


Varieties of Risk¹

PHILIP A. EBERT 

University of Stirling

MARTIN SMITH

University of Edinburgh

IAN DURBACH

University of Cape Town & St Andrews

1. Introduction

The notion of risk plays a central role in economics, finance, health, psychology, law and elsewhere, and is prevalent in managing challenges and resources in day-to-day life. In recent work, Duncan Pritchard (2015, 2016) has argued against the orthodox probabilistic conception of risk on which the risk of a hypothetical scenario is determined by how probable it is, and in favour of a *modal* conception on which the risk of a hypothetical scenario is determined by how modally close it is. In this article, we use Pritchard's discussion as a springboard for a more wide-ranging discussion of the notion of risk. We introduce three different conceptions of risk: the standard probabilistic conception, Pritchard's modal conception, and a normalcy conception that is new (though it has some precursors in the psychological literature on risk perception). Ultimately, we argue that the modal conception is ill-suited to the roles that a notion of risk is required to play and explore the prospects for a form of *pluralism* about risk, embracing both the probabilistic and the normalcy conceptions.

2. Stage setting: risk judgments

Our focus is on *risk judgments* of the following kind:

- 1 Given our evidence, the *risk* of a plane crash is very *low*.
- 2 Given our evidence, there's a *high risk* of food poisoning at this restaurant.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

¹ The paper was jointly devised and written by Ebert and Smith. Survey design is due to Ebert, Smith, and Durbach. Ebert and Durbach are responsible for survey recruitment and survey analysis.

Often when we make a risk judgment, we would omit any mention of evidence and simply say something like:

- 3 The risk of a plane crash is very low.
- 4 There's a high risk of food poisoning at this restaurant.

We take the view that a risk judgment always implicates a body of evidence, which we refer to as the *background evidence*. In cases where the background evidence is not made explicit, we take it to be supplied by the context of utterance and, in typical cases, to be the evidence possessed by the one making the judgment. That is, we are inclined towards a contextualist semantics for utterances such as 3 and 4, on which their truth conditions feature an evidence parameter, the value of which is fixed by the context. The semantics of such utterances is not, however, our primary concern here.

As well as making categorical risk judgments such as the above, we often make comparisons. While we may judge that the risk of a plane crash is very low, we may also judge that there is a *higher* risk of a car crash on the way to the airport. As well as judging that there's a high risk of food poisoning at a particular restaurant, we might also judge that there is a *lower* risk of food poisoning at the restaurant next door.

Moreover, while we often speak about the riskiness of feared *events*, such as plane crashes, food poisoning, etc., we can also assess the risk of *states of affairs*. For instance, before drilling into the wall of a 1970s West Australian house, one might assess the risk that the wall contains asbestos, or jurors in a criminal trial, when contemplating a guilty verdict, might consider the risk that the defendant is innocent, or a mountaineer may ponder the risk that the snow conditions are unfavourable for a climb. Here, we treat *propositions* as the primary bearers of risk, with the riskiness of an event or state of affairs corresponding to the riskiness of the proposition that the event occurs, or the state of affairs obtains. As well as making judgments about the risk of specific events and states of affairs, people also assess the risk of *activities* or *decisions*, saying things like 'Drilling into this wall is risky', 'It would be risky to attempt a climb under these conditions'. These judgments are important to understanding the connections between risk and decision making but we put them to one side here.

3. The probabilistic account of risk

According to the probabilistic account of risk, the *risk* of a proposition P is determined by the *probability* of P—the higher the probability, the higher the risk. On this view, the risk of P is higher than the risk of Q just in case P is more probable than Q is. The probability here should be interpreted as *evidential* probability—probability conditional upon the background evidence.

We follow Pritchard in treating the probabilistic account as the orthodox account of risk (Pritchard, 2015, section 1). It is important to note, however, that this is somewhat different to the definition of risk that has become standard in professional risk management and in some economics textbooks. On this definition, risk is equated with expected disvalue: the probability of an outcome multiplied by a measure of how severe or detrimental it would be. There is some evidence to suggest that this is best regarded as a technical definition (and a relatively recent one) which doesn't directly connect with our

ordinary risk judgments (Boholm et.al. 2016).² In any case, our focus here will be on risk comparisons made in cases where severity is held constant and the two accounts of risk generate the same predictions.

It is widely acknowledged that the intuitive risk judgments that people are inclined to make don't always align with what the probabilistic account would predict. In research dating back to at least the 1970s psychologists have identified a range of ways in which people's intuitive judgments systematically deviate from what the probabilistic account would sanction. These results have been treated not as evidence against the probabilistic account but, rather, as revealing important heuristics and biases which guide our judgments about risk and probability. For example, Kahneman & Tversky propose a number of heuristics that underlie risk judgments, such as the availability and representation heuristics, which in many cases lead to systematic deviations from what the probabilistic account would predict to be correct (Kahneman & Tversky 1973, Tversky & Kahneman 1974; for an overview, see Kahneman 2011).

Under the influence of the dual process theory of reasoning (e.g. Epstein 1994, Slovic 1996, and many others; for an overview, see Evans 2011), more recent work has emphasised the role that *emotions* play in risk judgments and highlighted the way in which the determinants of *cognitive* risk judgments and *emotive* risk judgments are often very different (Loewenstein, et.al. 2001, Finuncane et.al. 2000, Slovic, et.al. 2002, Slovic, et.al. 2004). It has been suggested that:

“feelings about risk are largely insensitive to changes in probability, whereas cognitive evaluations do take probability into account. As a result, feelings about risk and cognitive risk perceptions often diverge, sometimes strikingly.” (Loewenstein, et.al. 2001, p.271)

Proponents of the ‘risk as feeling’ hypotheses adopt a dual system approach to risk judgments whereby some risk judgments will be the output of a broadly *cognitive* system, and predominantly determined by features such as estimated probabilities, while other judgments will be generated by a more *affective* system, and influenced by the ease or vividness with which the outcomes can be imagined, and by personal experience with the relevant outcome (Loewenstein, et.al. 2001). That risk judgments can deviate from the probabilistic norm is thus to be expected on a dual system approach.

In a series of recent articles, Pritchard (2015, 2016) has criticised the probabilistic account of risk and argued in favour of an alternative that he terms the *modal* account. Pritchard asks us to compare two cases in which the same bad outcome may occur in two very different ways which are stipulated to have the same probability. Given these stipulations, the probabilistic account predicts that the risk in the two scenarios ought to be equal but, according to Pritchard, one scenario is “clearly far *more risky* than the

² It's clear that the severity of various possible outcomes is one factor governing judgments about the riskiness of activities or decisions. We think it is less clear that this factor is involved in ordinary judgments about the risk of events or states of affairs. Think of two negative events that might occur during a plane flight: (i) the plane might crash, killing all on board, (ii) one might be seated next to an irritating person. If risk is calculated as the product of probability and severity, it's plausible that (i) ought to be regarded as the higher risk event. Yet it would seem at best rather misleading to say “there is a greater risk of the plane crashing than of me being seated next to an irritating person.” We won't pursue this issue further here.

other” (Pritchard 2015, p. 441, our italics). Pritchard’s cases are as follows (adapted from Pritchard 2015, p.441):

(Bomb 1) An evil scientist has rigged up a large bomb, which he has hidden in a populated area. If the bomb explodes, many people will die. There is no way of discovering the bomb before the time it is set to detonate. The bomb will only detonate, however, if a set of six specific numbers between 1 and 49 come up on the next national lottery draw. The odds of these numbers appearing is fourteen million to one. It is not possible to interfere with this lottery draw.³

(Bomb 2) Same as above, however the bomb will only detonate if a series of three highly unlikely events obtains. First, the weakest horse in the field at the Grand National, Lucky Loser, must win the race by at least ten furlongs. Second, the worst team remaining in the FA Cup draw, Accrington Stanley, must beat the best team remaining, Manchester United, by at least ten goals. And third, the Queen of England must spontaneously choose to speak a complete sentence of Polish during her next public speech. The odds of this chain of events occurring are fourteen million to one.

According to Pritchard, (Bomb 1) is riskier than (Bomb 2) and, if forced to choose, we ought to prefer (Bomb 2) to (Bomb 1). We admit there may be some temptation to judge as Pritchard does, but would be cautious in drawing any immediate conclusions about the viability of the probabilistic account. Even if this judgment is widely shared (more on this below), this would seem to put Pritchard’s example in the same category as other existing examples in which risk judgments have been shown to deviate from what the probabilistic account predicts. Nonetheless, Pritchard invests his thought experiment with an added significance: as putting pressure on the probabilistic account of risk, rather than exposing a heuristic or bias which can affect our judgments.

