The Philosophical Quarterly Vol. 0, No. 0 *2024*
ISSN 0031-8094 https://doi.org/10.1093/pq/pqae131

A unified theory of risk*

Jaakko Hirvelä and Niall J. Paterson

University of Helsinki, Finland

A novel theory of comparative risk is developed and defended. Extant theories are criticized for failing the tests of extensional and formal adequacy. A unified diagnosis is proposed: extant theories consider risk to be a univariable function, but risk is a multivariate function. According to the theory proposed, which we call the unified theory of risk, the riskiness of a proposition is a function of both the proportion and the modal closeness of the possible worlds at which the proposition holds. It is argued the theory passes the tests of extensional and formal adequacy.

Keywords: risk; modality; epistemic risk; modal theory of risk; normic theory of risk.

I. Introduction

Risk judgements play a pivotal role in human societies. In the case of practical action and deliberation it is often crucial to know what kind of risks one's actions or decisions incur. Risk judgements also feature in an important way in normative practices. Often it is proper to blame someone for imposing unnecessary risks on us, such as transporting hazardous materials without taking proper safety measures, even if nothing bad happens. Many think that knowledge is incompatible with epistemic risk, since to know something one must be safe from error (Sosa 1999; Williamson 2000; Pritchard 2005; Hirvelä 2019, 2022).[1] The common ground in the literature on risk is that risks are unwanted events that have some chance of occurring (Pritchard 2015; Gardiner 2020; Hansson 2023). But what is this notion of chance?

Correspondence to: Jaakko Hirvelä, jaakko.hirvela@helsinki.fi
*The authors are listed in order of contribution.

[1] For criticism of one or more of these views, see Paterson (2020, 2022) and Hirvelä & Paterson (2021).

The prevalent theory of risk is probabilistic. Recently, this theory has been criticized for being extensionally inadequate, and rival theories have been developed and defended (Pritchard 2015, 2016; Ebert, Smith, & Durbach 2020). But the theories offered to replace the probabilistic theory are themselves extensionally inadequate and have undesirable formal properties. The aim of this paper is to offer a novel theory of risk that is both extensionally adequate and does not have those undesirable structural properties. On the theory proposed, the riskiness of a proposition is a function of both (a) the modal closeness of the worlds where the proposition holds and (b) the proportion of the worlds in which it holds throughout the space of possible worlds. Because it accommodates the intuitively correct aspects of the extant theories, we call it the *unified theory of risk*.

Before we begin, two points are worth noting. *First*, following Ebert et al. (2020), we assume risk to be primarily attributable to propositions, not events. We understand statements of the form 'the risk of *e* is *x*' to be elliptical for statements of the form 'the risk that *e* occurs is *x*'. This is because risk judgements range over outputs of the Boolean operations. It is sensible to ask whether there is a high risk of either a hurricane *or* a volcanic eruption. If risk is primarily attributable to events, it would not be sensible to attribute risk to the outputs of the Boolean operations, because events do not serve as inputs to those operations. But such attributions are sensible. So risk should be primarily understood as a relation between propositions, and only derivatively between events. For ease of exposition we will sometimes talk of the risk of events rather than propositions.[2]

Secondly, the unified theory, like the other theories we will discuss, does not seek to accommodate the intuitive view that a proposition with greater disutility is more risky, *ceteris paribus,* than a proposition with lesser disutility.[3] A bet that could lose all you hold most dear is more risky, all else being equal, than a bet that could lose a pair of socks. We idealize away from the value-related features of risk in part because our primary interest is to pin down the structural properties that risk has due to its relation to 'chance', and in part because we think that there are good arguments against the idea that the disutility of a proposition could increase its risk value.[4] That being said, we do think that only unwanted events are candidates for being risky, and in Section IV we will spell out how the unified theory can be modified to account for the effect disutility intuitively has on risk. To simplify the discussion, we treat all events as being equally disadvantageous in what follows.

[2]Navarro (2019) has argued that risk judgements are essentially forward looking, in that they refer to future events. We set this complication aside. The theory we propose can accommodate Navarro's idea by restricting the domain of quantification to propositions that hold in the future.

[3]Buchak (2013) and Gardiner (2020) hold that the disutility of an event increases its riskiness.

[4]See Ebert et al. (2020: 434) for an argument against the idea that the disutility of a proposition could increase its riskiness.

Here is the plan. In Section II, we introduce the probabilistic theory of risk and its recent critique. In Section III, we introduce the modal and normic theories of risk, and offer novel criticism of both. In Section IV, we offer a new account of risk and argue that it is superior to the extant proposals.

II. Risk as probability

According to *the probabilistic theory of risk*, the risk of a proposition *p* is just the probability of *p*. The relevant notion of probability may be unconditional, or it may be conditional, say on a body of evidence. We remain neutral between these formulations, the discussion that follows applies to both with equal force.

The probabilistic theory has many virtues. *First*, it assigns numerically precise risk values to propositions. *Secondly*, those risk values are well behaved when used as inputs to Boolean operations. Just as with probabilities, the risk of either a volcanic eruption or a hurricane is the sum of the risk of a volcanic eruption and the risk of a hurricane minus the risk of a volcanic eruption and a hurricane. *Thirdly*, just as we should expect, the probabilistic theory assigns all, and only necessary truths the maximal risk value, namely 1, and all, and only necessary falsehoods the minimal risk value, namely 0. *Fourthly*, it provides a straightforward, cardinal interpretation of the comparative risk relation. The risk of a proposition *p* is greater than the risk of *q* if, and only if, the probability of *p* is greater than the probability of *q*. Ranking propositions by their risk tells us not only which propositions are more risky than others, but also how much more risky they are. The amount by which one proposition is riskier than another is just the difference between their probabilities, that is, the greater probability minus the lesser. *Fifthly*, since the risk value distribution will conform to the axioms of the probability calculus, the theory can straightforwardly be implemented into decision theory. We take this to be a particularly attractive feature of the view, because risk plays a central role in decision theoretic frameworks.

Despite its virtues, the probabilistic theory has met growing resistance. This is for two reasons: the probabilistic theory does not underlie our ordinary practices that relate to risk (Redmayne 2008; Williamson 2009; Pardo 2018; Smith 2018, 2021; Ebert et al. 2020; Gardiner 2020; Moss 2022), and there seem to be cases in which the theory delivers wrong verdicts (Redmayne 2008; Pritchard 2015; Smith 2016). For instance, in a court of law naked statistical evidence of guilt is not sufficient to condemn the accused, even if the probability of the accused being guilty is high (Moss 2022; Pardo 2018; Smith 2018, 2021). This is puzzling if the probabilistic theory of risk is correct, since the court should condemn the accused if the risk of wrongful conviction is low enough. Moreover, there are cases where the probability of a wrong verdict is the same, but yet the court should find the defendant guilty or liable only in one of the cases, since the risk of wrongful conviction is not the same across the cases. This point is often made vivid with the following pair of cases.

BUS 1: A bus causes injury to a pedestrian, but no one sees which company the bus belongs to. 75% of the buses that operate on the street where the accident occurred are owned by the Blue Bus Company. No other evidence is available.

