

An Objection to the Modal Account of Risk

Martin Smith

In a recent paper in this journal Duncan Pritchard responds to an objection to the modal account of risk pressed by Ebert, Smith and Durbach (2020). In this paper, I expand upon the objection and argue that it still stands. I go on to consider a more general question raised by this exchange – whether risk is ‘objective’, or whether it is something that varies from one perspective to another.

Keywords: risk, modal account, objectivity

1 Risky Advice

Duncan Pritchard has recently defended a theory of risk that he calls the *modal account* (Pritchard, 2015, 2016). Consider the familiar picture of possible worlds arrayed in space, with their proximity to the actual world representing their similarity to it. Clustered close to the actual world are those worlds that are very similar, next come worlds that are almost, but not quite, as similar and so on until we reach remote worlds that are very different from actuality. On the modal account, the risk of a proposition is determined by the closeness of the closest worlds in which it obtains – the closer these worlds the higher the risk, and the more distant these worlds the lower the risk.

Since no world can be closer or more similar to the actual world than it is to itself, the modal account has the consequence that any proposition which is *true* will be at maximal risk of being true. In Ebert, Smith and Durbach (2020, section 4) it was argued that this would prevent the notion of risk from playing one of its central roles – namely, that of guiding decision making under conditions of uncertainty¹. This was illustrated with the following example:

Suppose one is about to drill into a wall in a West Australian house built in the 1970s, and is considering 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, one cannot take a view on the *risk* that the wall contains asbestos without already taking a view on whether it *does* contain asbestos.

(Ebert, Smith and Durbach, 2020, p441).

According to Pritchard, this objection trades on a confusion between the *actual* risk of a proposition and what would be a *reasonable assessment* of the risk, given the available evidence (Pritchard, 2022, pp10, 12). Pritchard concedes that if there really is asbestos in the wall then the modal risk that there is asbestos in the wall will be maximally high – but, in order to reasonably judge that there is a low modal risk that the wall contains asbestos, one doesn’t have to ‘settle’ the question of whether it does (Pritchard, 2022, p10). Suitably interpreted, I think this claim is correct – but it does nothing to neutralise the objection to the modal account of risk.

Let’s put risk aside for a moment. Suppose I am the person who is trying to decide whether to drill into the wall, with the possibility that it contains asbestos weighing on my mind. Suppose you

¹ Similar objections to the modal account have been presented by Gardiner (2020, section 5), Fratantonio (2021, section 4.1) and Newton (2022, section 3). See also Smith (2018, pp1204-1205).

offer the following advice: ‘As long as there’s no asbestos in any of the walls in the neighbourhood, you should go ahead and drill’. On the one hand this advice isn’t *bad* or *misleading* – in the circumstance in which there is no asbestos in any walls in the neighbourhood, it really would be fine to drill into the wall. On the other hand, though, the advice is of no *use* when it comes to making a decision, and would likely prompt something like the following reply: ‘That’s no help! I don’t even know whether there is asbestos in *this* wall, let alone the other walls in the neighbourhood!’ Since I haven’t made the judgment that the wall before me does not contain asbestos, I’m obviously in no position to make the stronger judgment that no walls in the neighbourhood contain asbestos, so I’m left in the same predicament as before.

One could point out, at this juncture, that there’s a distinction between the actual distribution of asbestos in the neighbourhood and what I might reasonably judge the distribution to be. One might claim that, in order to reasonably judge that there is no asbestos in any wall in the neighbourhood I don’t need to settle the question of whether there is asbestos in the wall before me. Depending on what we think is involved in ‘settling a question’ this claim may be correct. But does this mean that your advice really is helpful after all? I think it’s pretty clear that it doesn’t.

If I’m in a position to reasonably judge that no walls in the neighbourhood contain asbestos then even if I haven’t *settled the question* of whether the wall before me contains asbestos I must at least be in a position to reasonably judge that it doesn’t. The wall before me is one of the walls in the neighbourhood. So here’s the problem: If I can’t reasonably judge that the wall before me does not contain asbestos then I can’t reasonably judge that no walls in the neighbourhood contain asbestos. And if I *can* reasonably judge that the wall before me does not contain asbestos then there is no reason to care about the other walls in the neighbourhood. Either way, your advice is no use to me – it’s either impossible to act on, or it’s redundant.

