RISKY BELIEF

Martin Smith

University of Edinburgh

In this paper I defend the claim that justification is closed under conjunction, and confront its most alarming consequence – that one can have justification for believing propositions that are unlikely to be true, given one's evidence.

Key Words: Justification, Belief, Conjunction Closure, Single Premise Closure, Risk Aggregation

I   THE RISK AGGREGATION ARGUMENT

Closure principles in epistemology state that some positive epistemic status – such as knowledge or justification – is closed under a given logical operation or relation. One of the most discussed closure principles states that justification is closed under conjunction:

If one justifiably believes P_1 and justifiably believes P_2 … and justifiably believes P_n, then one has justification for believing $P_1 \wedge P_2 \wedge … \wedge P_n$.

Call this *Justification Conjunction Closure* (JCC)[1]. While many have found this principle attractive[2], there is a powerful reason for rejecting it – conjunction *aggregates the risk of falsity*. That is, a conjunction (in general) is at a higher risk of falsity than its individual conjuncts.

Suppose I've invited 10 people to a party: Peter, Shelley, Shannon, Timmy, Benny, Summer, Teeway, Garrett, Pepita and Pam. Suppose each person has replied saying that they will attend, and I take them at their word. Suppose I know each person to be reliable and trustworthy – so I'm justified in believing that Peter will attend, in believing that Shelley will attend, in believing that Shannon will attend etc. Would I be justified in believing the *conjunction* of these propositions – that Peter and Shelley and Shannon and Timmy and Benny and Summer and Teeway and Garrett and Pepita and Pam will attend? The problem is that although there is a low chance that Peter won't turn up, and a low chance that Shelley won't turn up etc., the chance that *some person or other* won't turn up may be quite high. Although each invitee is individually reliable, it is quite probable that a mishap will befall at least *one of them* – illness, car trouble, family emergency etc. –

---

[1] Conjunction closure principles of this kind were first thrust into the spotlight as a result of the lottery and preface paradoxes – see Kyburg (1961, pp197-198) and Makinson (1965). Both paradoxes are connected to the primary argument against JCC that I consider here, but I won't discuss them (not in the body text at any rate – but see n3 and n15).

[2] Defences of JCC are mounted by, for instance, Kaplan (1995, 1996, chaps. 3 and 4), Ryan (1996), Nelkin (2000) and Smith (2016, chap. 4, 2018, forthcoming). The idea that 'idealised' belief sets are closed under conjunction is enshrined in many formal representations of belief (see for instance Hintikka, 1962, Routley and Routley, 1975, Spohn, 2012, chap. 4). Principles akin to JCC are also a common feature of systems of nonmonotonic inference (see for instance Kraus, Lehmann and Magidor, 1990, rule (7) on p179).

and they will be prevented from attending. As a result, if I were to believe that Peter and Shelley and Shannon and ... will attend, then I would be at a high risk of believing something false.

Let's attach some numbers to the example. Suppose that, for each invitee, my evidence makes it 90% likely that they will attend – so it's 90% likely that Peter will attend, it's 90% likely that Shelley will attend and so on. Suppose, further, that these propositions are independent – so the probability of any one person attending is unaffected by who else attends. In this case, the probability, given my evidence, that Peter and Shelley and Shannon and … will attend is approximately *35%*. That is, there is about a 65% chance, given my evidence, that this proposition is false – in which case, presumably, it wouldn't be a sensible thing to believe.

Consider the following:

> If a proposition P is more than 50% likely to be false, given one's evidence, then one does not have justification for believing P.

This principle seems difficult to deny – how could one possibly have justification for believing something if it's more likely to be false than true? If anything, one would seem to be better off believing its *negation*. Thorn (2017) remarks, of such a principle, that it is '… about as secure as any conceptual claim (outside of logic and mathematics) could be. It is nigh absurd to claim that an agent may rationally believe a proposition in a case where the agent's rational personal probability for the proposition is below 0.5…' (Thorn, 2017, pp105-106). If we accept this principle then I must lack justification for believing that Peter and Shelley and Shannon and ... will come to the party, and JCC fails. Call this the *risk aggregation* argument against JCC[3,4].

---

[3] Versions of the risk aggregation argument can be found in, for instance, Kyburg (1961, p197), Foley (1993, chap. 4, 2009), Christensen (2004, chap. 3) and Thorn (2017). For largely historical reasons, the argument is often woven into discussions of the *preface paradox* – but it can be treated as separate. The preface paradox begins in much the same way as the risk aggregation argument – that is, we imagine a set of propositions P_1, P_2, …, P_n that one justifiably believes and which, given one's evidence, are individually likely to be true, but unlikely to all be true (in the classical presentation, these are supposed to be claims in some well researched non-fiction book, but it doesn't especially matter what they are – they could equally well be propositions about who will attend an upcoming party etc.) Given this set up, JCC implies that one must have justification for believing $P_1 \land P_2 \land … \land P_n$, (justification for believing that 'every claim in the book is true') even though this has a low evidential probability. Instead of pushing back on this by simply highlighting the low evidential probability of $P_1 \land P_2 \land … \land P_n$, in the preface paradox we now introduce a new supposition – namely, that one is justified in believing $\sim P_1 \lor \sim P_2 \lor … \lor \sim P_n$ (justified in believing that 'some claim in the book is false'). By another application of JCC we can then derive the absurd result that one must have justification for believing an out-and-out contradiction – $(P_1 \land P_2 \land … \land P_n) \land (\sim P_1 \lor \sim P_2 \lor … \lor \sim P_n)$ ('every claim in the book is true and some claim in the book is false').