BUS 2: A bus causes injury to a pedestrian, who witnesses at court that she was hit by a Blue Bus Company bus. All bus companies have an equal share of traffic on the street. It's known that eyewitness testimony is only 75% reliable. No other evidence is available. (Redmayne 2008: 281)

In these cases, the evidential probability that a Blue Bus Company bus hit the pedestrian is the same, but intuitively the company can be found liable for the accident in light of the evidence only in one of the cases. Many think that the verdict of liability is correct only in one of the cases because there is a greater risk of false verdict of liability in the other case.[5] The probabilistic theory cannot accommodate this fact, since the evidential probability that the blue bus hit the pedestrian is the same across the cases. Though legal cases are much discussed in this connection, the fundamental problem seems to be that we cannot hold others *responsible* on the basis of purely probabilistic evidence, and hence actions like punishment that often serve to express blame, are not appropriate on such grounds either (Littlejohn 2020).

The inadequacy of the probabilistic theory of risk has also been challenged from a purely epistemic perspective. Many authors have recently argued that knowledge either is, or at least requires, a safe belief (Sosa 1999; Williamson 2000, 2009; Pritchard 2005, 2012, 2016; Lasonen-Aarnio 2010). These authors hold that knowledge is incompatible with the risk of believing what is false. Thus, risk and knowledge are plausibly inversely connected. But minimizing the risk of false belief, when risk is understood as a function of probability, is not sufficient to make a belief safe in the way required by knowledge (Pritchard 2005; Williamson 2009) as the famous lottery puzzle demonstrates (Nelkin 2000). Even though it is extremely probable that your lottery ticket is a loser, you cannot know that it is a loser simply on the basis of the odds involved (Hawthorne 2004; Pritchard 2005). Pritchard (2015) has also presented cases across which the probability of some proposition p is the same but the risk of p intuitively differs. Consider:

BOMB 1: An evil scientist has rigged a bomb to explode if ticket #65 wins the lottery. The winning ticket is determined by a lottery draw, which involves six balls falling out of a lottery machine. The odds that ticket #65 wins the lottery is 1: n.

[5] For empirical evidence of intuitions regarding the bus cases, see Wells (1992).

BOMB 2: An evil scientist has rigged a bomb to explode if, (*a*) the weakest team remaining in the FA cup beats the best team in the cup by ten goals, and (*b*) the weakest horse in the field at the Grand National wins the race by at least 2000 meters, and (*c*) the Queen of England spontaneously chooses to speak a complete sentence of Polish during her next public speech. The odds of *a*, *b*, and *c* occurring jointly is 1: *n*. (Pritchard 2015: 441)[6]

Pritchard intuits that in Bomb 1 the risk that the bomb explodes is greater than in Bomb 2. But since the probability that the bomb explodes is 1: *n* in both cases, there should not be a difference in its riskiness. But since equiprobable propositions can differ in their riskiness the probabilistic theory is extensionally inadequate.

Another reason to doubt the probabilistic theory of risk is that the concept of risk precedes the concept of probability. Risk plays an important role in our judgements of knowledge, responsibility, and safety, all of which are far older than the 17th century concept of probability that underlies the probabilistic theory of risk. It is not then surprising that probability and risk, safety, and knowledge come apart (Williamson 2009: 19).

It may be objected that the problems that the probabilistic theory faces are illusory since our intuitions and practices that relate to risk are products of heuristics, such as the availability bias (Kahneman and Tversky 1973, 1974).

There are two problems with this interpretation. *First*, heuristic reasoning patterns tend to occur in absence of knowledge of the facts relevant to their non-heuristic analogues. For example, in the case of availability bias, an agent who undergoes a heuristic judgement based on some available evidence is likely to change their judgement when availed to a more expansive evidence set. In the examples given, in contrast, the probabilities are known to be the same in both cases, and yet the intuition that there is a difference in risk prevails.

Secondly, the interpretation merely assumes that the probabilistic theory of risk is correct, and sees our tendency to deviate from its recommendations as a failure of rationality (Gigerenzer 1996; Pritchard 2015). But if we are interested in what risk is, we cannot assume at the start of the investigation that the probabilistic theory of risk is correct. Rather, the fact that our judgements and practices of regulating risk deviate from the predictions of the probabilistic theory of risk should be taken as evidence that the notion of risk employed by us in general is distinct from the probabilistic one.

While these reasons don't suffice to show conclusively that the intuitions about the bus and bomb cases are not the product of mistaken heuristics, they do give us a reason to search for an alternative explanation of the cases.

[6] In Pritchard's original case, the probability that the bomb explodes is set to 1:14 000 000. Ebert et al. (2020) have argued that the probability of *a* & *b* & *c* is much lower. For this reason, we have not fixed the value of *n*. The reader is free to estimate it on their own.

III. The modal and normic theories of risk

The probabilistic theory is extensionally inadequate. Recently two modal theories have been developed in response.

Pritchard (2015, 2016) endorses *the modal theory*, according to which risk is solely a function of modal closeness. That is:

Modal theory: For all propositions *p*, *q*, and worlds *w*, *p* is riskier than *q* if, and only if the closest world where *p* is true is closer to the actual world than any world where *q* is true. (Pritchard 2015: 447)

As usual, closeness between possible worlds is understood in terms of similarity. To determine the risk of a proposition *p,* we look to the closest world from the actual world at which *p* holds. The closer the world is to the actual world, the greater the risk of *p*. A proposition *p* is more risky than a proposition *q* just in case the closest world at which *p* is true is closer to the actual world than the closest world at which *q* is true.

Ebert et al. (2020) endorse *the normic* theory, according to which the risk of a proposition is solely a function of the normalcy of the world at which the proposition holds. Normalcy is understood in terms of the *need to provide explanations*. It would be normal for a car engine to start when the keys are turned in the ignition, in that it would not call for an explanation. But if the car did not start, it would be abnormal, since it would call for an explanation (Ebert et al. 2020: 443). The normic theory relativizes the risk of a proposition to a body of evidence. This gives us the following condition on comparative risk:

Normic theory: For all propositions, *p*, *q*, worlds *w*, and bodies of evidence E, *p* is riskier than *q* given E if, and only if: the most normal world where E and *p* are true is more normal than any world where E and *q* are true. (Ebert et al. 2020: 444)[7]

The modal and normic theories are not extensionally equivalent with the probabilistic theory. A proposition that holds at a normal, or at a modally close world, may be improbable. The probability that a given isotope of uranium-238 decays during the next year is low (roughly 1:4.5 billion), but the closest world at which the isotope decays during this interval is extremely close to the actual world, since decaying is a spontaneous process. Similarly, it would be perfectly normal for the isotope to decay during the interval, even though it is highly improbable.