Turning back to the notion of risk, suppose you now offer me the following, ostensibly more sensible, piece of advice: ‘As long as there is a low risk that the wall contains asbestos, you should go ahead and drill’. On the modal account, however, this advice turns out to be *just as useless* as the advice about walls in the neighbourhood – and for essentially the same reason. On the modal account, as Pritchard concedes, if the wall contains asbestos then it follows that the risk that the wall contains asbestos is maximally high. By contraposition, if the risk that the wall contains asbestos is low it follows that the wall does not contain asbestos². Since I haven’t made the judgment that the wall before me does not contain asbestos, I’m obviously in no position to make the stronger judgment that that there is a low modal risk that the wall contains asbestos.

Here’s another way to look at it. On the modal account, what it means for there to be a ‘low risk’ that the wall contains asbestos is for the wall to not contain asbestos in any close worlds. So when you say ‘As long as there’s a low risk that the wall contains asbestos you should go ahead and drill’ this could be paraphrased as ‘As long as the wall does not contain asbestos in any close worlds you should go ahead and drill’. But surely that would invite the rebuke ‘That’s no help! I don’t even know if there

² In addition to the fact that it prevents the notion of risk from playing its action-guiding role, this claim is independently very implausible. Here are a few quick considerations: If one person says to me ‘There’s no asbestos in the wall’ and another person says to me ‘There’s a low risk that there is asbestos in the wall’ it’s very strange to think that it’s the *second* person who is making the stronger, more committal claim. Suppose I say ‘There’s no asbestos in the wall – in fact, there’s a low risk that there’s asbestos in the wall’. If the second part of this were logically stronger than the first then this should sound OK – but it really doesn’t. Finally, ‘There’s a low risk that there is asbestos in the wall, but there might be’ sounds perfectly felicitous while ‘There is no asbestos in the wall, but there might be’ does not. If the latter is logically weaker than the former, then it ought to be at least as acceptable.

is asbestos in the wall in the *actual* world, let alone the other close worlds'. If I can't reasonably judge that the wall does not contain asbestos in the actual world then I can't reasonably judge that the wall does not contain asbestos in any close worlds. And if I *can* reasonably judge that the wall does not contain asbestos in the actual world then why would I care about the other close worlds?

While I have presented this as a conversation, in which advice is offered and rebuffed, the central problem for the modal account can be put quite simply: It should be possible for me to reasonably judge that there is a low risk that the wall contains asbestos, and to use this in my deliberations, without being in a position to reasonably judge that the wall does not contain asbestos. This is crucial in order for the notion of risk to play its action-guiding role – but the modal account rules it out³.

In the example that we've been considering, we are of course assessing the risk of a state of affairs which either does or doesn't already obtain. One might think that things are different when we are assessing the risk of an event which might or might not happen in the *future*. But if we are willing to suppose that future contingents have truth values, then the same reasoning will go through, irrespective of whether the target proposition concerns the future or the present⁴. Suppose I'm worried that my application for a loan will be refused. On the modal account, there is a low risk, at the present time, that my application will be refused just in case it's true in all nearby possible worlds, at the present time, that my application won't be refused. But if it's true that my application *will* be refused at the actual world then, obviously, this condition cannot be met. As a result, we still have the same problem; I cannot reasonably judge that there is a low risk that my application will be refused unless I am in a position to reasonably judge that my application won't be refused.

2 Is Risk 'Objective'?

This completes the main argument of the paper but, before concluding, I will turn to an important issue that is raised by Pritchard's discussion. Pritchard, as we've seen, draws extensively upon a distinction between a risk judgment being *true* and a risk judgment being *reasonable*, given limited evidence. For most judgments – like judgments as to whether there's asbestos in the wall or whether my application will be refused etc. – this kind of distinction is very familiar and straightforward. It's one thing for a person's evidence to suggest that the world is a particular way and it's another thing for the world to objectively be that way. But when we make a risk judgment are we really attempting to describe the way the world objectively is?

³ Could Pritchard draw a distinction between the ordinary, action-guiding conception of risk and a more specialised notion of risk, which aligns with the conditions set out by the modal account, and is suited to certain kinds of theoretical work? As well as claiming that the modal account captures our ordinary risk judgments, Pritchard also has certain theoretical roles in mind for modal risk – such as figuring in the analysis of knowledge (for instance, Pritchard, 2016, section III) – and it may be that these two aspects of his view can be teased apart. I don't propose to investigate this in detail here, but if, as I've argued, our ordinary risk judgments don't have the truth conditions laid down by the modal account, then we should be cautious about using the language of risk to describe these conditions, even in more technical settings.

⁴ And what if we *deny* that future contingents have truth values? This option would not, I think, offer any immediate relief for the modal theorist. On the contrary, if future contingents cannot be assigned truth values then, unless the modal account is modified in some way, it would appear to predict that they cannot be assigned risks either.