A number of heuristics and biases that psychologists have identified could potentially come into play in the kinds of cases Pritchard describes. One thing that we might observe is that the conditions for triggering a detonation in (Bomb 1) seem far *easier to imagine* than the conditions described in (Bomb 2). To imagine that the weakest horse in the Grand National wins by at least 10 furlongs or that the worst team remaining in the FA cup draw beats the best team remaining by at least 10 goals or that the Queen spontaneously chooses to speak a complete sentence of Polish during her next public speech plausibly involves the construction of a rich accompanying narrative. In contrast, to imagine six particular numbers coming up in the next National Lottery draw requires no particular narrative. In a range of studies, psychologists have shown that there is a positive correlation between the ease with which an event can be imagined or recalled, and how probable we estimate the event to be—a phenomenon variously labelled as the *availability heuristic* (Kahneman & Tversky 1973, Tversky & Kahneman 1974) or the *simulation heuristic* (Kahneman & Tversky 1982).

Additionally, one might hypothesise that ease of imagining may exert an influence on intuitive risk judgments even when probabilities are made explicit. That is, one might hypothesise that the greater ease with which one can imagine the scenario as described in

³ We take it that Pritchard’s intent here is that the bomb would be triggered by the presence of 6 particular numbers in the lottery draw, irrespective of their order. That is, taking order into account, there are 720 possible lottery outcomes that would serve to trigger the bomb. This, at any rate, is what gives us a probability close to the fourteen million to one figure.

(Bomb 1) might lead one to assign it a higher risk than (Bomb 2) even though the probabilities are stipulated to be equal.⁴

Further, according to the *competence hypothesis* (Heath & Tversky 1991), preference to bet, when the relevant options are judged equally probable, can depend on how knowledgeable and competent people take themselves to be with respect to the bets in question. Heath & Tversky showed that sports fans prefer to take bets on sporting events rather than on chance events, even when they judge the probabilities to be equal, since success in the former would be attributable to one's knowledge and competence as opposed to mere luck. Given that (Bomb 1) is a scenario in which success or failure is simply due to luck, someone who takes himself to be somewhat knowledgeable about the conditions in (Bomb 2) may be expected to prefer that option. In the highly likely event that a subject thereby succeeds in saving lives, that subject may expect to deserve some credit.⁵

Another confounding factor is the plausible variation, across the two scenarios, of one's confidence about the probability assignments in question. The probability of a lottery outcome can be straightforwardly determined, given the properties of the lottery. However, it is much less clear how to determine the probabilities of the triggering conditions in (Bomb 2). Pritchard simply stipulates a value for the probability of an explosion in (Bomb 2), but it is natural to take this as much more speculative and less certain than the corresponding value in (Bomb 1).⁶ Gärdenfors & Sahlin (1982) and Goldsmith & Sahlin (1983) highlight a number of ways in which preferences in an equiprobable choice scenario are affected by differences in one's confidence about the probabilities. In a finding that is particularly significant for the present discussion, Goldsmith and Sahlin observe that when asked to choose between equiprobable bets under lose-or-not-lose conditions (where only a neutral or adverse outcome is in prospect), some subjects prefer bets where the probabilities are judged to be less certain. Such a mechanism offers another possible explanation for a preference for (Bomb 2).⁷

Finally, some (including the authors) might well think that the stipulated probabilities in (Bomb 2) are unreasonably high. The probability that the weakest horse in the Grand National wins by at least 10 furlongs *and* the worst team remaining in the FA cup draw

⁴ Indeed, in cases with affect rich outcomes—and maybe the Bomb scenario is one such outcome—studies show that subjects tend to give less weight to explicitly given probabilities when assessing the risk (Rottenstreich & Hsee 2001).

⁵ These considerations nicely tie in with Pritchard's observation that in Bomb 1 one would tensely await the lottery results while in Bomb 2 there would be no similar cause for alarm. Compare our appendix for further discussion of Heath & Tversky's hypotheses in the context of our survey results.

⁶ Instead of drawing on higher-order confidence, one may allude to Knight's (Knight 1921) famous distinction between *risk* as measurable and *uncertainty* as non-measurable and observe that the first scenario involves a clearly identifiable risk, while the second scenario involves uncertainty. Also relevant here is the well-known ambiguity paradox (Ellsberg 1961) which highlights that the way in which people experience uncertainty involves aspects that are not easily captured by the relevant probabilities. Further discussion of this issue would lead us too far afield. However, see Einhorn & Hogarth (1985) who discuss the interrelationship between Ellsberg's insights and the issue of higher-order probabilities alluded to in the main text.

⁷ Goldsmith & Sahlin (1983) also found that some subjects, rather than exhibiting a blanket preference for less certain probabilities, exhibited a preference for less certain probabilities when the probability of loss was high and a preference for more certain probabilities when the probability of loss was low. In the bomb scenarios, where the probability of an adverse outcome is low, subjects who chose in accordance with this pattern might be expected, all else equal, to prefer (Bomb 1).

beats the best team remaining by at least 10 goals *and* the Queen spontaneously chooses to speak a complete sentence of Polish during her next public speech might reasonably be regarded to be far lower than the stipulated value of 1 in 14 million.⁸ Given a more realistic estimate of this probability, the probabilistic account will straightforwardly predict that the risk in (Bomb 2) is lower than in (Bomb 1), offering another possible explanation for the intuition.

These hypotheses are admittedly speculative, yet they offer potential ways of explaining Pritchard's judgment about the bomb scenario in a way that is compatible with the probabilistic account of risk. In so far as Pritchard intends to challenge the probabilistic account as a descriptively adequate theory of ordinary risk judgments, hypotheses of this kind need to be considered. Moreover, as mentioned above, it is also crucial, in this context, that Pritchard's judgment is widely shared.

To investigate this further, we put together a short survey using the original vignettes from Pritchard (2015). The details of our survey (recruitment, vignettes, results) can be found in the appendix, which includes a more thorough discussion of the results. However, to summarise the main findings: while the claim that people tend to prefer (Bomb 2) over (Bomb 1) receives support from our survey, the claim that people judge (Bomb 2) to be less risky is not supported, with the majority of subjects judging that the risk of the bomb detonating, in the two scenarios, is equal (See Figure 1).

These results don't, on their own, show that Pritchard's judgment about the bomb cases is wrong, and don't necessarily undermine his case against the probabilistic account. After all, Pritchard could maintain that his premises don't need to be supported by experimental surveys. We won't engage here in a wider debate about philosophical methodology. Rather, more modestly, we think that what the survey results suggest is that, even in the kinds of cases to which Pritchard draws attention, the probabilistic account offers a reasonable, though far from perfect, prediction as to how people judge risk.⁹

Rather than relying upon judgments about particular examples, an alternative approach to challenging the probabilistic account is to argue that this notion of risk fails to make sense of certain pervasive patterns of reasoning. Consider what we call *checklist reasoning* about risk. It is often natural to assess the risk of an event or state of affairs by separately assessing the risk of various possible ways in which it could occur or obtain. In checklist reasoning, an event or state of affairs is judged to be low risk on the grounds that every possible way in which it could happen is regarded as being low risk. Suppose

⁸ We were recently quoted 5000 to 1 odds for a bet on an outcome similar to Pritchard's first condition ("Newport County to draw Man City in the FA Cup and beat them by at least 10 goals. We would be happy to quote 5,000/1 for a maximum bet of £10." email correspondence on 1/2/2018 with www.williamhill.co.uk). Additionally, other bookmakers refused to offer quotes on the second condition, with one citing the "extremely improbable nature" of the event as a reason to not offer us any odds. So, assuming similar odds for this bet, and assuming that the two conditions are independent, we already arrive at a 1 in 25 million probability for both conditions being met. If we assume the same probability for the third condition and treat it as mutually independent of the other two, we derive a probability of 1 in 125,000,000,000 (1 in 125 billion) for the bomb actually exploding in (Bomb 2). We could, of course, change the conditions in (Bomb 2), or supply a more realistic probability value and alter the lottery in (Bomb 1) to restore the match in probabilities. It may be that these changes wouldn't significantly alter people's judgments about the cases—but we regard this as something of an open question, which is the subject of further planned experimental work.

⁹ Thanks to an anonymous referee for pressing us on this point.