The fact that the modal and normic theories are extensionally inequivalent with the probabilistic theory turns out to be a virtue of the former theories, at

[7] Ebert et al. (2020) end up advocating risk pluralism, on which there are multiple equally legitimate notions of risk. We return to risk pluralism later.

least with respect to Pritchard's bomb cases. Unlike the probabilistic theory, the modal and the normic theories are able to accommodate the intuition that the risk that the bomb goes off in BOMB 1 is greater than it's in BOMB 2.

Consider the modal theory. The closest world at which the bomb goes off in BOMB 1 is comparatively close. In order for ticket #65 to win the lottery in BOMB 1 only a few coloured balls need to be dropped from the lottery machine in a different configuration. In contrast, more has to change about the world for the Queen to speak Polish than for the lottery balls to fall in a different configuration. So the closest world in which the bomb goes off in BOMB 2 is further away from the actual world, than the closest world in which the bomb goes off in BOMB 1 (Pritchard 2015: 442). So the risk is greater in BOMB 1 than in BOMB 2.

Consider the normic theory. The most normal world at which the bomb goes off in BOMB 1 is comparatively normal. If ticket #65 were to win the lottery in BOMB 1, no explanation would be required. In contrast, an explanation would be required if the bomb were to go off in BOMB 2. That the Queen spoke a sentence of Polish calls out for an explanation. So the most normal world in which the bomb goes off in BOMB 2 is less normal than the most normal world in which the bomb goes off in BOMB 1. So the risk is greater in BOMB 1 than in BOMB 2 (Ebert et al. 2020: 444).

The modal and the normic theories of risk allow us to hold onto plausible connections between risk minimization and epistemic concepts such as knowledge and justification (Pritchard 2016; Smith 2016). The modal theory of risk vindicates a safety condition on knowledge if knowledge requires minimizing the risk of false belief. The normic theory of risk is able to explain why we cannot believe with justification that a lottery ticket is a loser simply on the basis of the odds involved if epistemic justification requires minimizing the risk of false belief.[8] Furthermore, both theories have been argued to deliver the correct verdicts in the bus cases (Pritchard 2018: 118–9; Smith 2018: 1208). Which theory, if either, should we accept?

Ebert et al. (2020: 441) argue that the modal theory is extensionally inadequate because it assigns all propositions that are true at the actual world the maximal risk value. This is because any reasonable similarity measure will take the actual world to be maximally similar to itself. For any actually true proposition, the closest world at which that proposition holds is the actual world. Because no world is closer to the actual world than the actual world, on the modal theory no proposition is more risky than any actually true proposition. Hence, the modal theory implies that there is no distinction between a proposition being true and it being maximally risky.

[8] For discussion of how normic and modal accounts of knowledge contrast with each other, see Mortini (2024).

Ironically, a somewhat similar objection can be raised against the normic theory. This is because part of the reason the objection holds against the modal theory is what we call the *privileged world problem*, according to which a proposition's truth value at a single world can make it maximally risky. For the modal theory, this is the actual world. On the normic theory any maximally normal world will do the trick. After all, if *p* is true at a maximally normal world then there won't be a more normal world in which some other proposition is true. Hence, no proposition can be more risky than *p* if *p* is true at a maximally normal world. And since *p* can be a proposition that is true only at a single world, a proposition's truth value at a single world can make it maximally risky.[9]

The privileged world problem targets both the modal and normic theory. While only the modal theory collapses the distinction between truth and maximal risk, both the modal and normic theories are rendered extensionally inadequate in three ways due to the fact that they suffer from the privileged world problem.

First, they will violate the idea that all and only necessary truths are assigned the maximal risk value.[10] Since contingently true propositions can be true at either a maximally normal world, or the actual world, both theories entail that contingently true propositions can be maximally risky. This means that a contingent proposition can be as risky as a necessary truth. Moreover, these contingent propositions can be intuitively extremely low risk propositions. Take the conjunction of all propositions that obtain at the privileged world. The modal and normic theories entail that the risk of the conjunction is maximal, but it seems that the *negation* of the conjunction has a much higher risk than the conjunction.

Secondly, they violate the intuition that true or normal propositions can be less risky than propositions that are merely possibly true. Suppose a dice is thrown and lands on a '6', and that the actual world belongs to the set of the most normal worlds. Even though the result was a '6', and '6' was as normal as any other result, the risk that it would have been '$1 \vee 2 \vee 3 \vee 4 \vee 5$' is greater than the risk that it would be '6'. The modal and normic theories cannot accommodate this fact.

[9]Note this argument does not require that the normic ordering over worlds is strongly centred, i.e., that there exists a unique maximally normal world. This is because the risk of a proposition, on the normic view, is determined by the normalcy of the most normal world at which it holds. If there is only a single maximally normal world at which the proposition holds, irrespective of whether that world is uniquely maximally normal, the proposition will be counted as maximally risky. We would like to thank an anonymous reviewer for raising this concern.

[10]It may be wondered whether a theory of risk ought to assign risk values to necessary truths at all, especially when the necessity of the proposition is known. After all, it seems strange to assert that the risk that Hesperus is Phosphorus is maximal, given that we know Hesperus to be Phosphorus. Although we are amenable to this suggestion, we can consider necessary truths to be a kind of limit case for formal models of risk. We would be disposed to seek an explanation for the seeming infelicity of risk assertions concerning known necessary truths by appealing to norms of assertion. Thanks to an anonymous reviewer for raising this point.

Thirdly, they violate the intuition that distinct but actually true or normal propositions can differ in their risk. Consider throwing two dice: one six sided, one twelve sided. For both throws if the die lands on 1 you will lose a substantial financial sum. Now suppose that, unbeknownst to you, both will actually land on 1, and that there is nothing abnormal about this. The modal and normic theories are forced to conclude that throwing the first die is equally risky to throwing the second. But intuitively, throwing the six sided die is more risky than throwing the twelve sided die. Essentially the problem is this: the modal and normic theories entail that we cannot make risk comparisons between propositions that hold at the actual world, or at some of the most normal worlds. But intuitively such propositions can differ in risk value.

The normic theory has the unintuitive consequence that two mutually exclusive propositions can be maximally risky. Suppose that you flip a fair coin that is guaranteed to land either 'heads' or 'tails'. It would be perfectly normal for the coin to land 'heads' and it would be perfectly normal for the coin to land 'tails'. The set of the most normal worlds contains both 'heads' and 'tails' worlds. Therefore, both propositions are maximally risky. However, there are far riskier propositions than flipping 'heads' with a fair coin. For example, it's far riskier to throw 'heads $\vee$ tails' than 'heads'. But the normic theory cannot accommodate this, since it would be equally normal to flip 'heads' or 'heads $\vee$ tails' with the fair coin.[11]

The modal theory avoids this consequence. If p and q are mutually exclusive and jointly exhaustive, then one of them is true in the actual world, making that proposition maximally risky. This entails that either p, or q, is maximally risky, depending on which of them obtains actually. However, this consequence seems to be just as bad. If you win a fair lottery, it was maximally risky that you would win it! Furthermore, the risk of winning the lottery was the same as the risk of either not winning the lottery or winning it, since the disjunction is true at the actual world. But clearly the disjunction is more risky than either disjunct!