For those who are primarily interested in setting up a reductio of JCC, it's not clear that this extra detour is of much dialectical help. The claim that one lacks justification for believing $P_1 \land P_2 \land … \land P_n$ (because it is more than 50% likely to be false, given one's evidence) is all that is needed to refute JCC – and it seems a lot more secure than the claim that one *has* justification for believing $\sim P_1 \lor \sim P_2 \lor … \lor \sim P_n$. This latter claim could only be motivated, I suspect, by a view on which justification is *equated* with high evidential probability, which anyway entails the former claim, but is considerably stronger. The risk aggregation argument can also be used to put pressure on a conjunction closure principle for knowledge and, in this setting, is more often set out as a free-standing argument – see, for instance, Hawthorne (2004, p47), Lasonen-Aarnio (2008) and Leuenberger and Smith (2021).

When it comes to formulating a response to this argument, our options are very limited. One could quibble over whether I would really be justified in believing that Peter will attend, that Shelley will attend etc. – perhaps I don't really have enough evidence for these propositions in order to justifiably believe them. This strikes me as a rather implausible view though – and, more importantly, it fixates on something that is incidental to the argument. We could easily imagine even stronger evidence in favour of these propositions – maybe it's the night of the party and each invitee has called to tell me they're on the way etc. – and then imagine that there are more invitees in order to restore the intended structure of the case.

If one adopted an infallibilist stance, on which we can only be justified in believing propositions that are made *certain* by our evidence, then no argument of this kind could get up and running of course. This would seem a drastic step to take in order to salvage JCC however. It's natural to worry that an infallibilism of this kind would lead directly to scepticism[5] in which case JCC would only end up being true in a trivial way. If we are not prepared to countenance infallibilism, then the only alternative for saving JCC is to insist that we can, sometimes, have justification for believing propositions that are unlikely to be true, given our evidence. And, as suggested above, this seems like a nonstarter.

II   WHAT NOW?

Rather than straightaway trying to devise a strategy for resisting the risk aggregation argument, I will turn instead to a question that is raised by its *acceptance*. If the party case really is a counterexample to JCC then, although I justifiably believe that Peter will attend, and justifiably believe that Shelley will attend, and justifiably believe that Shannon will attend etc., I should *refrain*

---

[4] In talking about 'risk', I have been taking for granted a standard probabilistic picture on which the *risk* that a proposition P is false, given one's evidence, is equal to the *probability* that P is false. I am inclined to think that this is not the only legitimate conception of risk however – there are others, including conceptions on which a conjunction *cannot* be at a higher risk of falsity than each of its conjuncts (see Ebert, Smith and Durbach, 2020, sections 4 and 5). While this is potentially relevant to the issues under discussion, there is no quick response to the risk aggregation argument in the offing here. It's clear that the probabilistic picture does present *one* perfectly legitimate conception of risk. Furthermore, in spite of the name, we don't actually need to use the term 'risk' in setting up the risk aggregation argument – it could be formulated directly in terms of evidential probabilities.

[5] Any proposition that is a part of one's evidence will, of course, count as certain, given one's evidence. As a result, one way to halt this slide into scepticism is to adopt a generous conception of when a proposition might count as part of one's evidence, such as that defended by Klein (1995) or Williamson (2000, chaps. 9, 10). Another option is to adopt a contextualist treatment of certainty ascriptions, such as that defended by Beddor (2020), which allows for a proposition to be truly described as certain, in some contexts, even if one hasn't eliminated all possibilities of error (Beddor proceeds to analyse evidence in terms of certainty, reversing the usual order). Either of these approaches would make certainty more attainable, and infallibilism more viable. But even if it doesn't lead to scepticism, I'm inclined to think that infallibilism would still leave us with a need for some new property – signalling that a belief is not a legitimate target for criticism or condemnation – that *is* capable of attaching to beliefs that fall short of certainty. If the infallibilist maintains that this property is closed under conjunction, then the risk aggregation argument immediately resurfaces and will have to be dealt with anew. If, on the other hand, the infallibilist concedes that this property is not closed under conjunction, then it looks as though the bump is simply being moved from one part of the carpet to another. I won't pursue this further here – for more discussion (and criticism) of infallibilism about justification see Smith (2018, sections 2 and 4).

from believing that Peter and Shelley and Shannon and … will attend, and should at the very least suspend judgment about this proposition. Here's the question: How exactly do I *go about* following this piece of advice? This might at first seem an odd thing to ask – after all, I'm not really being advised to *do* anything. At most, I'm being advised to *not* do something – namely, hold the conjunctive belief. On closer inspection though, following this advice is anything but straightforward.

I currently believe that the front door and the back door are locked, having just checked them. While I clearly remember checking the front door and forming the belief that it's locked, and I clearly remember checking the back door and forming the belief that it's locked, I don't remember ever forming the conjunctive belief that the front door *and* the back door are locked. That is, I don't remember going through any *inference* in which I put these two propositions together and derived the conjunction – yet I hold the conjunctive belief nonetheless. One might suggest that there was an implicit or unconscious inference that took place – that the conjunctive belief arrived unbidden, as it were. Perhaps. And yet my own impression, in reflecting upon this simple case, is that the conjunctive belief never really 'arrived' at all – the state of believing the conjunction is just *the same thing* as the state of believing the individual conjuncts.