Pritchard is clear that the modal account treats risk ‘as an objective feature of the world’ (Pritchard, 2022, p10). This is in keeping with earlier work in which he describes risk as an ‘objective phenomenon’ (Pritchard, 2015, p440). While it’s not entirely clear to me what Pritchard intends with these remarks, one thing that he could be claiming is that, once we fix upon a given world and time, risk is a *monadic property* of propositions, like metaphysical necessity or possibility. If one person judges that it is necessary that gold is an element and another person judges that it is not necessary that gold is an element then only one of these judgments could be true – though they could both be reasonable given the evidence on which they are based. But it’s clear, on reflection, that risk cannot work like this.

Suppose a fair coin is flipped and A and B have both placed a bet on heads. A is ignorant of the result of the flip and judges that there is a high risk that the coin landed tails. B, on the other hand, sees the coin land heads and judges that there is a low risk – or no risk – that the coin landed tails. If risk is a monadic property of a proposition then the proposition that the coin landed tails could only have one of these properties, and only one of these judgments could be true. But which one? Presumably we wouldn’t want to say that A is the person who is getting things right, since B clearly has superior evidence. If I’ve also bet on heads and I want to know about the risk that the coin landed tails it’s clear that B is the person I ought to ask and not A.

So we can’t maintain that A is the one making the correct risk judgment here. But if B is the person who is correctly judging the risk, then this would set a troubling precedent. If a proposition P is false then a person who has ideal or complete evidence regarding P would judge that the risk of P is maximally low. And if a proposition P is true then a person with ideal or complete evidence regarding P will judge that the risk of P is maximally high. If we treat these judgments as authoritative then risk effectively collapses into truth – any true proposition will be at maximal risk and any false proposition will be at minimal risk.

We only got into this bind, however, because we were attempting to treat risk as if it were a monadic property attaching to a proposition. The alternative is to treat it as a *relation* between a proposition and a body of evidence. On this way of thinking, A’s judgment and B’s judgment could *both be true* – the proposition that the coin landed tails has a high risk relative to A’s evidence and a low risk relative to B’s evidence⁵. Risk can still be ‘objective’ in the sense that the level of risk conferred

⁵ B, upon overhearing A’s judgment, might say something like ‘A’s got it wrong – there’s no risk that the coin landed tails’. If that’s a sensible thing for B to say – and it would seem to be at first blush – then it puts some pressure on the claim that A’s and B’s judgments are both true. This is analogous to the so-called ‘eavesdropper’ cases that have bedevilled the literature on epistemic modals (see, for instance, Egan, Hawthorne and Weatherson, 2005, Hawthorne, 2007, von Fintel and Gillies, 2008, MacFarlane, 2011, Phillips and Mandelkern, 2020). To take a familiar example, suppose Holmes and Watson are at Baker St using a wire to listen in on a conversation between Moriarty and his underlings. If Moriarty says ‘Holmes might be in Paris right now attempting to foil our plans’ it would seem appropriate for Watson to remark ‘That’s not true’.

There is considerable disagreement as to what, if anything, the eavesdropper cases show about the semantics of epistemic modals. But what’s important for present purposes is that no-one, as far as I’m aware, has concluded from these cases that epistemic necessity and possibility must be monadic properties of propositions, possessed independently of anyone’s evidence or perspective. Pritchard (2015, p440) appeals to something close to an eavesdropper case in support of his claim that risk is ‘objective’. He writes: ‘If one is unaware of the inherent propensity for dynamite to explode when tossed in the air, then one will not judge that juggling with dynamite is a risky activity. Nevertheless it clearly *is* a risky activity...’ But Pritchard doesn’t explain in detail what he takes the case to show, and doesn’t draw the connection with the literature on epistemic modals. The coming material on the ‘semi-objective’ reading of risk judgments can, I think, straightforwardly explain the relevant intuitions.

upon a proposition by a body of evidence is not determined by anyone's views or opinions on the matter, but it can't be objective in the sense of being independent of evidence.

Risk judgments can, of course, be explicitly relativised to particular bodies of evidence. That is, it is perfectly felicitous to say something like 'Given what I've heard, there's a high risk of food poisoning at this restaurant', 'Given what we know about this medication, there's a low risk that you would have an adverse reaction' etc. While it's possible that unqualified risk judgments express something different or more 'objective' than judgments like these, an alternative hypothesis is that a body of evidence *always* figures in the truth conditions for a risk judgment – it's just that there are some cases in which the evidence need not be specified, as it will be clear from the context. When someone makes an unqualified risk judgment it would often be natural to assume that the relevant evidence is the evidence possessed by the speaker.