The modal and normic theories have undesirable formal properties. We note two. *First*, they are unable to tell how much riskier a proposition is compared to another one.

Ebert et al. (2020: 441, 444) have argued the modal and normic theories can be enhanced to give an ordinal ranking of risks. To do this, one would need to endorse the limit assumption, according to which, for each world w, and proposition p, there is a unique smallest p-permitting sphere. A sphere is said to be p-permitting if it contains a world at which p is true. The system of

[11] The normic theory can avoid this consequence if the system of normalcy spheres is strongly, rather than weakly centred (Lewis 1973: sec. 1.7). However, if the system of normalcy spheres is strongly centred, thus guaranteeing that there is a unique most normal world, then the normic theory has the same problem that the modal theory has with mutually exclusive and jointly exhaustive propositions.

spheres is required to be at least non-empty, nested, and closed under unions and non-empty intersections (Smith 2016: 154). On this picture, the size of the smallest p-permitting sphere, centred on w, gives the similarity or normalcy 'rank' of p. The smallest possible similarity/normalcy rank that proposition can have is 0. The smaller the similarity/normalcy rank of a proposition the riskier it is.

But ranking worlds in the above way only yields an ordinal, not a cardinal ranking of risks (Smith 2016: 168). A fully fledged theory of comparative risk ought to tell us both whether a proposition is more risky than another, and how much riskier it is. But the modal and normic theories are only able to say whether a proposition is riskier than another.

Secondly, the modal and normic theories validate the following inference pattern, called *checklist reasoning*:

1. The risk of p is low.
2. The risk of q is low.
3. The risk of $p \vee q$ is low.

Since a disjunction is true at any world in which either disjunct is true, the risk of a disjunction on the modal and normic theories is always equal to the risk of the most risky disjunct. If p is riskier than q, then the risk of $p \vee q$ is equal to the risk of p.

If checklist reasoning is invalid, it would not always be safe to ignore small risks, since small risks could potentially add up to big ones. But if checklist reasoning is a valid pattern of inference then it's valid for all risk values. If small risks never add up, then big risks never add up either. Therefore, when considering whether checklist reasoning is a valid pattern of inference, we ought to have in mind the following inference pattern, which the modal and normic theories also validate:

1. The risk of p is n.
2. The risk of q is n.
3. The risk of $p \vee q$ is n.

This generalized pattern is clearly invalid. Suppose Jani takes the following bet. If you throw a '6' with a fair six-sided dice, Jani will buy you a Big Mac. Leinonen takes another bet. If you throw either a '6' or a '5', he will buy you a Big Mac. Intuitively, Leinonen is at greater risk of losing the bet. But because the world where you throw a '6' is just as normal, or close to the actual world, as the world where you throw a '6' or a '5', the modal and normic theories count both bets as equally risky. These theories entail that a disjunction can never be riskier than either disjunct. But that is plainly false. Risks can add up.

Ebert et al. (2020) argue that validating checklist reasoning is a virtue of the modal and normic theories, because it is implied by *de minimis* risk management, a practice in which risks that are deemed extremely low are ignored in future risk calculations (Peterson 2002). When developing the unified theory, we show that it's possible to make good sense of *de minimis* risk management while rejecting the validity of checklist reasoning.

IV. A unified theory of risk

The probabilistic theory is extensionally inadequate, and the modal and normic theories are extensionally inadequate and have undesirable formal properties. It may be thought that by embracing a pluralist theory, the best aspects of both theories can be kept. Ebert et al. (2020) endorse risk pluralism on which the probabilistic and the normic theories codify two equally legitimate notions of risk. Our view is that pluralism is best avoided. We lack the space to fully argue against risk pluralism, but we note two central problems for such theories. *First*, a pluralist approach faces a normative problem: Which notion of risk should we use when evaluating the risk of a given event? *Secondly*, the pluralist faces a formal problem: Given that the different theories entail different structural properties for risk, how can we calculate risks of complex events involving inferences regarding normic and probabilistic risks? Without an answer to these problems the prospects of the pluralist theory are bleak.[12]

A preferable strategy is to develop a monist theory of risk that unifies our judgements of risk, whilst maintaining the desirable formal properties of the probabilistic view. In what remains we develop such a view: the *unified theory of risk*.

The unified theory starts with the observation that probabilistic, modal, and normic theories tap into something essentially correct about risk. The probabilistic view is correct insofar as it holds that, *ceteris paribus,* increases in the proportion of possible worlds at which a proposition holds increases its risk.[13] The modal and normic theories are correct insofar as they hold that, *ceteris paribus,* increases in the similarity/normalcy of the most similar/normal world at which a proposition holds increases its risk. The problem with all of the views is that they consider risk to be a *univariable* function, of either the proportion, modal closeness or normalcy of the worlds at which a proposition holds. The unified theory obtains the best aspects of both theories by considering risk to be a *multivariate* function. The risk of a proposition p is a

[12]See Mace & O'Sullivan (2024) for pluralist account that aims to solve the first of these problems.

[13]If we subscribe to the principle of indifference, we can straightforwardly equate the probability of a proposition with the proportion of possible worlds in which that proposition holds. However, it should be noted that the principle of indifference is viewed with suspicion, primarily because it leads into Bertrand's paradox (Bertrand 1889). While the principle is suspicious many are committed to its truth, since it's entailed by the principal principle (Hawthorne et al., 2015). We don't take a stand on this principle.

function of both the proportion and modal closeness of the worlds at which *p* is true.[14]

Here is the plan. *First*, we give a model in which the consequences of the theory can be examined with formal precision. *Secondly*, we use the model to demonstrate that the unified theory is extensionally adequate and has desirable formal properties.

We start with a *frame*, which is an ordered pair $<W, A>$, where W is a set and A is a binary accessibility relation over members of W.[15] Informally W is conceived as the set of metaphysically possible worlds. Possible worlds are individuated in terms of the propositions that are true at those worlds. Like the probabilistic theory, we make the idealization that the space of possibilities is finite.

To this frame we add a similarity measure $\$$, which is a function from ordered pairs of members of W to the half open unit interval (0,1].[16] Informally, $\$$ tells us the relative similarity between two worlds. The greater the value of $\$(w, w^i)$, the greater the similarity between w and w^i is from the perspective of w.

We add a set of functions R to the frame, which maps values of $\$$ for each world from the perspective of an evaluation world to the half open unit interval (0, 1]. Informally R is conceived as assigning *risk values* to each world from the perspective of an evaluation world. The risk values of members of W (from the perspective of an evaluation world w) are generated from the similarity values as follows:

Let $\$_w$ be the multiset of all $\$(w, x)$.[17] Then let $R_w = x/\Sigma_{y \in \$w}\, y$, where $x = \$(w, x)$ for some x such that $\$(w, x)$. Basically, R_w is a normalization of the values $\$_w$.