Furthermore, if I were to say 'The front door is locked' and then say 'The back door is locked' you wouldn't hesitate to describe me as believing that the front door *and* the back door are locked. It's not as though you need additional evidence to the effect that I've gone through some inferential procedure before you can make the belief report – it seems that you already have all of the evidence that you need. If you are willing to ascribe the belief that the front door is locked and ascribe the belief that the back door is locked, then you will be willing to ascribe the belief that the front door and the back door are locked.

Although there are more conjuncts at play in the party example, many of the above observations would still seem to apply. If I had asserted that Peter is coming to the party and that Shelley is coming to the party and that Shannon is coming to the party etc. you would naturally describe me as believing that Peter and Shelly and Shannon and … are coming to the party – without needing to ascertain whether I'd gone through an inference[6]. And *I* would regard myself as believing this too, even if all that I really remember is receiving each reply and forming each belief in turn. In both the door-locking and party examples, we have a set of beliefs that concern a single topic of interest, and that were formed under very similar circumstances within a relatively short period of time – and, in both examples, believing the conjunction appears to be *automatic*.

If I'm not to believe that Peter and Shelley and Shannon and … are coming to the party, as recommended by the JCC-denier, it seems that I would after all have to *do* something in order to prevent this – though it's unclear what would be effective. Perhaps I could mentally affirm to myself that Peter is coming and that Shelley is coming etc. but refuse to mentally affirm that Peter and

---

[6] Not only would you describe me as believing that Peter and Shelley and Shannon and … are coming to the party, you would also describe me as having *asserted* this – even if I only ever uttered the words 'Peter is coming to the party', 'Shelley is coming to the party' etc. and never connected them together with 'and's. Some philosophers might assign significance to whether one actually voices the word 'and', but our ordinary practice of reporting assertions almost certainly does not. Suppose you sue me for defamation, claiming that I called you a liar and a cheat. I could hardly defend myself by conceding that I called you a liar and conceding that I called you a cheat but insisting that I never asserted the conjunction!

Shelley and Shannon and … are coming, and pledge that I'm really suspending judgment about this. I'm unsure whether this would make any difference to what I actually believe but, even if it did, can we really accept that *this* is where the risk aggregation argument ultimately leads – to some internal monologue that I have to recite?

Furthermore, it is not just the conjunction of all *ten* propositions that has an evidential probability of below 50% – in fact, given our stipulations, conjoining any seven or more will be enough for that. Am I supposed to consider *every such conjunction* and assure myself that I'm suspending judgment about it – that I'm suspending judgment on whether Peter and Shelley and Shannon and Timmy and Benny and Summer and Teeway are coming, and I'm suspending judgment on whether Shelley and Timmy and Summer and Teeway and Garrett and Pepita and Pam are coming, and so on? It turns out that there are no fewer than 176 different conjunctions in this category!

What if I were to actively *deny* the proposition that everyone will show up – repeat to myself that it's *not the case* that Peter and Shelley and Shannon and … will attend? The first thing to note about this is that it goes beyond any advice generated by the risk aggregation argument. The advice, once again, is to *refrain from believing* that Peter and Shelley and Shannon and … will attend. To believe the negation of this proposition would seem an excessive way to try and bring this about. And even if I do find it reasonable to deny the proposition that everyone will show up, it doesn't seem as though I could apply this strategy to all of the propositions that I am supposed to refrain from believing. Consider again a proposition involving just seven invitees – say the proposition that Shelley and Timmy and Summer and Teeway and Garrett and Pepita and Pam will all attend. Am I supposed to deny this proposition too – even though it's about 48% likely to be true? Finally, just because I deny the proposition that everyone will show up, that doesn't mean that I can't *also* believe it. It's not at all impossible for a person to have conflicting attitudes towards a single proposition. If I were to assert that Peter will attend and that Shelley will attend and that Shannon will attend… and then assert that it's not the case that Peter and Shelley and Shannon and … will attend, most people would, I imagine, find this very confusing – and might well interpret me as *both believing and denying* that everyone will attend[7].

---

[7] We don't just ascribe beliefs on the basis of what people *assert* of course – but also on the basis of how they *behave*. But it doesn't look as though my behaviour could ever provide a basis for ascribing each of the conjunct beliefs without also providing a basis for ascribing the conjunctive belief – it's plausible that any behaviour which could be rationalised by the former ascription could *also* be rationalised by the latter. Suppose, for instance, that I set ten places at the dinner table. This could be interpreted as ten separate actions rationalised, in turn, by each of the conjunct beliefs – he believes that Peter is coming, so he's setting a place for him, he believes that Shelley is coming, so he's setting a place for her etc. Alternately, this could be interpreted as a single act rationalised by the conjunctive belief – he believes that all ten people are coming, so he's setting ten places at the table.

I could, of course, perform some actions which suggest that I believe each of the conjuncts (like setting ten places at the table) and other actions which suggest that I deny the conjunction or believe something inconsistent with it (like only getting enough food for nine). This would be a bit like asserting each of the conjuncts while asserting the negation of the conjunction, and would I think be just as perplexing for an observer. Rather than providing a basis for ascribing the conjunct beliefs without ascribing the conjunctive belief, this would, for the reasons just given, be better seen as grounds for ascribing a *conflicted* attitude towards the conjunction.