We have arrived at a picture on which all risk judgments involve a body of evidence, which can be made explicit in the judgment but is otherwise supplied by the context and would, as a default, be the evidence possessed by the speaker⁶. This is, I think, a good first attempt at understanding how risk judgments work – but it can't be the full story. Imagine now that, unbeknownst to A and B, the coin that is flipped is actually double-headed. In this case I think there is an interpretation of A's utterance 'There's a high risk that the coin landed tails' on which it's *false*. What this shows, I suggest, is that risk judgments are sometimes assessed relative to a set of propositions that extends *beyond* the evidence of the speaker – or any salient body of evidence – and includes certain significant facts about the world. The reason that A's judgment turns out to be false on this 'semi-objective' reading is that there is no risk that the coin has landed tails relative to A's evidence *and* the truth about whether the coin is double-headed⁷.

⁶ There are several ways in which evidence relativity can be built into the semantics for risk judgments. On a *contextualist* approach, a judgment like 'There's a high risk that the coin landed tails' will express different propositions in different contexts of utterance. Roughly speaking, in a context of utterance *c* this sentence will express the proposition *given E_c there is a high risk that the coin landed tails* where *E_c* is the evidence supplied by *c* (which would typically be the evidence possessed by the speaker). In this case the judgments made by A and B will not contradict one another – as they express propositions about different bodies of evidence. On a *relativist* approach 'There's a high risk that the coin landed tails' will express the same proposition in any context of utterance, but the truth of the proposition can only be assessed relative to a body of evidence supplied by the context in which the assessment is taking place. In this case A's judgment and B's judgment will contradict one another, in that they could never be simultaneously true relative to a single context of assessment, but they could each be true relative to A's and B's respective contexts of assessment.

These approaches mirror two prominent options that have been explored in the literature on epistemic modals. While I cleave to a broadly contextualist position in the main text, this is largely for illustration – my main conclusion is just that any truth conditions for risk judgments must implicate evidence *in some way*. Whether this ultimately leads us down a contextualist path, or in some other direction, is not something on which I mean to take a principled stand. What I do think, though, is that risk judgments ought to receive the same broad semantic treatment as epistemic modals, where contextualism could be described as the standard or default view (von Fintel and Gillies, 2011, Anderson, 2014, p597, Phillips and Mandelkern, 2020, p4). For contextualism about epistemic modals see for instance Kratzer (1977), DeRose (1991), von Fintel and Gillies (2008, 2011), Dowell (2017). For relativism see for instance Egan, Hawthorne and Weatherson (2005) and MacFarlane (2011). For other alternatives, such as non-truth-conditional approaches, see for instance Yalcin (2011) and Moss (2015, partic. chaps. 2-3).

⁷ What I have termed 'semi-objective' readings of risk judgments are analogous to what Hawthorne (2007) calls 'danger-theoretic' readings, and Dorr and Hawthorne (2013) call 'constrained' readings, of epistemic modals. Hawthorne (2007) suggests that the availability of such readings can help to explain eavesdropper cases in a way

But even in the double-headed coin case, a purely speaker-centric evidential reading of A's judgment is still available. Consider the following dialogue:

A: There's a high risk that the coin has landed tails.

Me: But it could be a double-headed coin.

A: Yes, it *could* be – but there's a high risk that it's a fair coin and that it landed tails.

Even though the coin is in fact double-headed, I don't think that we would take exception to anything that A says here. In light of the overall conversation it is natural to take A's risk judgments as pertaining specifically to his own evidence. One also has the option of forcing a speaker-centric evidential reading by prefacing a risk judgment with 'Given my current evidence...' or words to that effect. If A says 'Given my current evidence, there's a high risk that the coin landed tails' then that is clearly true, irrespective of whether the coin is in fact double-headed.

Even on a speaker-centric evidential reading, it may still be possible for a risk judgment to be reasonable, though false. But these cases won't be characterised by an ignorance of worldly facts (the risk is, in effect, relativised to the facts of which the speaker is *not* ignorant). Rather, these will be cases in which the speaker makes a (reasonable) mistake about the bearing of their own evidence. So if I reasonably judge 'There is a low risk that there is asbestos in the wall' then this utterance could be false – if I'm misunderstanding the significance of my own evidence. But it won't be made false by the fact – if it is one – that there is actually asbestos in the wall.

On a semi-objective reading, however, the truth and reasonableness of a risk judgment can come apart in something closer to the way that Pritchard envisages. On this reading, the truth of 'There's a low risk that there is asbestos in the wall' will be hostage to certain facts that go beyond my evidence, and about which I could make a reasonable mistake⁸. But even on a semi-objective reading, the truth about whether there is asbestos in the wall will not *itself* enter into the set of facts relative to which the risk is assessed. Otherwise my judgment would entail that there is no asbestos in the wall and, for the reasons given in the last section, there is no reading of a risk judgment on which that entailment holds. Low risk possibilities do, on occasion, turn out to be actual.