Here is a concrete example of how to calculate risk values for worlds in a frame. Consider the following frame where $W = \{w, w^1, w^2, w^3, w^4\}$ and $\$_w$ has the following values: $\$(w, w) = 1$, $\$(w, w^1) = 0.7$, $\$(w, w^2) = 0.7$,

[14]Compare de Grefte's (2020) theory of luck. One of the primary differences between the unified theory of risk and de Grefte's account of luck is that for de Grefte the degree to which *p* is a matter of luck depends on the closeness of the *closest* world where *p* is true and the probability of *p* (2020: 246). The unified theory of risk, in contrast, holds that risk of *p* is partially a function of the modal closeness of all the possible worlds where *p* is true. Billot et al. (2005) and Gilboa et al. (2010) hold that similarity weighted frequencies to *constitute* probabilities. We think that probability is a useful concept in itself, and would hesitate to equate it with risk. See Kratzer (2012: 41–3) for a way of generating a probability measure from comparative possibilities.

[15]A can be understood as having access to all worlds which are compatible with one's knowledge at the evaluation world. We don't commit to any specific interpretation of A. For simplicity, all members of W are taken to be accessible from the evaluation world.

[16]The half open unit interval (0,1] includes all the numbers that are > 0 and ≤ 1. The reason why similarity values fall within (0,1] is that while two worlds can be maximally similar to each other when they are identical, two worlds cannot be completely dissimilar to each other. Propositions that are necessarily true are true in all possible worlds, and hence all possible worlds have something in common.

[17]Multisets differ from sets in that they can contain the same member multiple times.

$\$(w, w^3) = 0.5$, $\$(w, w^4) = 0.1$. To get the risk value distribution of members of W, first sum up all the members of $\$_w$:

$$1 + 0.7 + 0.7 + 0.5 + 0.1 = 3$$

Then, for each member of $\$_w$, divide it with the sum previously obtained:

(1 / 3), (0.7 / 3), (0.7 / 3), (0.5 / 3), (0.1 / 3)

These divisions give the risk value distribution of members of W. In this particular frame, the risk values are distributed as follows: $R_w(w, w^1) = 1/3$, $R_w(w, w^2) = 0.7/3$, $R_w(w, w^3) = 0.7/3$, $R_w(w, w^4) = 0.5/3$, and $R_w(w, w^5) = 0.1/3$.

The risk value distribution that R_w generates, respects the following conditions, proof of which is in Appendix 1:

(i) NORMALIZATION: if $\forall w \in W$, $w \vDash p$ then $\Sigma_{wi \vDash p \in W} R_w(w, w^i) = 1$
(ii) ADDITIVITY: if $\neg\exists w \in W$, $w \vDash (p \;\&\; q)$ then $\Sigma_{wi \vDash p \vee q \in \mathrm{W}} R_w(w, w^i) = \Sigma_{wi \vDash p \in \mathrm{W}} R_w(w, w^i) + \Sigma_{wi \vDash q \in W} R_w(w, w^i)$
(iii) CLOSENESS: $\forall w \in W$, if $\$(w, w^i) > \$(w, w^j)$, then $R_w(w, w^i) > R_w(w, w^j)$

Given that the outputs of R_w fall within the half open unit interval (0, 1] the risk value of any given world is > 0 and ≤ 1. NORMALIZATION states that if a proposition is true at all accessible worlds, then the sum of the risk values of those worlds is 1. ADDITIVITY states that for any two propositions p, q, that are mutually exclusive, the sum of the risk values of the worlds in which either p is true or q is true is the sum of the risk values of worlds in which p is true plus the sum of the risk values worlds in which q is true. CLOSENESS states that when assigning risk values to worlds, worlds that are more similar to the evaluation world w receive a greater risk value than worlds that are less similar to w.

In this framework, worlds are given *risk values* and propositions are given *risk scores*. The risk score of p at the evaluation world w is obtained by summing the risk values of all the worlds in which p is derivable. The risk of p is straightforwardly equated with the risk score of p. That is:

$$\mathrm{R}(p) = \sum_{wi \vDash p \in W} R_w(w, w^i)$$

The comparative risk relation is understood in terms of inequality between risk scores:

$$\mathrm{R}(p) > \mathrm{R}(q) \text{iff} \sum_{wi \vDash p \in W} R_w(w, w^i) > \sum_{wi \vDash q \in W} R_w(w, w^i)$$

A proposition p is more risky than a proposition q just in case the risk score of p is greater than the risk score of q. Two propositions are equally risky just in case they have the same risk score. If p is riskier than q, then how much riskier p is given by subtracting the risk score of q from the risk score of p.

The risk score 1 is the highest possible risk score that a proposition can receive, and 0 is the lowest score. The risk score 1 indicates maximal risk that the proposition is true and 0 that there is no risk that the proposition is true.[18,19]

The unified theory doesn't share any of the troublesome formal properties that the modal and normic theories entail. Consider the following model, in which W consists of five members: @, w^1, w^2, w^3, and w^4. Assume that all of these worlds are accessible from @ and that the risk values of members of W and truth values of propositions p, q, r, and s are distributed as follows:

World	Risk value	Value of p	Value of q	Value of r	Value of s
@	0.4	1	1	0	0
w^1	0.2	1	1	1	0
w^2	0.2	0	1	0	0
w^3	0.1	0	1	0	1
w^4	0.1	0	1	0	0

where 1 represents that the proposition is true and 0 that it is false.

It is easy to prove that the unified theory doesn't fall prey to the privileged world problem, since there is no world such that any proposition that holds at it is maximally risky.

Proof: Let w be the supposed privileged world. By the law of excluded middle w is either a p or a $\neg p$ world. In the above model, the risk score of p is 0.6 and the risk score of $\neg p$ is 0.4. Since either p or $\neg p$ is true at w, the fact that a proposition is true at w does not make it maximally risky.

On the unified theory, checklist reasoning is an invalid pattern of reasoning.

Proof: Suppose that 0.2 is a low risk score, and that any risk score higher than 0.2 isn't a low risk score. Therefore, the risk of r (0.2) is low, and the risk of s (0.1) is low. However, the risk score of $r \vee s$ is 0.3, since the disjunction is true

[18] One of us argues that we can gain a greater insight into the structure of luck via the unified theory of risk since luck and risk are negatively correlated (Hirvelä 2024).

[19] It's possible to modify the unified theory to accommodate the idea that more detrimental events are *ceteris paribus* more risky than less detrimental events. To do so we need to introduce a function that assigns a disutility value for each world on the half-open unit interval (0, 1]. The risk value of each world would then be multiplied with its disutility. The disutility weighted risk value of a world is the product of its disutility and risk value divided by the sum of the products of all disutility and risk values of each world. The disutility weighted risk score of a proposition is the sum of the disutility weighted risk values of the worlds in which the proposition is derivable.

at w^1 and at w^3. Since 0.3 isn't a low risk score, checklist reasoning is invalid on the unified theory.