Regardless of what one thinks of the present example, there are clear cases in which believing one proposition means that one automatically believes another distinct proposition, and no amount of mental effort or self-talk is going to change that. We have been focussing on cases of *Conjunction Introduction*, but consider for a moment its counterpart – *Conjunction Elimination*. If I believe a conjunction – say that snow is white and grass is green – then I also believe the conjuncts – that snow is white and that grass is green. I could say to myself 'I believe that snow is white and grass is green, however I'm suspending judgment on whether grass is green' – but that wouldn't make it *true*. And if I say to myself 'I believe that snow is white and grass is green, but I deny that grass is green' that would, at most, make it the case that I both affirm *and* deny the proposition that grass is green. Believing that grass is green is just *part of what it is* to believe that snow is white and grass is green.

If the above reflections are right then simply *telling myself* that I'm denying, or suspending judgment about the proposition that Peter and Shelley and Shannon and … will attend is not the answer. One way in which I really could stop believing this proposition is, of course, to stop believing one of the conjuncts – perhaps I could stop believing that, say, Shelley will attend and suspend judgment on this instead. If Shelley got in touch to say that she might not make it after all, or if I heard that she's been unwell etc. then it might be perfectly reasonable to give up my belief that she'll attend. But to stop believing that Shelley will attend *simply to avoid believing a conjunction* seems absurd in the extreme – and quite unfair to her. Furthermore, to stop believing just one of the conjuncts wouldn't be enough – to ensure that I don't believe a conjunction that is too risky I would have to choose four invitees in which to lose faith, even though they all gave positive replies and are just as reliable as the others. Finally, the idea that I'm obliged to do this is inconsistent with the claim that I *justifiably believe* that Peter will attend, that Shelley will attend, that Shannon will attend etc. – which the JCC-denier accepts. If I were only permitted to form six of these beliefs then, once I'd received the first six replies and formed the corresponding beliefs, I would not be justified in forming any more, even when further replies came in. Once I have formed all ten beliefs, and it is granted that they are all justified, the opportunity to make this kind of recommendation has passed.

While the risk aggregation argument might seem to clinch the case against JCC, my aim in this section has been to sound an initial note of caution. This argument is not merely a piece of abstract reasoning – it generates concrete advice about what one ought to believe. While this advice may concern a hypothetical scenario, it can, and will, carry over into many actual situations as well. At best, it is unclear how this advice might be followed. At worst, the advice may actually be *impossible* to follow, which suggests something seriously amiss with the argument that gave rise to it[8].

---

[8] One might insist that *any* advice as to what one should or shouldn't believe is bound to be problematic on the grounds that our beliefs are not under our direct voluntary control (the classic defence of 'doxastic involuntarism' is given by Alston, 1985, section IV, 1988 – and the issue has been much discussed since then). Whatever truth there is in this, if the arguments I offer here are on the right track, then the advice that issues from the risk aggregation argument raises a special concern. The worry is not (just) that we have no notion of how to get ourselves into the recommended belief state – the worry is that, in some cases, the recommended belief state is not even a possible way that one's beliefs *could be*. Put differently, the worry is not that the advice is excessively demanding, but that the advice is *incoherent*.

III   BELIEVING CONJUNCTIONS

Opponents of JCC seem to take for granted a particular picture of what it is to believe a conjunction – a picture on which this always involves some kind of deliberate mental act, over and above believing the conjuncts. It is difficult to find any explicit endorsement of such a picture however – though David Armstrong may come close in *Belief, Truth and Logic*, where he seems to suggest that believing a conjunction is a matter of 'bringing together' one's beliefs in the conjuncts (Armstrong, 1973, chap. 8, section II). But even Armstrong concedes that an inference is not *always* needed for a conjunctive belief to exist (p107).

At the other end of the spectrum, we have a view defended by Richard and Valerie Routley, who claim that there is no genuine distinction between believing conjuncts separately and believing them conjointly (Routley and Routley, 1975, §9). As they put it 'The view of conjunction as a sort of mental glueing operation which one can fail to perform is an error. It is not as though believing (A & B) is something over and above believing A and believing B; it is not a further act…' (Routley and Routley, 1975, p211, Douven 2002, p395 and Weiss, 2019, pp44-45 also express sympathy with such a view). In effect, Routley and Routley endorse a conjunction closure principle not for justification, but for *belief itself*:

   If one believes P_1 and believes P_2 … and believes P_n, then one believes $P_1 \wedge P_2 \wedge … \wedge P_n$.

But this principle – which we could call BCC – goes too far. While the Routley/Routley picture fits well with the party and door-locking examples, there are other cases in which it breaks down. Among other things, the principle implies that anyone who holds inconsistent beliefs must also believe a contradiction – which seems not to be correct[9].

David Lewis reports that he once believed that Nassau St ran roughly east-west, believed that the railroad ran roughly north-south and believed that Nassau St and the railroad were roughly parallel (Lewis, 1982, p436). But Lewis never believed the impossible proposition that Nassau St runs roughly east-west *and* the railroad runs roughly north-south *and* Nassau St and the railroad are roughly parallel. The most natural explanation for what is going on here is actually an Armstrongian one: While Lewis believed these three propositions individually, this was before he 'brought them together' within his own mind. Indeed, when he first came to entertain the conjunction of these propositions, his response was to stop believing the first two (and replace them with the belief that both Nassau St and the railroad run roughly northeast-southwest).