Survey Response (total number)




Figure 1. Survey Response of 105 subjects listing the number of respondents judging preference for (Bomb 1), preference for (Bomb 2), or no preference (left) and providing comparative risk judgments about the two scenarios (right).

you are concerned that you will be late for an important work meeting and I try to reassure you as follows:

You’ve left enough time to get to work, so there are really only two ways in which you could be late: if the car breaks down or you get stuck in a serious traffic jam. The risk of the car breaking down is low—the car is still relatively new and you’ve just had it serviced. And the risk of you getting stuck in a traffic jam is low too—it’s usually quiet on the roads at this time, and there are no reports of traffic problems. So, you needn’t worry, the risk of you being late to work is low.

Formally, *checklist reasoning* involves the following inference pattern:

- 1 There is a low risk of P.
- 2 There is a low risk of Q.
- 3 There is a low risk of $(P \vee Q)$.

Not only is this an intuitive pattern, but in certain settings it may be that something approaching this reasoning is formally prescribed. Checklist reasoning is implicated, for instance, in the practice of *de minimis* risk management, on which risks that are deemed to be suitably low are ignored (see Comar, 1979, Mumpower, 1986, for critical discussion, see Peterson, 2002). More precisely, *de minimis* risk management involves ranking adverse possibilities in terms of their risk and specifying a low, but non-zero, level of risk to serve as the *de minimis* threshold. Those possibilities which fall below the threshold are disregarded, while those which are above the threshold are subjected to a comprehensive risk analysis in order to determine whether measures should be taken to protect against them. This approach may have been first explicitly employed by the US Food and Drug Administration in the 1960s, and is still widely used in the regulation of health and environmental risks (see Peterson, 2002).

Suppose P and Q both fall below the *de minimis* risk threshold. If one disregards the possibility of P and takes no measures against it, and disregards the possibility of Q and

takes no measures against it, then one has in effect disregarded the possibility of $P \vee Q$ and taken no measures against it. But, unless checklist reasoning is valid, there is no guarantee that the risk of $P \vee Q$ is below the de minimis level. Consider again the above example. Suppose the risk of a car breakdown and of a traffic jam are both reckoned to be below the de minimis threshold. On the de minimis approach, both of these possibilities are disregarded and no measures are taken against them. Since these are the only ways in which you could be late for the meeting, no measures are taken against this possibility. And yet, if checklist reasoning fails, then the risk that you will be late for the meeting may well be *above* the de minimis level, in which case the de minimis approach recommends that we *do* take measures against it, or at least seriously consider the option of doing so. Hence, without the validity of checklist reasoning, it is doubtful that this approach to risk management is coherent.¹⁰

Another setting in which checklist reasoning seems to play a normative role is in the context of *legal fact-finding*. Many of the rules of criminal procedure are designed to minimise the risk of convicting an innocent person. The high standard of proof for criminal trials—beyond reasonable doubt—is intended to ensure that a defendant cannot be convicted unless the risk that he or she is innocent, given the presented evidence, is very low. And yet, most criminal charges will involve a number of essential elements, and the established legal practice is to apply the standard *separately* to each element (see, for example, Allen & Jehl 2004, Allen 2008, Spottswood 2016).

For instance, to be guilty of theft, in many jurisdictions, a defendant must have:

- i) taken property from the rightful owner.
- ii) intended to permanently deprive the owner of that property.
- iii) lacked the owner's permission.

A defendant will be innocent of theft if any one of these conditions is not met. Under prevailing legal practice, a defendant will be convicted if each of these conditions is proved beyond reasonable doubt. But this merely ensures that there is a low risk that (i) the defendant did not take property from the rightful owner, there is a low risk that (ii) the defendant did not intend to permanently deprive the owner of that property, and there is a low risk that (iii) the defendant had the owner's permission. To conclude from this that there is a low risk that the defendant is in fact innocent of the crime requires *checklist reasoning*.

On the probabilistic account of risk, however, checklist reasoning is an invalid inference pattern. The probability of $P \vee Q$ can be higher than both the probability of P and the probability of Q .¹¹ If I roll a fair die, the probability that I will roll a 6 is $\frac{1}{6}$ and the probability that I will roll a 5 is $\frac{1}{6}$. The probability that I will roll either a 5 or a 6 is $\frac{1}{3}$.

¹⁰ Checklist reasoning may also be implicit in the practice of calculating *risk scores*—widely used in IT, medical, or transport contexts—by summing (weighted) scores for individual risk factors. In particular, a commitment to checklist reasoning may be evident for scoring systems in which risk factors can receive a score of 0 when the risk is deemed to be suitably low. Of course, other systems for calculating risk scores don't have this feature. For a critical discussion of the use of risk scores in risk management, see (Hubbard 2009). For an example of a risk matrix that appears to capture something approaching checklist reasoning in the context of IT security, see (Open Group Standard, 2013, p.26)—a document whose aim is to “provide a set of standards for various aspects of information security risk analysis” (op.cit., p.1).

¹¹ Probability is *additive*, which is to say that, if P and Q are incompatible propositions, $\Pr(P \vee Q) = \Pr(P) + \Pr(Q)$. More generally, for any P and Q , $\Pr(P \vee Q) = \Pr(P) + \Pr(Q) - \Pr(P \wedge Q)$.

According to the probabilistic account of risk, even if the risk of P is low and the risk of Q is low, it needn't follow that the risk of $P \vee Q$ is low.

Here then are two established practices which are arguably at odds with the probabilistic account of risk. Importantly, we don't mean to suppose that such practices are beyond criticism—indeed, we are open to the idea that they may need reform. It seems to us, however, that the existence of these practices provides good motivation to consider other notions of risk that may make different predictions about checklist reasoning. We will turn to such a notion in what follows.

4. The modal account of risk

As mentioned above, Pritchard uses the Bomb example not merely to put pressure on the probabilistic account of risk, but to motivate an alternative which he terms the *modal* account of risk. In spite of the odds against it, the triggering event in (Bomb 1) is, according to Pritchard, something that could *easily* happen—all that is needed is for a few coloured balls to fall in the right way at the right time. In contrast, the triggering conditions in (Bomb 2) all seem to be very far-fetched—it couldn't easily happen that the weakest horse in the Grand National wins by at least 10 furlongs or that the worst team remaining in the FA cup draw beats the best team remaining by at least 10 goals or that the Queen spontaneously chooses to speak a complete sentence of Polish during her next public speech (or so we would tend to think).

One way to formalise the idea of *easy possibility* is by appealing to an ordering of possible worlds reflecting how similar they are to the actual world—an idea familiar to philosophers since Lewis (1973, 1979). The degree of similarity of another world is determined by how much would need to change about the actual world in order to bring it into conformity with the world in question (Pritchard, 2015, p. 443f). As Pritchard points out, very little would need to change in order for a set of six numbers to come up in the next national lottery draw—all that is required is that six coloured balls fall in a particular configuration. Thus, there is a very similar or *close* world in which this obtains. More generally, it seems that any lottery outcome could obtain as easily as any other. Amongst the most similar worlds in which the lottery is run will be worlds in which each outcome obtains, including those which feature the six specified numbers. In contrast to this, a great deal would need to change (so we might suppose) in order for the weakest horse in the Grand National to win by at least 10 furlongs or the worst team in the FA cup draw to beat the best team by at least 10 goals or the Queen to choose to speak a complete sentence of Polish during her next public speech. As such, the most similar worlds in which these events occur are very dissimilar or *distant* worlds.

According to the modal account of risk, the risk of a proposition P is determined by the similarity of the most similar worlds in which P is true: the more similar these worlds, the higher the risk. On this view, the risk of P is higher than the risk of Q just in case the most similar worlds in which P is true are more similar than the most similar worlds in which Q is true. Naturally, this world ordering should be restricted to worlds in which the background evidence holds, and should exclude worlds which are inconsistent with the background evidence.

As a result, on the modal account of risk (Bomb 1) and (Bomb 2) have very different associated risks. The most similar worlds in which the bomb detonates in (Bomb 1) are much closer to actuality than the most similar worlds in which the bomb detonates in

(Bomb 2). Accordingly, the risk of the bomb detonating is higher in (Bomb 1) than (Bomb 2), despite the probabilities being the same in both cases.

While Pritchard doesn't make this point, it is interesting to observe that the modal account does validate checklist reasoning. Any world in which $P \vee Q$ is true is either a world in which P is true or a world in which Q true. As such, the most similar worlds in which $P \vee Q$ is true will either be the most similar worlds in which P is true, or the most similar worlds in which Q is true, or a mixture of the two.¹² Suppose again that I throw a fair die. The most similar worlds in which I throw a 5 or a 6 cannot be more similar than the most similar worlds in which I throw a 5 and more similar than the most similar worlds in which I throw a 6. After all, any world in which I throw either a 5 or a 6 must be either a world in which I throw a 5 or a world in which I throw a 6. On the modal account, therefore, the risk of $P \vee Q$ cannot be higher than the risk of P and the risk of Q —it will, in fact, be equal to the higher of these risks. Hence if, on the modal account, the risk of P is low and the risk of Q is low, it follows that the risk of $P \vee Q$ is low—just as required by checklist reasoning.