One might object that since the unified theory invalidates checklist reasoning it cannot make sense of *de minimis* risk management. As we have argued, invalidating checklist reasoning is a virtue of a theory of risk. Risks can add up. While the unified theory invalidates checklist reasoning, we can still make sense of *de minimis* risk management. We hold that ignoring low risk propositions is done for pragmatic reasons. For example, if a law firm wishes to assess the risk of taking on a client, they may consider only a subset of the known risks, since considering all of the known risks would be laborious and cognitively demanding. Although the risk assessment would be worse, from an epistemic standpoint, it would nevertheless be more pragmatic given the firm's operational resources. In our model, we might artificially assign low risk propositions a risk score of 0 for the purposes of future calculations. If both r and s are assigned the risk score 0 in all future calculations, then the risk score of $r \vee s$ is 0 according to the unified theory. Therefore, the unified theory can make perfect sense of *de minimis* risk management without granting that checklist reasoning is a valid pattern of inference.

Recall that on the normic theory, two mutually exclusive events can be maximally risky. The unified theory doesn't share this problematic feature.

Proof: Let p and q be mutually exclusive propositions and the risk score of $p =$ 1. If the risk score of $p = 1$, then necessarily, p is true at every world. Since p and q are mutually exclusive propositions, there is no p & q world. Since every world in W is a p-world, no world in W is a q-world. Therefore, the risk of q is 0, and hence q is not maximally risky. Two propositions that are mutually exclusive are never both maximally risky.

Recall that on the modal theory if two propositions are mutually exclusive and jointly exhaustive, then one of them is maximally risky. The unified theory doesn't share this problematic feature.

Proof: Let p and q be mutually exclusive and jointly exhaustive propositions, and the risk score of p be 0.6. Given that p and q are mutually exclusive and jointly exhaustive propositions, q is true in all, and only the worlds where p is false. The risk score of q is simply the sum of the risk values of the possible worlds where q is true. Since the sum of the risk values of the worlds at which p is false is 0.4, the risk score of q is 0.4.

Recall that the modal and the normic theory yielded only an ordinal ranking of risks. While these accounts enable us to say that a p is riskier than q, they cannot tell us how much riskier p is. The fact that the modal and the normic theories fail to give a cardinal ranking of risks makes these accounts inapt for many applications where the notion of risk plays a significant role, such as decision theory. For example, there is no straightforward way to determine whether one should ϕ or ¬ϕ, if ϕ-ing entails that the risk of *p* is 56 and ¬ϕ-ing entails that the risk of *q* is 40, while *p* three times more desirable than *q*.

The unified theory, in contrast, yields a cardinal ranking of risks. It tells us how much riskier a proposition is compared to another one. This feature allows us to implement the unified theory into decision-theoretic frameworks. There are two dominant decision-theoretic frameworks: the causal and the evidential one. The causal expected utility of an act A is the weighted average of the utility of each possible causal outcome of the act, multiplied by the probability of the act-state dependency hypothesis, where the act causes the outcome. The evidential expected utility of an act A is the weighted average of the utility of a possible state of the world and A, multiplied by the probability of the state conditional on A. Since the risk scores that the unified theory generates fall within the unit interval, and obey ADDITIVITY and NORMALIZATION, probabilities can straightforwardly be replaced with risks from a formal perspective when doing decision theory.

Unlike the modal and normic theories, the unified theory is able to count lotteries that differ only in their probability as differing in their risk. Consider again two lotteries with different odds, where a single outcome will cause a bomb to explode.[20] On the unified theory the risk of an event increases if the proportion of worlds where it obtains increases. Since the ratio of possible worlds where the bomb explodes is greater in one lottery as opposed to another, the outcomes are assigned different risks on the unified theory. It's worth noting that the unified theory is the only theory that can deliver intuitively correct verdicts across all of these cases.

Finally, let us return to consider the cases that the probabilistic theory struggled with. Assuming that neither bomb actually explodes, the unified theory is able to deliver the intuitively correct verdict that the risk that the bomb explodes is greater in BOMB 1 than in BOMB 2. Even though the explosions are equiprobable, the possible world where the bomb explodes in BOMB 1 is much closer to the actual world than the possible world where the bomb explodes in BOMB 2. Since worlds that are closer to the actual world carry greater weight than worlds that are further away, the unified theory predicts that the risk that the bomb explodes is greater in BOMB 1 than in BOMB 2.

But if we stipulate that *a*, *b,* and *c* obtain in the actual world, the unified theory will predict that the risk that the bomb explodes is greater in BOMB 2 than in BOMB 1. We don't think that this is a problem. The intuition that the risk of the bomb exploding is greater in BOMB 1 is guided by what we take the actual world to be like. If we add details to BOMB 2, such that *a*, *b*, and *c*

[20]See Ebert et al. (2020: 446) for discussion of such a case. Their response to the case is to adopt pluralism about risk.

are not far away possibilities, but modally close events, the relevant intuition is reversed.

Now consider the bus cases. Of the theories considered so far, only the probabilistic theory and the normic theory straightforwardly entail a precise verdict on both cases. The probabilistic theory entails that the risk that a blue bus hit the pedestrian is equally high in both BUS 1 and in BUS 2, since the probability of it is the same across the cases. The normic theory entails that the risk that the pedestrian was not hit by a blue bus in BUS 1 is high, since given the evidence available, it would be normal that the pedestrian was hit by another bus. But the risk that the pedestrian was not hit by a blue bus in BUS 2 is low, since there would have to be some special explanation as to why the eyewitness testimony was false. However, the information that we have regarding BUS 1 is not sufficient to determine whether the risk that a blue bus hit the pedestrian is high or low according to the modal and the unified theory of risk. This is because knowing the mere probability of a proposition does not allow us to determine the modal spread of the worlds in which the proposition is true, and hence the modal and the unified theory of risk do not deliver a precise verdict of the case. For all we know the risk that a blue bus hit the pedestrian might be low or high, and hence it is plausibly not permissible to find the blue bus company liable in BUS 1. The situation in BUS 2 is, however, radically different since we have reliable testimony that directly links the blue bus company bus to the accident (Pritchard 2018: 118). The evidence is not merely probabilistic, but also makes it reasonable to think that the worlds in which a blue bus hit the pedestrian are modally closer than the possible worlds where a bus operated by another company hit the pedestrian. We say 'makes it reasonable' since the evidence doesn't entail it. After all, the testifier could be a pathological liar. But this is hardly a problem for the unified or modal theory of risk. After all, any theory that aims to capture an objective fact might fail to deliver the correct verdict in a given case due to impoverished evidence. In such cases, we have to make reasonable estimates about how to fix the values of certain parameters. The modal and the unified theory of risk is no exception to this general phenomenon.[21]

Before concluding, we consider two objections. The first objection states that the unified theory is a probabilistic theory of risk because it satisfies Kolmogorov's axioms of probability. But in addition to Kolmogorov's axioms, the unified theory of risk vindicates CLOSENESS. There are some risk distributions that are possible on the probabilistic theory of risk that the unified theory of risk deems as impossible. One way to think about the difference between the unified and the probabilistic theory is that the unified theory gives a principled answer to the question of how prior probabilities ought to be assigned.

[21] Fratantonio (2021) argues that the prospects of epistemic solutions to the proof paradox are bleak.