The risk aggregation argument also generates advice as to what one should and shouldn't *assert* – and, as suggested in footnote 6, this advice is beset by the very same concern, even though assertions are, presumably, subject to direct voluntary control. In some cases, asserting a series of propositions serves as *one way* of asserting their conjunction, in which case there is simply no possibility of doing the former without doing the latter, regardless of whatever control we enjoy over our own assertions.

[9] Routley and Routley insist that this principle *does* apply in cases of inconsistent beliefs (pp213-214, §11). Indeed, they come close to using their account of conjunctive belief as a means of arguing from the (plausible) claim that people often hold inconsistent beliefs to the (less plausible) claim that people often believe contradictions. While I don't think the Routley/Routley account can ultimately be used in this way, I do think there are many cases of conjunctive belief that it captures more faithfully than the Armstrongian alternative.

Lewis takes this to be an illustration of how one's belief corpus can be fragmented or *compartmentalised*. In Lewis's view, compartments are like encapsulated belief sets, within a single mind, that may not manifest or activate simultaneously. While the beliefs in a single compartment are always integrated into a cohesive whole, integrating the beliefs of separate compartments requires mental activity and effort and may be an occasion for belief revision, as in the case of the beliefs about Nassau St and the railroad (Lewis, 1982, section II). If all putative counterexamples to BCC involved beliefs held in separate compartments, then that would suggest the following qualified version of the principle:

If one believes P_1 and believes P_2 … and believes P_n and these beliefs are all part of a single compartment, then one believes $P_1 \wedge P_2 \wedge … \wedge P_n$.

This principle is endorsed by Evnine (1999, section 7) and also, perhaps, by Stalnaker (1984, chap. 5).

While I am somewhat sympathetic to this principle, I won't attempt to defend it here, or to further explore the topic of compartmentalisation (for some further discussion see Schotch and Jennings, 1980, Stalnaker, 1984, chap. 5, Evnine, 1999, section 7, Spectre, 2018, Leuenberger and Smith, 2021). In a way, my aim here is more modest than to defend any qualified version of BCC. My only claim about BCC is that it has *some* genuine instances – that there are some cases in which one automatically counts as believing $P_1 \wedge P_2 \wedge … \wedge P_n$ whenever one believes P_1 and believes P_2 … and believes P_n. Provided these instances are of the right kind, this is all that will be needed to answer the risk aggregation argument against JCC.

IV   PSYCHOLOGICAL GAPS

Another closure principle that has attracted attention from epistemologists states that justification is closed under single premise deductive consequence:

If one justifiably believes P and P entails Q, then one has justification for believing Q.

We might call this *Justification Single Premise Closure* (JSPC)[10]. Unlike JCC, JSPC raises no worries about risk aggregation – if P entails Q then the risk that Q is false, given one's evidence, can be no higher than the risk that P is false[11]. Disputes over JSPC have tended to focus on a different sort of worry – that mundane, everyday propositions can entail substantial 'heavyweight' propositions, like the negations of radical sceptical hypotheses. It's obvious that I can justifiably believe, say, that I'm currently wearing shoes. But this entails, among other things, that I'm not a disembodied brain in a vat currently being supplied with an appearance as of wearing shoes. When it comes to this 'heavyweight' proposition, however, one might question whether I have justification for believing it

---

[10] It was Fred Dretske, in his paper 'Epistemic operators' (1970), who sparked epistemologists' continuing interest in single premise closure principles. While Dretske's main concern in the paper was to argue against a single premise closure principle for *knowledge*, he also considers, and rejects, principles that are much closer to JSPC.

[11] For an argument that single premise closure principles are in fact subject to risk aggregation concerns see Lasonen-Aarnio (2008). For discussion see Smith (2013) and Tang (2018).

– perhaps on the grounds that my evidence fails to rule out its relevant alternatives, or that my evidence seems to be neutral between its truth and falsity.

Whatever one makes of this argument against JSPC, co-opting the reasoning from section II in order to respond to it would not seem like a particularly promising strategy. That is, we cannot complain that a failure of JSPC in this case would give rise to impossible advice. It clearly *is* possible for one to believe that one is wearing shoes without also believing that one is not a disembodied brain in a vat. Whether this is *still* possible once one has considered both propositions in tandem and recognised the entailment is somewhat less clear, but there is at least an in principle *psychological gap* between the premise and the conclusion – and the advice to believe the premise while suspending judgment about the conclusion is at least *coherent*.

Not all single premise deductive entailments open up a psychological gap however. As discussed in section II, there is no gap between believing $P \wedge Q$ – that snow is white and grass is green, say – and believing Q – that grass is green. Rather, the latter belief is simply a part of the former. Even Armstrong denies that any mental act is required to take one from a belief in $P \wedge Q$ to a belief in Q – Conjunction Elimination, as he puts it, does not represent any 'movement of thought' (Armstrong, 1973, p107). As a result, it would make no sense to deny an instance of JSPC that involved Conjunction Elimination. If believing Q is part of what it is to believe $P \wedge Q$, then it would make no sense to grant that the latter belief is justified, while denying that the former is. If one believes Q without justification, then one is epistemically obliged to give the belief up. And yet, to stop believing Q is to stop believing $P \wedge Q$ – so how then can we maintain that the latter belief is epistemically permitted?