Although it has a number of intriguing features, the modal account of risk faces a difficulty which raises questions as to whether it can play one of the fundamental roles that we require from a notion of risk. One thing we can immediately observe is that no world counts as more similar to the actual world than it is to *itself*. It follows that, on the modal account, any event which *actually happens* must be at maximally high risk of happening, and any state of affairs which *actually obtains* must be at maximally high risk of obtaining. To put the point slightly differently, if P is true, it follows that P could easily be true—it is contradictory to say “ P is true but P couldn't easily be true”. If P is true, then there is a maximally similar world—the actual world—at which P is true, in which case, according to the modal account, there is then a maximal risk of P .

This, we submit, can give rise to a potentially serious concern about the modal notion of risk. Suppose one is about to drill into a wall in a West Australian house built in the 1970s, and is wondering about the risk that the wall contains asbestos. On the modal account, if the wall really does contain asbestos, then the risk is maximally high. In this case, there is a maximally similar world—the actual world—in which the wall contains asbestos. If, on the other hand, the wall does not contain asbestos, then, according to the modal account, the risk will be lower—the closest worlds in which this is true will be somewhat distant from actuality, depending upon further facts of the case. In any event, on the modal account it seems that one cannot make a judgment about the *risk* that the wall contains asbestos without taking a view as to whether it *does* contain asbestos.

Similarly, consider again the example of jurors in a criminal trial who are contemplating a guilty verdict and reasoning about the risk that the defendant is innocent. Suppose the prosecution produced two independent eyewitnesses to the crime who were willing to identify the defendant as the culprit. What is the risk that the defendant is innocent—that he did not commit the crime—given the testimonial evidence? If the defendant did commit the crime, and the eyewitnesses are reliable and truthful, then the modal account predicts that

¹² If we impose what Lewis termed the *limit assumption* for world similarity, stating that any set of worlds must have maximally similar members, then each world could be assigned a similarity ‘rank’—the most similar worlds will be assigned rank 0, the next most similar assigned rank 1 etc. (see Lewis, 1973, Smith, 2016, chap. 8) A proposition P could also be assigned a similarity rank $s(P)$ according to the rank of the most similar worlds in which it is true. A lower rank will, on the modal account, indicate a higher risk.

the risk is low. In this case, the most similar worlds in which the defendant is innocent will be very dissimilar—not only will the facts of the case be different in these worlds, but the eyewitnesses will be lying or mistaken. What if, on the other hand, the defendant *is* in fact innocent, and the eyewitnesses *are* lying or mistaken? In this case, on the modal account, the risk that the defendant is innocent given the testimony, is maximally high. After all, in this case, the most similar world in which the defendant is innocent is the *actual* world.¹³

Again, on the modal account, it seems that one cannot assess the risk that the defendant is innocent without already taking a stand on whether he is innocent or guilty: if he is guilty, the risk is low and if he is innocent the risk is maximally high. This seems to be of little help when it comes to actually making a decision (for related discussion see Smith, 2018, pp. 1204-1205).¹⁴

Importantly, this issue also affects Pritchard's reasoning about the Bomb scenarios above. His reasoning about these scenarios effectively takes it for granted that the triggering conditions in (Bomb2) are not met—namely, that the weakest horse in the Grand National *won't* win the race by at least ten furlongs, and the worst team in the FA Cup draw *won't* beat the best team by at least ten goals, and the Queen *won't* choose to speak a complete sentence of Polish during her next public speech. If we were in a world in which the triggering conditions in (Bomb 2) are met, the modal account would predict that the risk in (Bomb 2) is at least as high as the risk in (Bomb 1) or even higher.

A proponent of the modal account could point out that there are some cases in which one can judge a proposition P to be *high risk* without the need to take a view on whether P is true. In (Bomb 1), for instance, the modal account does allow us to judge that there is a high risk that, say, (14, 6, 32, 20, 12, 41) will be the winning numbers, while remaining neutral on whether they will be the winning numbers. Irrespective of whether these numbers come up in the actual world, we know that there is a very similar world in which they do. While this much should be granted, it remains the case that if the modal account is correct, then one cannot judge that there is a *low risk* of P without thereby presupposing that P is false.

Of course, even if P is actually true, one could still, on the modal account, be in an evidential position in which it is *reasonable to judge* that there is a low risk of P. That is, a defender of the modal account could insist that, even if the wall does contain asbestos or the defendant is innocent, one could still reasonably judge these propositions to be low risk. Importantly, however, on the modal account, these judgments could never be true—and this is in stark contrast to what the probabilistic account predicts. On the probabilistic account, the risk that the wall contains asbestos or the risk that the defendant is innocent, is purely a function of the available evidence, and does not depend on whether the wall actually contains asbestos, or whether the defendant is actually innocent. Now, while the jury is still out on whether the modal account can deal with this difficulty in a

¹³ Another possibility is that the eyewitnesses are lying or making some mistake, but the defendant is guilty nonetheless. Depending on just how the details are filled in, the modal account may predict that there is a somewhat higher risk of innocence in this case, as fewer facts would need to be altered in order to reach worlds at which the defendant is innocent.

¹⁴ According to Pritchard, the modal account of risk is the flipside of his modal account of luck, on which the luckiness of an event or state of affairs is gauged by the similarity of the most similar worlds in which it fails to obtain (see Pritchard, 2016, section II). When one assesses the luckiness of an event or state of affairs one typically presupposes that the event has occurred or the state of affairs obtains (in this sense risk judgments are importantly different to luck judgments). As such, the fact that the actual world counts as maximally similar poses less of a problem for a modal account of luck.

satisfying manner, we think there is an alternative notion of risk that can reproduce some of the benefits of the modal account without facing this particular problem.

5. The normic account of risk

Given a possible proposition, one question we can consider is how probable it is that the proposition is true. Another question we might consider is how easy or difficult it would be for that proposition to be true. As the preceding discussion illustrates, our answer to the second question is not determined by our answer to the first. We might judge that the bomb could more easily explode in (Bomb 1) than (Bomb 2) even though the probabilities of an explosion are equal. A third question that we might consider is how *normal* or *abnormal* it would be for a proposition to be true. It is somewhat natural to judge that it would be more abnormal for the bomb to detonate in (Bomb 2) than in (Bomb 1), even though the probabilities are equal. Although it might be highly unlikely for a particular set of 6 numbers—say (14, 6, 32, 20, 12, 41)—to come up in the next National Lottery draw, there is a sense in which there would be nothing *abnormal* about this occurrence. After all, some set of six numbers has to come up, and it may just as well be this as any other. In contrast, it would be highly abnormal for the weakest horse in the Grand National to win the race by at least ten furlongs, and the worst team in the FA Cup draw to beat the best team by at least ten goals, and the Queen to choose to speak a complete sentence of Polish during her next public speech.

But what exactly is this notion of normalcy which is being appealed to here? Sometimes when we describe an event as abnormal we mean only that it is relatively rare or infrequent. In this sense, of course, it would count as abnormal for (14, 6, 32, 20, 12, 41) to be the numbers that come up in the National Lottery draw (and any other outcome would count as abnormal too). Other times, though, when we describe a situation as abnormal, we are not merely making a claim about frequencies. If the lights in my house suddenly start to flicker, or my car fails to start when I turn my key in the ignition and I remark “that’s not normal”, I’m not simply pointing out that this is something infrequent—part of what I’m suggesting is that there needs to be some specific *explanation* for what is occurring. Perhaps there’s a problem with the electricity supply or the wiring in my house, perhaps the car battery is dead or has been disconnected, etc. Whatever the truth, there has to be *more to the story*—and this is what we flag in characterising these events as abnormal.

It may be highly unlikely for (14, 6, 32, 20, 12, 41) to be the numbers that come up in the National Lottery draw, but this wouldn’t require a special explanation of any kind—all the lottery outcomes are on an explanatory par. On the other hand, it’s natural to think that the triggering conditions in (Bomb 2)—the weakest horse winning the Grand National by at least ten furlongs, the worst team in the FA Cup draw beating the best team by at least ten goals, and the Queen choosing to speak a complete sentence of Polish during her next public speech, are each events that would require explanation.

