguments I have presented here. However, when it comes to decision-making, the problems that I have raised are not asymmetrically distributed. On the contrary, they are lopsided towards the non-standard accounts or risk. Indeed, there seems to be no clear benefit for the non-standard account of risk when it comes to decision-making. Thus, as a theory of normative decision-making is concerned, there is no room for plurality on the account(s) of risk.

Nevertheless, that does not mean that there is no role for alternate accounts of risk in other areas. For example, non-standard accounts of risk may be beneficial when it comes to understanding risk perception among lay people (see Ebert et al., 2020, for some empirical data and Pritchard, 2015, for a more theoretical discussion based on other research). This raises the question of whether, and if so how, lay people's risk perception should be taken into consideration in an acceptable normative theory of decision-making. However, that is the topic for another paper.

REFERENCES

Bradley, R., & Drechsler, M. (2014). Types of uncertainty. *Erkenntnis* 79(6): 1225–1248. https://doi.org/10.1007/s10670-013-9518-4.

Ebert, P. A., Smith, M., & Durbach, I. (2020). Varieties of Risk. *Philosophy and Phenomenological Research* 101(2): 432–455. https://doi.org/10.1111/phpr.12598.

Hansson, S. O. (2023). Risk. *The Stanford Encyclopedia of Philosophy* (Summer 2023 Edition). https://plato.stanford.edu/archives/sum2023/entries/risk/.

Hopster, J. (2021). What are socially disruptive technologies? *Technology in Society* 67, 101750. https://doi.org/10.1016/j.techsoc.2021.101750.

Lundgren, B., & Stefánsson, O. (2023). Can the normic de minimis expected utility theory save the *de minimis* principle? *Erkenntnis*.

Pritchard, D. H. (2015). Risk. *Metaphilosophy* 46: 436–461. https://doi.org/10.1111/meta.12142.

Pritchard, D. H. (2016). Epistemic risk. *Journal of Philosophy* 113: 550–571. https://doi.org/10.5840/jphil20161131137.

Pritchard, D. (2022). In defence of the modal account of legal risk. *Synthese* 200, 290. https://doi.org/10.1007/s11229-022-03693-z.

Smith, M. (2022). Decision theory and De Minimis Risk. *Erkenntnis*. https://doi.org/10.1007/s10670-022-00624-9.

Smith, M. (2023). An objection to the modal account of risk. *Synthese* 201, 156. https://doi.org/10.1007/s11229-023-04134-1.