I briefly noted above some possible motivations for denying JSPC: In order to have justification for believing P, perhaps it's necessary that one's evidence rules out all of the relevant alternatives to P, or that it favours P over its negation. Perhaps it's possible that one's evidence rules out all of the relevant alternatives to P without ruling out all of the relevant alternatives to an entailed proposition Q, or possible that one's evidence favours P over its negation, but is neutral when it comes to an entailed proposition Q[12]. Here is the crucial point: If either of these motivations

---

[12] Talk of 'relevant alternatives' in epistemology also traces back to Dretske (1970) who defended a relevant alternatives account of knowledge. We are now perhaps accustomed to the practice of first defending an account of knowledge, or a condition upon knowledge, and then using this to guide discussion of closure principles. This isn't the structure of Dretske's original argument however. For Dretske (as I read him) the claim that knowledge is not closed under single premise deductive consequence is meant to be motivated largely by examples, with the relevant alternatives theory being introduced after the fact, as a way of systematically explaining the underlying structure that these examples share.

A relevant alternatives account of *justification*, which leads directly to the failure of JSPC, is outlined by Blome-Tillmann (2015). On Blome-Tillmann's view, we lack what he calls 'RA-justification' for believing the negations of radical sceptical hypotheses, but have an 'entitlement' to believe them nonetheless. For applications of the relevant alternatives framework to other kinds of positive epistemic status see Gardiner (2021).

Wright (2004, 2014), Pedersen (2009) and Coliva (2012) are amongst those who maintain that we can justifiably believe ordinary, everyday propositions, but lack justification for believing the negations of radical sceptical hypotheses, on the grounds that they are not supported or favoured by our evidence. According to these authors, although we are not permitted to adopt an attitude of belief toward the negations of radical sceptical hypotheses, we are permitted to adopt a weaker attitude, such as 'trust' or 'acceptance'.

were to apply in certain 'no-psychological-gap' cases – cases in which JSPC couldn't possibly fail – then this would cast doubt upon the motivation *itself*, even when it comes to cases that *do* involve a psychological gap.

Suppose that, once we've filled in the details of 'relevance' and 'ruling out', it turns out that there is a case in which one justifiably believes $P \wedge Q$ and in which one's evidence rules out all of the relevant alternatives to $P \wedge Q$, even though one's evidence fails to rule out all of the relevant alternatives to Q. What this would show is that, with 'relevance' and 'ruling out' understood in these specific ways, the relevant alternatives condition could not be necessary for justification after all, and could not be deployed in any argument against JSPC[13]. While no-psychological-gap cases don't give us reason to accept JSPC in full generality, they *can* be used to deflect arguments against the principle – if these arguments fail to steer clear of such cases.

The situation with JCC is very similar. In fact, the risk aggregation argument against JCC has the same basic structure as a relevant alternatives argument against JSPC, and has the same vulnerability. In the risk aggregation argument, the claim that one cannot have justification for believing a proposition that is less than 50% likely, given one's evidence, plays the same role as a relevant alternatives condition. After all, this claim also places a putative necessary condition on justification – and it is a condition that can be met by each of a series of propositions P_1, P_2, …, P_n even if it is not met by $P_1 \wedge P_2 \wedge … \wedge P_n$.

If we had a case in which one automatically counted as believing $P_1 \wedge P_2 \wedge … \wedge P_n$ upon believing P_1 and believing P_2 … and believing P_n, then it would make no sense to deny JCC in this case – it would make no sense to grant that the latter beliefs are all justified while denying that the former belief is. If one does not have justification for believing $P_1 \wedge P_2 \wedge … \wedge P_n$ then one is epistemically obliged to give this belief up. And yet, to stop believing $P_1 \wedge P_2 \wedge … \wedge P_n$ is to stop believing P_1 or P_2 … or P_n – and the claim that one is epistemically obliged to do *this* is inconsistent with the stipulation that these beliefs are all justified. If this were also a case in which the probability of $P_1 \wedge P_2 \wedge … \wedge P_n$, given one's evidence, was below 50%, then this would show that there is something wrong with the risk aggregation argument against JCC[14].

---

[13] Whether extant accounts of relevance and ruling out are subject to this particular difficulty is not something that I propose to investigate here. It is worth noting, though, that Dretske's original account of relevance may well allow for cases of this kind to arise. Dretske suggests that a relevant alternative to a proposition is an alternative that *might* have been realised if the proposition had not been true, and an irrelevant alternative is one that *would not* have been realised if the proposition had not been true (Dretske, 1970, p1021). Consider a proposition $P \wedge Q$. It seems possible that, while there are alternatives in which Q is false, none of these would be realised in the event that $P \wedge Q$ is false. Clearly, though, some of these alternatives *would* be realised in the event that *Q* is false. This opens up the possibility that one has ruled out all of the relevant alternatives to $P \wedge Q$, without having ruled out all of the relevant alternatives to Q. This would present a problem for Dretske, who explicitly denies that single premise closure can fail in cases of Conjunction Elimination (Dretske, 1970, p1009).