As well as ordering possible worlds with respect to how similar they are, we may also order possible worlds with respect to how *normal* they are, with the normalcy of a proposition linked to the normalcy of the most normal worlds in which it is true (Smith, 2010, 2016). Amongst the most normal worlds in which the lottery takes place will be worlds in which each possible set of numbers comes up, including those that will trigger an explosion in (Bomb 1). In contrast, the most normal worlds in which the trigger conditions of (Bomb 2) are met are highly abnormal.

According to the *normic* account of risk, the risk of a proposition P is determined by the normalcy of the most normal worlds in which P is true—the more normal these worlds, the higher the risk. On this view, the risk of P is higher than the risk of Q just in case the most normal worlds in which P is true are more normal than the most normal worlds in which Q is true. Once again, this world ordering should be restricted to worlds in which the background evidence holds, and should exclude worlds which are inconsistent with the background evidence. According to the normic account, the risk of the bomb detonating is higher in (Bomb 1) than (Bomb 2), despite the probabilities being the same in both cases. The most normal worlds in which the bomb detonates in (Bomb 1) are more normal than the most normal worlds in which the bomb detonates in (Bomb 2).

While the normic account and the modal account may make the same prediction about the bomb examples, there are a number of significant differences between the two accounts. First and foremost, while the actual world must be maximal in a world similarity ordering, the same is not true of a world normalcy ordering. While the actual world must count as maximally similar to itself, it will not count as maximally normal—after all, the actual world is witness to any number of abnormal events and states of affairs. Put differently, truth does not entail normalcy and there is nothing contradictory in saying ‘P is true but P wouldn’t normally be true’.

As a result, like the probabilistic account, the normic account of risk makes it perfectly possible to assess the risk of a proposition without taking a prior view on its truth. Unlike the probabilistic account, however, the normic account validates checklist reasoning. Any world in which $P \vee Q$ is true is either a world in which P is true or a world in which Q is true. As such, the most normal worlds in which $P \vee Q$ is true will either be the most normal worlds in which P is true, or the most normal worlds in which Q is true, or a mixture of the two.¹⁵ The most normal worlds in which I throw a 5 or a 6 cannot be more normal than the most normal worlds in which I throw a 5 and more normal than the most normal worlds in which I throw a 6. Any world in which I throw either a 5 or a 6 must be either a world in which I throw a 5 or a world in which I throw a 6. On the normic account, the risk of $P \vee Q$ will be equal to the maximum of the risk of P and the risk of Q. Hence, on the normic account, if the risk of P is low and the risk of Q is low, it follows that the risk of $P \vee Q$ is low. The validation of checklist reasoning is a by-product of linking the risk of a proposition P to the highest position, in an ordering, of worlds in which P is true. Making the actual world maximal in this ordering is irrelevant for this purpose.

Previously, we alluded to the availability heuristic as a potential explanation for why a subject might judge the risk in (Bomb 1) to be greater than that in (Bomb 2). Originally, a subject was said to use the heuristic “whenever he estimates the frequency or probability by the ease with which instances or associations come to mind.” (Tversky & Kahneman 1973, p.208) The description of this heuristic has since become somewhat broader (Schwarz & Vaughn 2002) and is often taken to encompass “ease of imagination” as a

¹⁵ If we impose a limit assumption for world normalcy, then each world could be assigned a normalcy rank – the most normal worlds receiving rank 0, the next most normal receiving rank 1 and so on (Smith, 2018, chap. 8). A proposition P could also be assigned a normalcy rank $n(P)$ according to the rank of the most normal worlds in which it is true. A lower normalcy rank will indicate, on the normic account, a higher risk.

basis for judging risk—another phenomenon that has been experimentally demonstrated (Sherman, et.al. 1985).

Importantly, one who uses ease of imagining as a way of making risk judgments might be expected to conform to checklist reasoning. In order to imagine a situation in which $P \vee Q$ holds one must either imagine a situation in which P holds or a situation in which Q holds. The ease with which one can imagine $P \vee Q$ would in general be equal to the ease with which one can imagine P or the ease with which one can imagine Q —whichever one finds easier. If one judges that P is low risk, on the grounds that it is difficult to imagine, and judges that Q is low risk, on the grounds that it is difficult to imagine, one would be expected to make the very same judgment about $P \vee Q$.

It is *prima facie* plausible that there is a connection between the amount of explanation that an event or state of affairs would require, and the ease with which it can be imagined. What makes it comparatively difficult to imagine the weakest horse in the Grand National winning by at least 10 furlongs, and the worst team in the FA cup draw beating the best team by at least 10 goals, and the Queen choosing to speak a complete sentence of Polish during her next public speech is that these scenarios require some explanation and, as such, there is pressure to imagine a rich accompanying narrative in order to make them intelligible. In contrast, no explanation is needed in order for six particular numbers to come up in the next national lottery draw, and no imagined narrative is needed to supplement our imagination of this event—as Pritchard might say, we need only imagine a few coloured balls falling in a particular configuration.

Indeed, some psychologists have drawn an explicit link between the ease with which a hypothetical event or state of affairs can be imagined and the ease with which it can be explained (Sherman et al. 1985).¹⁶ In addition, there is some evidence to suggest that asking subjects to imagine a hypothetical event or state of affairs and asking subjects to offer possible explanations for a hypothetical event or state of affairs have a similar effect upon their subsequent judgments about its probability or risk (Koehler, 1991).¹⁷ This raises the possibility of rethinking the role of the availability heuristic with respect to risk judgments. Rather than regarding it as an expedient, but potentially misleading, method for assessing probabilistic risk, ease of imagining could be seen as the *canonical* method for assessing risk of a different kind—namely, normic risk. We will return to this in the next section.

We have argued that the normic account of risk is able to emulate the purported advantages of the modal account while avoiding what we take to be its most serious difficulty. The normic account is not free of difficulties however, and we conclude this section by highlighting one—a problem that it shares with the modal account. The cases described in Pritchard's thought experiment demonstrate that the modal and the normic accounts of risk can discriminate between the riskiness of propositions that the

¹⁶ Others have drawn a more general link between the availability of a hypothetical scenario and its normalcy, construed in terms of its conformity to norms or routines (Kahneman & Miller, 1986). How this conception of normalcy relates to the present conception is something we leave for possible future work.

¹⁷ We can distinguish between imagining an event or state of affairs in the sense of simply picturing or visualising it, and imagining an event or state of affairs in the sense of envisaging how and why it could occur or obtain. It is only ease of imagining in the second sense which is linked to the ease with which a hypothetical event could be explained. Interestingly, there is some evidence to suggest that it is only ease of imagining in the second sense which is involved in the availability heuristic, with ease of sheer visualisation having no discernible effect on probability or risk judgments (Levi & Pryor, 1987).

probabilistic account reckons to be equally risky. But it is also possible for the probabilistic account to discriminate between the riskiness of propositions that the normic and modal accounts reckon to be equally risky. Consider the following new scenario:

(Bomb 3) An evil scientist has rigged up a large bomb, which he has hidden in a populated area. If the bomb explodes, many people will die. There is no way of discovering the bomb before the time it is set to detonate. The bomb will only detonate, however, if six specific numbers between 1 and 99 come up on the next national lottery draw. The odds of these six numbers appearing is roughly one billion to one. It is not possible to interfere with this lottery draw.

In (Bomb 3) the odds are very strongly against the bomb detonating—more strongly than in (Bomb 1). As such, the probabilistic account predicts that (Bomb 1) is the riskier scenario. It is unclear, however, whether either the modal or the normic account can accommodate this intuitive difference in risk. Pritchard's comments about (Bomb 1) would seem to carry over to (Bomb 3)—all that is required for the bomb to detonate is for six coloured balls to fall in a particular configuration. Any lottery outcome could obtain as easily as any other. Amongst the most similar worlds in which the lottery is run will be worlds in which each outcome obtains, including those which feature the six specified numbers. According to the modal account, the risk of an explosion in (Bomb 3) should be equal to the risk of an explosion in (Bomb 1).¹⁸

The normic account makes the same prediction. For six particular numbers to come up in a lottery draw requires no special explanation. All of the lottery outcomes are on an explanatory par. Amongst the most normal worlds in which the (Bomb 3) lottery is run will be worlds in which each outcome obtains, including those which feature the six specified numbers. Since the modal and normic accounts predict that (Bomb 3) and (Bomb 1) are equally risky, they will also predict that (Bomb 3) is *riskier* than (Bomb 2), even though the probability of the bomb detonating in (Bomb 3) is significantly lower. In this case, it seems questionable whether the prediction made by the modal and normic accounts is the correct one.¹⁹ Whatever we might take its benefits to be, we think it is clear that the normic account is incapable of providing everything that we require from a notion of risk.