Acknowledgements

This paper was presented at a *Varieties of Risk* seminar at the University of Edinburgh in September 2022. Thanks to all of those who participated in the discussion – Antony Duff, Philip Ebert, Claire Field, Giada Fratantonio, Mario Günther, Matt Jope, Sandra Marshall, Lilith Newton, Angie O'Sullivan, Petronella Randell, Rafal Urbaniak, Xintong Wei and Crispin Wright. Particular thanks to Antony Duff and to Crispin Wright for extensive email correspondence on the ideas in this paper. Thanks also to

that is compatible with a broadly contextualist semantics. A similar strategy could be used to explain eavesdropper cases for risk judgments – but I won't pursue this here.

⁸ Risk judgments might also permit readings that are still 'evidential' but non-speaker centric, in that the relevant evidence might include evidence possessed by others in a contextually salient group, or even evidence that the speaker *ought to have possessed* (for corresponding discussion about epistemic modals see for instance DeRose, 1991, von Fintel and Gillies, 2011, Anderson, 2014, section 3). If the reasonableness of a risk judgment is always determined by the speaker's *own* evidence, then this could present another way in which the reasonableness and truth of a risk judgment could come apart – and one which, once again, needn't involve any misjudgment, by the speaker, of what their own evidence shows.

three anonymous referees for this journal. Work on the paper was supported by the AHRC grant *Varieties of Risk* (AH/T002638/1).

References

- Anderson, C. (2014) 'Fallibilism and the flexibility of epistemic modals' *Philosophical Studies* v167(3), pp597-606
- DeRose, K. (1991) 'Epistemic possibilities' *Philosophical Review* v100(4), pp581-605
- Dorr, C. and Hawthorne, J. (2013) 'Embedding epistemic modals' *Mind* v122(488), pp867-913
- Dowell, J. (2017) 'Contextualism about epistemic modals' in Ichikawa, J. ed. *The Routledge Handbook of Epistemic Contextualism* (London: Routledge)
- Ebert, P., Smith, M. and Durbach, I. (2020) 'Varieties of risk' *Philosophy and Phenomenological Research* v101(2), pp432-455
- Egan, A., Hawthorne, J. and Weatherson, B. (2005) 'Epistemic modals in context' in Preyer, G. and Peter, G. eds. *Contextualism in Philosophy* (Oxford: Oxford University Press)
- Fratantonio, G. (2021) 'Evidence, risk and proof paradoxes: Pessimism about the epistemic project' *International Journal of Evidence and Proof* v25(4), pp307-325
- Gardiner, G. (2020) 'Relevance and risk: How the relevant alternatives framework models the epistemology of risk' *Synthese* v199(1-2), pp481-511
- Hawthorne, J. (2007) 'Eavesdroppers and epistemic modals' *Philosophical Issues* v17(1), pp92-101
- Kratzer, A. (1977) 'What 'must' and 'can' must and can mean' *Linguistics and Philosophy* v1(3), pp337-355
- MacFarlane, J. (2011) 'Epistemic modals are assessment-sensitive' in Egan, A. and Weatherson, B. *Epistemic Modality* (Oxford: Oxford University Press)
- Moss, S. (2018) *Probabilistic Knowledge* (Oxford: Oxford University Press)
- Newton, L. (2022) 'Epistemic risk and epistemic anxiety' *Synthese* v200
- Phillips, J. and Mandelkern, M. (2020) 'Eavesdropping: What is it good for?' *Semantics and Pragmatics* v13, 19:1-35
- Pritchard, D. (2015) 'Risk' *Metaphilosophy* v46(3), pp436-461
- Pritchard, D. (2016). 'Epistemic risk' *Journal of Philosophy*, v113(11), pp550-571
- Pritchard, D. (2022) 'In defence of the modal account of legal risk' *Synthese* v200(4)
- Smith, M. (2018) 'When does evidence suffice for conviction?' *Mind* v127(508), pp1193-1218
- von Fintel, K. and Gillies, T. (2008) 'CIA leaks' *Philosophical Review* v117(1), pp77-98
- von Fintel, K. and Gillies, T. (2011) "'Might' made right' in Egan, A. and Weatherson, B. eds. *Epistemic Modality* (Oxford: Oxford University Press)

Yalcin, S. (2011) 'Nonfactualism about epistemic modals' in Egan, A. and Weatherson, B. eds. *Epistemic Modality* (Oxford: Oxford University Press)