[14] In cases where I automatically believe $P_1 \wedge P_2 \wedge … \wedge P_n$ when I believe P_1, believe P_2 … believe P_n, should we think of this as a matter of metaphysical necessity or as a contingent psychological truth? As far as I can tell, none of the claims in the main text hinge upon the answer to this – but it is an interesting question in its own right. The comparison with cases of Conjunction Elimination may suggest the former answer – the fact that I believe that grass is green whenever I believe that grass is green and snow is white is not, I think, an observation about my *psychology* but, rather, about the *logic of belief*. The criteria for attributing a belief in a

Are there cases like this? I have attempted to argue that the party example from section I is one such. As should now be clear, though, it doesn't especially matter whether one accepts my take on this example. Provided there are *some* no-psychological-gap cases in which the falsity risk aggregates to over 50%, this will be enough to cast doubt upon the risk aggregation argument, even when applied to cases that do involve a psychological gap[15]. Anyone who wanted to avoid this objection would either have to insist that there is *always* a psychological gap between believing a series of propositions and believing their conjunction, endorsing a kind of Armstrongian picture (though actually more radical than even Armstrong contemplated) or insist that, while there are some no-psychological-gap cases, for some reason they can never involve a set of justified beliefs with an aggregated falsity risk of over 50%. Neither of these strategies strikes me as promising and, at the very least, anyone who wished to pursue them would have the burden of showing otherwise.

The reach of the risk aggregation argument is, in a way, too *wide*, too *indiscriminate* – a belief in virtually any conjunction will be at greater risk of falsity than the beliefs in the individual conjuncts. In order for the argument to engage, we have to ensure that there are enough conjuncts to force the probability of the conjunction below 50% but, provided this formal requirement is met, it doesn't matter at all what the conjuncts are about, or how one came to believe them, or whether the beliefs are held within a single compartment, or how much time elapsed between the formation of the beliefs etc. Whatever the factors that determine whether there is a psychological gap between believing a series of propositions and believing their conjunction (and I take no official stand here on what these factors are), the risk aggregation argument exhibits no sensitivity to them. The problem with the argument, in effect, is that it 'proves' too much – not only that JCC fails, but that JCC fails even in cases where it is nonnegotiable. As a result, it proves nothing.

V BELIEVING THE IMPROBABLE

Even if we accept the existence of no-psychological-gap cases, and agree that JCC could not fail in such cases, and grant that some of these cases involve an aggregated falsity risk of over 50%, it

---

conjunction subsume the criteria for attributing beliefs in the conjuncts. And yet, we can't say quite the same thing about no-psychological-gap cases of Conjunction Introduction, since these cases will always be characterised by certain contingent features. We can still say something similar though – that whenever the right features are present, the criteria for attributing a series of individual beliefs coincide with the criteria for attributing a single conjunctive belief. While I am sympathetic to this perspective on the no-psychological-gap cases, I won't attempt to defend it further here.

[15] As discussed in n3, in the standard version of the preface paradox we are asked to consider an author who writes a well-researched nonfiction book and who believes each claim within it. Even if all of these beliefs are justified it would be absurd, so the standard thinking goes, for the author to infer that every claim in the book is true. This inference would not, strictly speaking, be a simple instance of Conjunction Introduction – but it would in any case be an inference that clearly opens up a psychological gap. That is, it is clearly possible for the author to refrain from drawing such an inference, and from forming the belief that every claim in the book is true. For all that has been said here it may be perfectly *permissible* for an author to hold back from forming this belief. But why think that the author is *obliged* to hold back? Why think that it would be *wrong* for an author to affirm that every claim is true and to stand by that, unless evidence against particular claims is presented? If this stems from an instinct to criticise any belief that runs a high risk of falsity, then the present considerations show that this instinct is misplaced – something on which I will elaborate further in the final section.

might *still* seem puzzling or paradoxical to claim that one could have justification for believing something that is less than 50% probable, given one's evidence. If someone says 'It's unlikely, given my evidence, that P is true, but I continue to believe it nonetheless' doesn't that just sound *obviously irrational*? The principle that justification requires a 50%+ probability, given one's evidence, may never altogether lose its appeal[16]. In this final section, however, I will argue that this is part of a family of principles – connecting error risk with a certain liability for criticism or blame – which all share a common defect.

Consider a judge who, over the course of a long career, has convicted 100 defendants. Suppose that the judge has always been scrupulous in evaluating evidence and, throughout her career, has maintained a high evidential standard for conviction, only ever returning a guilty verdict in cases where the prosecution presented multiple independent sources of incriminating evidence which the defence was unable to counter or explain. For each one of her convictions, the chance that the defendant was really innocent, given her evidence, was very low. Nevertheless, given the sheer number of convictions that she handed down, the chance that she has convicted an innocent person *at some point* may be high. Let's attach some numbers to the example. Suppose that, for each of these 100 convictions, the probability that the defendant was guilty, given the judge's evidence, was 98%. If we suppose that these cases are independent, the probability, given the judge's evidence, that all 100 defendants were guilty is approximately *13%* – that is, there is around an 87% chance that she has convicted at least one innocent person.

Consider the following principle:

If a judge has acted in a way that is more than 50% likely, given her evidence, to result in the conviction of an innocent person, then this action is unjust.