6. Varieties of risk

We have set out three competing accounts of risk—the probabilistic, modal and normic accounts. All three of these accounts represent what we might term *monist* conceptions

¹⁸ We could alter the triggering condition in (Bomb 3) in such a way that only five numbers are involved in the lottery. This would leave the probability of an explosion lower in (Bomb 3) than (Bomb 1) but would arguably, bring it modally closer—as only five balls would need to fall at a particular time in order for the bomb to detonate. In this case the modal account (though not the normic account) would make the even more counterintuitive prediction that (Bomb 3) is *riskier* than (Bomb 1) in spite of the lower probability of an explosion. We plan to discuss this and other decision-theoretic concerns for the modal and normic accounts in further work.

¹⁹ Even if one wished to dispute whether the triggering conditions in (Bomb 2) are more modally distant or more abnormal than the triggering conditions in (Bomb 3), the intended structure of this example is clearly a possible one. As Pritchard points out, there are events with arbitrarily low probabilities that occur at very similar worlds (Pritchard, 2015, section 4)—and the same is true of maximally normal worlds. As such, there are events with arbitrarily low probabilities that the modal account and normic account will reckon to be more risky than the triggering conditions in (Bomb 2).

of risk. That is, all three accounts agree that there is one correct notion of risk and disagree on how best to analyse or characterise that notion.

Given the monist assumption, a defender of the probabilistic account is forced to appeal to a range of heuristics and biases in order to explain why intuitive judgments can *systematically* deviate from those that the account predicts to be correct. If someone uses, say, ease of imagining as a way of making a risk judgment, this is portrayed as a kind of mental short-cut, which is employed because it can, under the right circumstances, serve as a reliable guide to risk, probabilistically construed. When this method leads to risk judgments that are at variance with what the probabilistic account predicts, this is taken to show that the method can *systematically* mislead. When confronted with the practice of using checklist reasoning, a defender of the probabilistic account of risk must adopt either a revisionary approach, arguing that the relevant practices should be changed and the reasoning abandoned, or attempt to explain, from within the confines of the probabilistic perspective, why the use of such reasoning may be useful or acceptable, though strictly invalid.²⁰

Similar issues arise for other monist views. A proponent of the normic account may highlight the fact that it validates checklist reasoning and is able to legitimise the practices that employ it. Additionally, she may be able to account for some intuitive differences in risk that cannot be accommodated from a probabilistic perspective. And yet, as discussed in the previous section, the normic account is also insensitive to certain intuitive differences in risk, such as that between (Bomb 1) and (Bomb 3), which can be straightforwardly captured by the probabilistic account.

Rather than attempting to weigh the costs and benefits of competing monist accounts of risk, an intriguing alternative is to embrace a *pluralist* view on which there are multiple, equally legitimate, notions of risk. While we won't attempt a full elaboration or defence of risk pluralism, in what follows we sketch the bare bones of such a view and note possible avenues for future research. On the kind of pluralist view that we wish to consider, the probabilistic and normic accounts represent two distinct precisifications of our ordinary, intuitive notion of risk, each emphasising or picking up on different aspects of that notion. While we remain open to the possibility of further viable notions of risk, it is not our intent that such notions proliferate. In our view, the preceding criticisms of the modal notion of risk show, for instance, that it cannot serve as a viable precisification of our intuitive notion.²¹

A pluralist view of this kind invites a rethinking of the status of certain heuristics and biases involved in making risk judgments. On a monist view, if we observe two different conflicting methods for making risk judgments, we are forced to conclude that one of these (at least) represents a heuristic which is leading us astray. On a pluralist view, while it is still possible for risk judgments to be affected by heuristics and biases, we cannot automatically conclude that heuristics are at play whenever we

²⁰ So, for example, one could argue that in many cases of checklist reasoning the probability of the disjuncts is so low as to leave the probability of the disjunction below the threshold for low risk. Compare in this context Christensen (2004, p.79f) discussion of *the argument*.

²¹ Pluralism about risk is also suggested by Bricker (2018) who argues that ordinary risk judgments can "track different metrics given different initial conditions" (p.207, n2) (for related discussion, see also Ebert & Robertson (2013, p.55f.)). This has an obvious affinity with the view that we suggest here, though Bricker, unlike us, appears to regard the modal account as providing one such metric and delivering one viable notion of risk. In this context, note that Chalmers (2011) suggests that pluralism about philosophical notions—and *risk* clearly seems one such notion—should be the default position.

observe conflicting methods for making such judgments—for it may be that the methods are tracking different kinds of risk, each of which can be extracted from our ordinary notion.

Consider again the availability heuristic. On the monist probabilist view, to use “ease of imagination” as a way of judging risk is to employ an expedient, but imperfect, method for tracking probabilities. Given a pluralist outlook, however, this can be seen instead as an appropriate method for tracking something non-probabilistic—normalcy—which is equally present in our ordinary notion of risk. This also opens up a new perspective on the factors which have been shown to exert an influence on risk judgments. For example, highlighting the vividness of a bad outcome has been portrayed as a source of error in risk judgments, on the grounds that it can lead people to ignore or downplay probabilistic information (Tversky & Kahneman 1974). On the pluralist picture, rather than being something which clouds our thinking about risk, vividness could simply serve to emphasise or foreground certain non-probabilistic aspects of the ordinary notion of risk.

Of course, there are a number of substantial questions about how, on the view we have outlined, we should think about the ordinary notion of risk and its relation to the probabilistic and normic notions. One model is to regard the ordinary notion of risk as a “defective” concept in the sense of Scharp (2013, 2018)—a concept constituted by various principles which lead to inconsistent and conflicting results. In this case, the probabilistic and normic notions of risk could be seen as consistent concepts which can be extracted from the ordinary notion, and which could fruitfully replace it for a range of purposes—in Scharp’s terms, they could be seen not as attempts at conceptual analysis, but as the results of conceptual *engineering* (see also Cappelen, 2018). As said, more research will be required properly to develop our preferred form of risk pluralism. We conclude by highlighting three further consequences our distinctive approach will have for current philosophical and psychological discussion regarding the notion of risk and the rationality of human risk judgments.

First, let’s revisit our own survey regarding Pritchard’s bomb scenarios. While a significant proportion judge the two scenarios as equally risky—in line with the probabilistic account—it is worth highlighting that roughly one in three subjects disagree and judge (Bomb 2) to be less risky than (Bomb 1) as predicted by the normic account. For a monist probabilist, one in three subjects exhibits a defective intuitive judgment which would need to be explained away using heuristics and biases. For the pluralist, on the other hand, it might be that some of these subjects are steered, by the very contents of the vignette, to reason along the lines predicted by a normic notion of risk. This raises the important question of how and when vignettes trigger such variation. Further research, controlling for alternative explanations of this phenomenon, promises to be an exciting research project. Ultimately, such research may help us to determine more reliably when a normic or probabilistic notion could offer a better prediction of human risk judgment.²²

²² There is some evidence, partly alluded to above, suggesting that fact finders in legal trials don’t engage in probabilistic reasoning but, rather, reach a verdict through comparing the plausibility of competing narratives (Pennington & Hastie, 1991, Allen & Jehl, 2004, Allen, 2008, Pardo & Allen, 2008, Allen, 2014, section 4)—a process which could potentially be understood in terms of the management of normic risk. We won’t pursue this further here. Certain pervasive judgments about lotteries—in particular, the judgment that one cannot *know* that a particular ticket will lose, or that a particular set of numbers won’t come up, purely on the basis of the odds involved—could be explained in terms of the fact that such a belief would be subject to a high normic (though low probabilistic) risk of error. (Ebert, Smith and Durbach, 2018)

Second, the pluralist approach could potentially be combined with the dual system approach to human judgment about risk. Consider the suggestion by Loewenstein et al. (2001) that there seem to be two systems involved in our risk judgment: an affective, non-probabilistic system which is distinct from a cognitive probabilistic system. One possibility is to merge risk pluralism with a dual system approach, by portraying normic risk and probabilistic risk as the notions tracked respectively by the affective and cognitive systems (which may in turn involve a re-evaluation of the affective system as partly cognitive). Whatever one thinks of the prospects for such a view, it is clear that research in psychology regarding the role of anticipatory emotions (i.e. immediate visceral reactions such as fear, anxiety and dread) when making risk judgments could potentially be informed by the role that normalcy plays in judgment about risk.