Worlds that are modally closer ought to be assigned higher prior probabilities than worlds that are further away. And since modal distance cannot be understood in purely probabilistic terms, the unified theory is still at heart a modal theory of risk. Risk cannot be understood purely in terms of probability. It is a modal concept.

Secondly, one might wonder why proponents of the normic theory could not endorse the kind of model that we put forward, and substitute the similarity ranking of the worlds with a normalcy ranking. Wouldn't the resulting theory be better off, in that it would say that there is a greater risk that the bomb explodes in BOMB 1 than in BOMB 2, no matter whether *a*, *b*, and *c* obtain in the actual world?

We are happy if proponents of the normic theory endorse our model and substitute the similarity ranking of the worlds with a normalcy ranking. But we think this view would ultimately be mistaken. This is because, as we now argue, there is a reason to think that risk is not a function of normalcy, at least when normalcy is understood in terms of explanation as Ebert et al. (2020) do. Locating abnormal events that call for explanation is a paradigmatic way of finding new research questions in science. Often enough it turns out that our best scientific theories predict that these events are bound to occur in such and such settings. Hence, they are not low risk events, contrary to what the normic theory predicts.

Consider the following case. Suppose that you've never tried to break an uncooked spaghetti noodle by bending it from both ends, but have tried to break many other rod-like objects, such as pens, in this way. These other objects have always broken in exactly two pieces. Prior to bending the spaghetti noodle, you think that it would be highly abnormal if it broke into three or more pieces, since you have good inductive evidence for thinking that thin rod-like objects always break into exactly two pieces. Weirdly enough, if you try to break a spaghetti noodle in this way, it always breaks into three or more pieces, a phenomenon that reputedly perplexed Richard Feynman.[22] While the fact that the spaghetti breaks into three or more pieces is something that calls for an explanation given your evidence, it's nevertheless highly risky, and something that our best physics predict. We think that if you were to judge prior to breaking the noodle that it's highly risky that the noodle will break into three or more pieces your judgement would be correct (although it would be irrational, since your evidence indicates that it will break into exactly two pieces).

An anonymous reviewer offers the following reply on behalf of Ebert et al (2020). On the normic theory, risk ascriptions are relativized to bodies of

[22]If you're curious as to why uncooked spaghetti noodles always break into three or more pieces here is the scientific explanation: after the noodle has broken in two, the remaining pieces bend backwards, creating flexural waves that break the remaining pieces further (Audoly and Neukirch 2005).

evidence. Since it is highly abnormal given your evidence that the noodle will break into three or more pieces, you would be correct in judging that the risk that the noodle breaks into three or more pieces is low. A scientist who is availed to the relevant evidence concerning the physics of thin rod-like objects would be correct in judging that the risk that the noodle breaks into three or more pieces is very high. According to the objection the fact that distinct bodies of evidence give rise to contradictory risk assessments that are correct is a feature rather than a bug of the normic theory.

The reviewer is certainly correct in that the normic theory has this consequence. However, we don't think that it gets Ebert et al. (2020) out of the pickle. To see this, consider the following continuation of the case. Suppose that after breaking the first noodle you are surprised by the fact that the noodle broke into three or more pieces. Perplexed by this you start breaking more and more noodles. At some point you will have acquired enough inductive evidence for knowing that the next noodle you will break will break into three or more pieces. Even so, we contend, the fact that the next noodle will break into three or more pieces calls for an explanation. The event calls for an explanation even though you know it will obtain. Since it calls for an explanation the normic theory, as developed by Ebert et al. (2020), entails that the risk that the noodle breaks into three or more pieces is low. We find this implausible given that you know that it will break into three or more pieces.

Finally, it is worth stressing that the above argument challenges normic theories of risk only insofar as they understand risk in terms of explanatory requirements. It might be possible to understand normalcy in some other way, in which case normic accounts might be able to deal with this case.

V. Concluding remarks

We argued that extant theories of comparative risk are either extensionally or formally inadequate and offered a new, unified theory of risk. According to the unified theory, the riskiness of an event is a function of both (i) the proportion of worlds at which the event obtains and (ii) the modal closeness of the worlds at which the event obtains. The unified theory is preferable to existing alternatives because it survives the tests of extensional and formal adequacy and can be implemented into decision-theoretic frameworks.[23]

[23]We would like to thank Sven Bernecker, Adam Bricker, Daniel Drucker, Giada Fratantonio, Maria Hämeen-Anttila, Antti Kauppinen, Maria Lasonen-Aarnio, Jesús Navarro, Lilith Mace, Martin Smith, Timothy Williamson, two anonymous reviewers at *Philosophical Quarterly*, and the audiences at Seville and Cologne for helpful comments and discussion. This project was supported by the Kone Foundation under the grant 'The Structure and Nature of Risk, and the Research Council of Finland under the grant 'The Metaphysics, Ethics and Epistemology of Risk.

Appendix

In this appendix, we prove that R_w satisfies NORMALIZATION, ADDITIVITY, and CLOSENESS. Let us start with NORMALIZATION, which is the claim that if $\forall w \in W, w \vDash p \rightarrow \Sigma_{wi \vDash p \in W} R_w(w, w^i) = 1$.

Proof: Assume that $\forall w \in W, w \vDash p$. This entails that $\{w \mid w \vDash p\} = \{w \mid w \vDash p \vee q\} = W$.

$\Sigma_{wi \in W, wi \vDash p \vee q} \; R_w(w, w^i) = \Sigma_{wi \in W,} \; R_w(w, w^i)$. $\Sigma_{wi \in W,} \; R_w(w, w^i) = \Sigma_{y \in \$w} y / \Sigma_{y \in \$w} y = 1$.QED.

Let us prove that R_w validates ADDITIVITY which states that

if $\neg\exists w \in W, w \vDash (p \;\&\; q)$ then $\Sigma_{wi \vDash p \vee q \in W} R_w(w, w^i) = \Sigma_{wi \vDash p \in W} R_w(w, w^i) + \Sigma_{wi \vDash q \in W} R_w(w, w^i)$.

Assume that $\neg\exists w \in W, w \vDash (p \;\&\; q)$. Let $\$_{w, A}$ be the multiset $\$_w|_A \{\$(w, x) \mid x \in W \;\&\; x \vDash A\}$. Since $\neg\exists w \in W, w \vDash (p \;\&\; q) \rightarrow \{w \mid w \in W \vDash p\} \cap \{w \mid w \in W \vDash q\} = \emptyset$. $\Sigma_{wi \in W, wi \vDash p} R_w(w, wi) = \Sigma_{y \in \$w,p} \, y \;/\; \Sigma_{y \in \$w} \, y$. $\Sigma_{wi \in W, wi \vDash q} R_w(w, wi) = \Sigma_{y \in \$w,q} \, y \;/\; \Sigma_{y \in \$w} \, y$. $(\Sigma_{y \in \$w,p} \, y \;/\; \Sigma_{y \in \$w} \, y) + (\Sigma_{y \in \$w,q} \, y \;/\; \Sigma_{y \in \$w} \, y) = \Sigma_{y \in \$w,p} \, y + \Sigma_{y \in \$w,q} \, y \;/\; \Sigma_{y \in \$w} \, y = \Sigma_{y \in \$w,p \vee q} \, y \;/\; \Sigma_{y \in \$w} \, y = \Sigma_{wi \in W, wi \vDash p \vee q} R_w(w, w^i)$. QED.