Lastly, our discussion offers a new take on the so-called “rationality wars” in psychology. In a number of publications, Gigerenzer and others have challenged Kahneman and Tversky’s attitude of characterising various biases and heuristics as errors or fallacies, rendering human judgment based on them irrational.²³ In contrast, Gigerenzer suggests that evaluating the performances of human judgment against the norm of an ideal probabilistic reasoner is often inappropriate, since it fails to take into account the boundedness of human judgment and rationality. We do not want to delve into the details of this dispute, but it is important to note that it links up with one of our primary concerns here—the question of whether a probabilistic norm supplies the “gold standard” for risk judgments.

The pluralist conception of risk offers a different perspective on questions about the rationality of risk judgments. According to the pluralist, there are different precisifications of our intuitive notion of risk, each of which have their own normative standards which will overlap in some cases but diverge in others. If subjects judge in a way that goes against the standards of, say, the probabilistic account, this doesn’t necessarily render their judgment irrational. After all, they may be judging in accordance with the standards appropriate for the normic notion of risk.

There may be a sense in which this simply pushes the normative question one step back. Even if we accept that there are competing normative standards associated with the different notions of risk, we might ask whether these notions are always equally legitimate, or whether there are contexts in which only one of them should be regarded as the correct norm of reasoning and judgment. By adopting a pluralist view, we open up the possibility of distinguishing two kinds of mistakes that are relevant to the debate about the rationality of human risk judgments: the *normative* mistake of failing to meet the demands of a given standard, and the *meta-normative* mistake of adhering to the wrong standard in a given context. Addressing the meta-normative issue of which standard ought to be in play, in a given context, will have to be postponed to another occasion.²⁴

²³ To be fair, this is only one aspect of Gigerenzer’s criticism, which is wide-ranging (see Gigerenzer & Murray, 1987), (Gigerenzer 1991, 1994, 1996). For a response, see Kahneman & Tversky (1996). For critical discussion of the normative criticism alluded to above and a list of other concerns levelled against Kahneman, see Vranas (2000) with a response by Gigerenzer (2001) and Samuels et.al. (2002).

²⁴ We would like to thank the audiences of the staff research seminar in philosophy at the University of Stirling and the research seminar of the Behavioural Science Centre at the University of Stirling for their feedback. In addition, we are extremely grateful to the reviewers for their extensive critical and supportive comments on earlier versions of this manuscript. Finally, we would like to acknowledge funding from to the University of Stirling APC fund to cover the costs for open access publishing.

Appendix: Survey recruitment, design, and results

We recruited participants for our surveys using *prolific.ac* which is suggested to be preferable to other platforms (Peer et.al 2017; Palan & Schitter 2018). Participants were paid £0.30 per survey, resulting in an average pay of £9 per hour and so well above minimum wage in the UK. Given the specific content of the survey, we restricted participation to native English speakers and UK nationals. Participants were asked for personal information (gender, age, university degree, philosophy background). In total, we recruited 106 participants (72 female, 33, male, 1 other) and we deleted one participant given his/her background in philosophy. We received ethical approval from the University of Stirling (GUEP355) for this study.

We used the following vignette which is almost identical to the text in Pritchard (2015):

Background

An evil scientist has rigged up a large bomb, which he has hidden in a populated area. If the bomb explodes, many people will die. There is no way of discovering the bomb before the time it is set to detonate.

Scenario 1.

The bomb will only detonate if a certain set of numbers comes up in a lottery draw. The odds of these numbers appearing is 1 in 14 Million. It is not possible to interfere with this lottery draw.

Scenario 2.

The bomb will only detonate if a series of three highly unlikely events obtains. First, the weakest horse in the field at the Grand National, Lucky Loser, must win the race by at least ten furlongs. Second, the worst team remaining in the FA Cup draw, Accrington Stanley, must beat the best team remaining, Manchester United, by at least ten goals. And third, the queen of England must spontaneously choose to speak a complete sentence of Polish during her next public speech. The odds of this chain of events occurring are 1 in 14 million. It is not possible to interfere with the outcomes of any of the events in this chain.

We then asked participants about their *choice preference*:

Assuming you are amongst the population possibly affected by the bomb, do you:

- a. prefer scenario 1 over scenario 2.
- b. prefer scenario 2 over scenario 1.
- c. have no clear preferences.

We then asked them about their *risk judgment* (both questions were visible to the subjects):

Do you judge that:

- a. there is an equal risk of the bomb detonating in both scenarios.
- b. there is a higher risk of the bomb detonating in scenario 1 than scenario 2.
- c. there is a higher risk of the bomb detonating in scenario 2 than in scenario 1.




Figure 2. Survey Response of 105 subjects listing the number of respondents judging preference for (Scenario 1 over Scenario 2), preference for (Scenario 2 over Scenario 1), or judging (No Preference).

The results of our survey are as follows: (Figure 2) lists responses with respect to the preference question.

While both scenarios have the same probability of occurring, scenario 2 is clearly preferred over scenario 1 or the no-preference option. Expressing a preference for scenario 2 was significantly more popular than the other two options ($\chi^2 = 20.8$, $df = 2$, $p\text{-value} < 0.001$).

However, when it comes to risk judgments (Figure 3), subjects tended to judge both scenarios to be equally risky. This judgment was significantly more popular than the other options ($\chi^2 = 37.2$, $df = 2$, $p\text{-value} < 0.001$).

We assessed people's behaviour with respect to both their preference and risk judgments (Figure 4). In our survey, we presented the preference and the risk question on the same page and, while the preference question was listed before the risk question, both questions were transparent to the subjects while responding to either. Assessing a subject's risk judgment in light of their preference judgment offers some intriguing results. Subjects who expressed *no preference* for either scenario were most likely to go with the accompanying risk assessment and judge the two scenarios as *equally risky* (32 out of 33




Figure 3. Survey Response of 105 subjects listing the number of respondents judging (Scenario 1 riskier than Scenario 2), (Scenario 2 riskier than Scenario 1), or judging (Equal Risk)

Risk judgments given preference judgments




Figure 4. Survey Response of 105 subjects representing their risk judgments given the individual’s prior preference judgment.

subjects did so). Strikingly, we do not find this pattern for subjects who expressed a preference for one of the scenarios. Roughly 47% and 38% of those subjects who express a preference for Scenario 1 and Scenario 2 respectively, *still* judged the scenarios to be equally risky.

This observation challenges the assumption—arguably implicit in Pritchard’s discussion—that in these cases risk judgments are determined or predicted by a subject’s preference. Focusing on the subjects who express a preference for scenario 2 (55 in total), one would expect, if this assumption were correct, to find a statistically significant difference between the number of people who made the accompanying risk judgment that scenario 1 is more risky, and the number who judged the scenarios to be equally risky. This, however, is not the case: $\chi^2 = 2.6667$, $df = 1$, $p\text{-value} = 0.1025$. Similar considerations apply to subjects who express a preference for scenario 1 (17 in total): $\chi^2 = 0$, $df = 1$, $p\text{-value} = 1$. However, since absence of statistical evidence does not imply the absence of an effect, particularly with small sample sizes, caution is required when interpreting these results.

It is worth highlighting that these findings complement observations made by Heath & Tversky (1991) who showed that subjects with certain competences who judge scenarios to be equally likely, may still express a distinct preference for a scenario. Further research is planned to investigate why preference judgments and risk judgments appear mismatched in some cases but not in others, and what other possible determinants there are for risk and preference judgments.

Bibliography

Allen, Ronald J. (2008). “Explanationism all the way down.” *Episteme*, 5(3), 320–328.
— (2014) “Burdens of proof.” *Law, Probability and Risk*, 13(3-4), 195–219.
— and Sarah Jehl. (2004). “Burdens of persuasion in civil cases: Algorithms vs explanations.” *Michigan State Law Review*, 893–944.

- Boholm, Max, Niklas Möller, and Sven Ove Hansson. (2016) "The concepts of risk, safety and security: Applications in everyday language." *Risk Analysis*, 36(2), 320–338.
- Bricker, Adam M. (2018) "Do judgements about risk track modal ordering?" *Thought*, 7 (3), 200–208.
- Cappelen, Herman. (2018). *Fixing Language*. Oxford: Oxford University Press.
- Chalmers, David J. (2011). "Verbal disputes." *Philosophical Review*, 120(4), 515–566.
- Comar, Cyril. L. (1979). "Risk: a pragmatic de minimis approach." *Science*, 203(4378), 319.
- Christensen, David. (2004). *Putting Logic In Its Place*. Oxford: Oxford University Press.
- Ebert, Philip A. and Simon Robertson. (2013). "A plea for risk." *Royal Institute of Philosophy Supplement*, 73, 45–64.
- , Martin Smith and Ian Durbach. (2018). "Lottery judgments: A philosophical and experimental study." *Philosophical Psychology* 31(1), 110–138.
- Einhorn, Hillel J. and Robin M Hogarth. (1985). "Ambiguity and uncertainty in probabilistic inference." *Psychological Review*, 92(4), 433–461.
- Ellsberg, Daniel (1961). "Risk, ambiguity axioms, and the Savage axioms." *Quarterly Journal of Economics*, 75, 643–669.
- Epstein, Seymour. (1994). "Integration of the cognitive and the psychodynamic unconscious." *American Psychologist*, 49, 709–724.
- Evans, Johnathan S. B. T. (2011). "Dual-process theories of reasoning: Contemporary issues and developmental applications." *Developmental Review*, 31(2-3), 86–102.
- Finucane, Melissa L., Ali Alhakami, Paul Slovic, and Stephen M. Johnson. (2000). "The affect heuristic in judgments of risks and benefits." *Journal of Behavioral Decision Making*, 13, 1–17.
- Gärdenfors, Peter and Nils-Eric Sahlin. (1982). "Unreliable probabilities, risk taking and decision making", *Synthese*, 53, 361–386.
- Gigerenzer, Gerd. (1991). "How to make cognitive illusions disappear: Beyond 'heuristics and biases'." In W. Stroebe & M. Hewstone (Eds.), *European Review of Social Psychology*. Chichester, England: Wiley. 83–115.
- . (1994). "Why the distinction between single-event probabilities and frequencies is relevant for psychology and vice versa." In G. Wright & P. Avion (Eds.), *Subjective probability*. New York: Wiley. 129–162.
- . (1996). "On narrow norms and vague heuristics: a reply to Kahneman and Tversky (1996)." *Psychological Review*, 103(3), 592–596.
- . (2001). "Content-blind norms, no norms, or good norms? A reply to Vranas." *Cognition*, 81(1), 93–103.
- and David J. Murray. (1987). *Cognition as Intuitive Statistics*. Hillsdale, NJ: Erlbaum.
- Goldsmith, Robert W. and Nils-Eric Sahlin. (1983). "Second-order probabilities and risk in decision making." *Advances in Psychology*, 14, 455–467.
- Heath, Chip and Amos Tversky. (1991). "Preferences and beliefs: Ambiguity and competence in choice under uncertainty." *Journal of Risk and Uncertainty*, 4(1), 5–28.
- Hubbard, Douglas. W. (2009). *The Failure of Risk Management*. Hoboken, New Jersey: John Wiley & Sons.

- Kahneman, Daniel. (2011). *Thinking Fast and Slow*. New York: Farrar, Strauss and Giroux.
- and David T. Miller. (1986). “Norm theory. Comparing reality to its alternatives.” *Psychological Review*, 93(2), 136–153.
- and Amos Tversky. (1973). “On the psychology of prediction.” *Psychological Review*, 80, 237–251.
- . (1979). “Prospect theory: An analysis of decision under risk.” *Econometrica*, 47(2), 263–292.
- . (1974). “Simulation heuristic.” In D. Kahneman, P. Slovic, and A. Tversky (Eds), *Judgement Under Uncertainty: Heuristics and Biases*. Cambridge, UK: Cambridge University Press. 201–208.
- . (1996). “On the reality of cognitive illusions: A reply to Gigerenzer’s critique.” *Psychological Review*, 103, 582–591.
- Knight, Frank. H. (1921). *Risk, Uncertainty, and Profit*. Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin Company.
- Koehler, Derek. (1991). “Explanation, imagination, and confidence in judgment.” *Psychological Bulletin*, 110(3), 499–513.
- Levi, Ariel and John B. Pryor (1987). The use of the availability heuristic in probability estimates of future events: The effect of imagining outcomes versus imagining reasons. *Organizational Behaviour and Human Decision Processes*, 40, 219–234.
- Lewis, David. (1973). *Counterfactuals*. Oxford: Blackwell Publishers and Cambridge, MA: Harvard University Press. Reprinted with revisions, 1986.
- (1979). “Counterfactual dependence and time’s arrow.” *Noûs*, 13(4), 455–476.
- Loewenstein, George, Elke U. Weber, Christopher K. Hsee, and Ned Welch. (2001). “Risk as feeling.” *Psychological Bulletin*, 127(2), 267–286.
- Mumpower, Jeryl. (1986). “An Analysis of the *de minimis* strategy for risk management.” *Risk Analysis*, 6(4), 437–446.
- Pardo, Michael S., and Ronald J. Allen. (2008). “Juridical proof and best explanation.” *Law and Philosophy*, 27(3), 223–268.
- Palan, Stefan and Christian Schitter. (2018). “Prolific.ac – a subject pool for online experiments.” *Journal of Behavioural and Experimental Finance*, 17, 22–27.
- Peer, Eyal, Laura Brandimarte, Ssonam Samat, and Allesandro Acquisti. (2017). “Beyond the Turk: Alternative platforms for crowdsourcing behavioral research.” *Journal of Experimental Social Psychology*, 70, 153–163.
- Peterson, Martin. (2002). “What is a *de minimis* risk?” *Risk Management*, 4(2), 47–55.
- Pennington, Nancy and Reid Hastie. (1991). “A cognitive theory of juror decision making: The story model.” *Cardozo Law Review* 13, 519–557.
- Pritchard, Duncan. (2015). “Risk.” *Metaphilosophy*, 46(3), 436–461.
- . (2016). “Epistemic risk.” *Journal of Philosophy*, 113(11), 550–571.
- Open Group Standard (2013). “Risk Taxonomy (O-RT)”, Version 2.0, C13K, *The Open Group*. www.opengroup.org/bookstore/catalog/c13k.htm
- Rottenstreich, Yuval and Christopher K. Hsee. (2001). “Money, kisses and electric shocks: On the affective psychology of probability weighting.” *Psychological Science*, 12(3), 185–190.
- Samuels, Richard., Stephen Stich, and Michael Bishop. (2012). “Ending the rationality wars.” In R. Elio (Ed.), *Common Sense, Reasoning and Rationality*. New York: Oxford University Press. 236–268.

- Scharp, Kevin. (2013). *Replacing Truth*. Oxford: Oxford University Press.
- . (2018). “Philosophy as the study of defective concepts.” In A. Burgess, H. Cappelen & D. Plunkett (Eds.), *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.
- Schwarz, Norbert and Leigh Ann Vaughn. (2002). “The availability heuristic revisited: Ease of recall and content of recall as distinct sources of information.” In T. Gilovich, D. Griffin & D. Kahneman (Eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment*. New York: Cambridge University Press. 103–119.
- Sherman, Steven J, Robert B. Cialdini, Donna F. Schwartzman, and Kim D. Reynolds. (1985). “Imagining can heighten or lower the perceived likelihood of contracting a disease: The mediating effect of ease of imagery.” *Personality and Social Psychology Bulletin*, 11(1), 118–127.
- Sloman, Steven A. (1996). “The empirical case for two systems of reasoning.” *Psychological Bulletin*, 119(1), 3–22.
- Slovic, Paul, Baruch Fischhoff, and Sarah Lichtenstein. (1980). “Facts and fears: Understanding perceived risk.” In R.C. Schwin & W.A. Albers (Eds.), *Societal Risk Assessment: How safe is safe enough?* New York: Plenum Press. 181–212.
- . (1980). “The psychometric study of risk perception.” In V.T. Covello, J. Menkes & J. Mumpower (Eds.), *Risk evaluation and management*. New York: Plenum Press. 3–24.
- , Melissa L. Finucane, Ellen Peters, and Donald G. MacGregor. (2002). “The affect heuristic.” In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment*. New York: Cambridge University Press. 397–420.
- . (2004). “Risk as analysis and risk as feelings: Some thoughts about affect, reason, risk and rationality.” *Risk Analysis*, 24(2), 311–322.
- Smith, Martin. (2010). “What else justification could be.” *Noûs*, 41(1), 10–31.
- . (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford: Oxford University Press.
- . (2018). “When does evidence suffice for conviction?” *Mind*, 127(508), 1193–1218.
- Spottswood, Mark. (2016). “Unraveling the conjunction paradox.” *Law, Probability and Risk*, 15, 259–296.
- Tversky, Amos and Daniel Kahneman. (1974). “Judgment under uncertainty: Heuristics and biases.” *Science, New Series*, 185(4157), 1124–1131.
- . (1992). “Advances in prospect theory: Cumulative representation of uncertainty”. *Journal of Risk and Uncertainty*. 5(4), 297–323.
- Vranas, Peter B. M. (2000). “Gigerenzer’s normative critique of Kahneman and Tversky.” *Cognition*, 76(3), 179–193.