Finally, let us prove that R_W satisfies CLOSENESS, which states that: $\forall w \in W$, if $\$(w, w^i) > \$(w, w^j)$, then $R_w(w, w^i) > R_w(w, w^j)$. Assume $\$(w, w^i) > \$(w, w^j)$. $R_w(w, w^i) = \$(w, w^i) \;/\; \Sigma_{y \in \$w} \, y$. $R_w(w, w^j) = \$(w, w^j) \;/\; \Sigma_{y \in \$w} \, y$. Since $\$(w, w^i) \;/\; \Sigma_{y \in \$w} \, y > \$(w, w^j) \;/\; \Sigma_{y \in \$w} \, y$, it follows that $R_w(w, w^i) > R_w(w, w^j)$. QED.

References

Audoly, B. and Neukirch, S. (2005) 'Fragmentation of Rods by Cascading Cracks: Why Spaghetti Does Not Break in Half', *Physical Review Letters,* 95: 095505.

Bertrand, J. (1889) *Calcul Des Probabilités*. Paris: Gauthier-Villars.

Billot, A. et al. (2005) 'Probabilities as Similarity-Weighted Frequencies', *Econometrica*, 73: 1125–36.

Buchak, L. (2013) *Risk and Rationality*. Oxford: OUP.

de Grefte, J. (2020) 'Towards a Hybrid Account of Luck', *Pacific Philosophical Quarterly*, 101: 240–55.

Ebert, P. A., Smith, M., and Durbach, I. (2020) 'Varieties of Risk', *Philosophy and Phenomenological Research,* 101: 432–55.

Fratantonio, G. (2021) 'Evidence, Risk, and Proof Paradoxes: Pessimism about the Epistemic Project', *The International Journal of Evidence & Proof*, 25: 307–25.

Gardiner, G. (2020) 'Relevance and Risk: How the Relevant Alternatives Framework Models the Epistemology of Risk', *Synthese*, 199: 481–511.

Gigerenzer, G. (1996) 'On Narrow Norms and Vague Heuristics: a Reply to Kahneman and Tversky', *Psychological Review*, 103: 592–6.

Gilboa, I., Lieberman, O., and Schmeidler, D. (2010) 'On the Definition of Objective Probabilities by Empirical Similarity', *Synthese*, 172: 79–95.

Hansson, S. O. (2023) 'Risk', in E. N. Zalta and U. Nodelman (eds.) *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/sum2023/entries/risk/>. Accessed June 13, 2024.

Hawthorne, J. (2004) *Knowledge and Lotteries*. Oxford: OUP.

Hawthorne, J. et al. (2015) 'The Principal Principle Implies the Principle of Indifference', *The British Journal for the Philosophy of Science* 68: 123–31.

Hirvelä, J. (2019) 'Global Safety: How to Deal with Necessary Truths', *Synthese*, 196: 1167–86.

Hirvelä, J. (2022) 'Justification and the Knowledge-connection', *Philosophical Studies*, 179: 1973–95.
Hirvelä, J. (2024) 'The Metaphysics of Risk and Luck.' *Nous*.
Hirvelä, J. and Paterson, N. J. (2021) 'Need Knowing and Acting be SSS-Safe?', *Thought: A Journal of Philosophy*, 10: 127–34.
Kahneman, D. and Tversky, A. (1973) 'On the Psychology of Prediction', *Psychological Review,* 80: 237–51.
Kahneman, D. and Tversky, A. (1974) 'Simulation Heuristic', in P. Slovic and A. Tversky (eds.) *Judgment under Uncertainty: Heuristics and Biases,* pp. 201–8. Cambridge: CUP.
Kratzer, A. (2012) *Modals and Conditionals*. Oxford: OUP.
Lasonen-Aarnio, M. (2010) 'Unreasonable Knowledge', *Philosophical Perspectives,* 24: 1–21.
Lewis, D. (1973) *Counterfactuals*. Oxford: Blackwell.
Littlejohn, C. (2020) 'Truth, Knowledge, and the Standard of Proof in Criminal Law', *Synthese*, 197: 5253–86.
Mace, L. and O'Sullivan, A. (2024) 'Reverse-Engineering Risk', *Erkenntnis*.
Mortini, D. (2024) 'The Explanationist and the Modalist', *Episteme*, 21: 371–386.
Moss, S. (2022) 'Knowledge and Legal proof', in *Oxford Studies in Epistemology*, Vol. 7. Oxford: OUP.
Navarro, J. (2019) 'Luck and Risk', *Metaphilosophy,* 50: 63–75.
Nelkin, D. K. (2000) 'The Lottery Paradox, Knowledge, and Rationality', *Philosophical Review*, 109: 373–409.
Pardo, M. S. (2018) 'Safety vs. sensitivity: Possible Worlds and the Law of Evidence', *Legal Theory,* 24: 50–75.
Paterson, N. J. (2020) 'Non-Accidental Knowing'. *Southern Journal of Philosophy*, 58: 302–26.
Paterson, N. J. (2022) 'Safety and Necessity', *Erkenntnis*, 87: 1081–97.
Peterson, M. (2002) 'What Is a De Minimis Risk?', *Risk Management,* 4: 47–55.
Pritchard, D. (2005) *Epistemic Luck*. Oxford: OUP.
Pritchard, D. (2012) 'Anti-luck Virtue Epistemology', *Journal of Philosophy*, 109: 247–79.
Pritchard, D. (2015) 'Risk', *Metaphilosophy,* 46: 436–61.
Pritchard, D. (2016) 'Epistemic Risk', *Journal of Philosophy,* 113: 550–71.
Pritchard, D. (2018) 'Legal Risk, Legal Evidence and the Arithmetic of Criminal Justice', *Jurisprudence,* 9: 108–19.
Redmayne, M. (2008) 'Exploring the Proof Paradoxes', *Legal Theory,* 12: 281–309.
Smith, M. (2016) *Between Probability and Certainty - What Justifies Belief*. Oxford: OUP.
Smith, M. (2018) 'When Does Evidence Suffice for Conviction?' *Mind,* 127: 1193–218.
Smith, M. (2021) 'More on Normic Support and the Criminal Standard of Proof', *Mind*, 130: 943–60.
Sosa, E. (1999) 'How Must Knowledge Be Modally Related to What Is Known?' *Philosophical Topics,* 26: 373–84.
Wells, G. (1992) 'Naked Statistical Evidence of Liability: Is Subjective Probability Enough?', *Journal of Personality & Social Psychology,* 62: 739–52.
Williamson, T. (2000) *Knowledge and Its Limits*. Oxford: OUP.
Williamson, T. (2009) 'Probability and Danger', *The Amherst Lecture in Philosophy*, 4: 1–35.