This is an analogue of the principle that one could never be justified in believing a proposition that is more than 50% likely to be false, given one's evidence – and, at first glance, may seem just as obviously correct. And yet, although our judge is 87% likely to have convicted an innocent person, it's clear that she has not acted unjustly. By stipulation, each one of her convictions was based upon overwhelmingly strong incriminating evidence – and it makes no sense to condemn this 'set' of 100

---

[16] As noted above, the principle is widely regarded as obvious, and doubts about it are not really on the radar for most epistemologists. There are a few, however, who have questioned the principle (Kaplan, 1995, Jackson, forthcoming, Smith, 2016, section 4.3, see also Smith, 2021). Kaplan (1995) and Jackson (forthcoming, section 2.2) have suggested that it may be possible to rationally believe a proposition while rationally investing a credence of below 0.5 in that same proposition. That's pretty close to the claim that one can justifiably believe a proposition even though its probability, given one's evidence, is below 50% – and the considerations that Kaplan and Jackson adduce would seem to carry over to this claim as well.

Some have argued that the term 'belief' in ordinary discourse picks out a weak doxastic state that can be justified even if its content is relatively unlikely, given one's evidence (Hawthorne, Rothschild and Spectre, 2016, Mandelkern and Dorst, forthcoming). Those who defend this view sometimes draw a distinction between the ordinary weak notion of belief and the more demanding notion of *full* or *outright* belief that epistemologists have focussed on (Mandelkern and Dorst, forthcoming, section I) and which would still, presumably, be taken to satisfy the principle in the text. If one is attracted to this kind of picture then the right way to express my view, in the associated terminology, is that even justified *full* belief does not require an evidential probability of greater than 50%. I wouldn't sloganise this by saying that full belief is 'weak' though. Nothing that I've said here suggests that the evidential standards for justified belief are loose or lax – just that they are unrelated to evidential probabilities.

convictions while applauding each individual one. If a judge has genuinely acted in an unjust way, then we would have reason to reverse or undo the effects of this action. But we cannot undo these 100 convictions without quashing some particular convictions that the judge has handed down – and this is clearly something we *should not* do. If we overturn some of these convictions and not others then *this* would be unjust, given that the evidence in each case was equally strong, and if we overturn *all* of these convictions then this would defeat the very purpose of having a criminal justice system.   There is undoubtedly something troubling about the thought that the judge has likely convicted an innocent person – and it may even be a source of regret as she looks back upon her career – but it is not a basis for *blame* or *condemnation* of any kind.

To advise the judge to convict each one of these defendants, but hold back from convicting *all* of them is obviously nonsense. Having convicted each of the 100 defendants, there is no further decision for the judge to make. To convict each one of the defendants *just is* to convict them all – the acts, we might say, 'conjoin themselves'. This is not some peculiarity about acts of conviction. If one has brought it about that P_1 is true and brought it about that P_2 is true … and brought it about that P_n is true then one doesn't get to make some further decision on whether to bring it about that $P_1 \wedge P_2 \wedge … \wedge P_n$ is true – this has already been done. You can't criticise someone for making $P_1 \wedge P_2 \wedge … \wedge P_n$ true while praising them for making P_1 true and for making P_2 true … and for making P_n true.

Consider another principle that has a powerful initial appeal:

If a doctor acts in a way that is more than 50% likely, given his evidence, to result in a patient being misdiagnosed and prescribed the wrong treatment, then the doctor has acted negligently.

A doctor who refuses to run all of the proper tests, and who makes a diagnosis and prescribes treatment on the basis of a sheer hunch would be acting in just such a way and would, of course, be acting negligently. This is the sort of case that might first come to mind when considering this principle. But now imagine a doctor who always carefully considers a patient's symptoms and runs all practical tests before arriving at a diagnosis and prescribing treatment. If this doctor diagnoses and prescribes treatment to 100 patients then, even if his evidence makes it 98% likely that each prescription is correct, assuming independence there will be more than a 50% chance that one of these patients has ended up with the wrong diagnosis and treatment. As a result, the doctor's actions will fall within the scope of the above principle – but he has obviously done nothing negligent.

Perhaps it does sound irrational to say 'It is likely, given my evidence, that such-and-such is false, but I continue to believe it nonetheless'. But does it sound any *less* irrational to say 'It is likely, given my evidence, that my actions have led to the conviction of an innocent person, but we have no reason to revisit or reverse those actions'? Or to say 'It is likely, given my evidence, that my actions have led to a person being misdiagnosed and prescribed the wrong treatment, but we have no

reason to re-examine those actions'? And yet, in the cases I've described, these judgments – however they sound – must be accepted[17].

Discussions of conjunction closure have been premised on the assumption that we can simply decline to 'conjoin' the propositions that we believe. Because of this, the question of whether a principle like JCC holds has been framed in a very particular way; as a question about whether, and when, we should exercise this choice. Risk aggregation, correspondingly, has been portrayed as something that we can always avoid or control (so why not choose to do so?) If the arguments presented here are right, then these issues have been fundamentally mischaracterised. In some cases, we are no more free to believe each of a series of propositions, while declining to believe their conjunction, than we are free to bring about each of a series of outcomes, while declining to bring about *their* conjunction. Risk aggregation is an inevitable part of acting in the world – and an inevitable part, also, of forming beliefs about it.

Acknowledgements

---

[17] One might argue that the judge/doctor examples are not truly analogous to a situation in which one believes a conjunction that is likely to be false. After all, in these examples there is no *identifiable person* who is likely, given the evidence, to have been wrongfully convicted/misdiagnosed – it's merely likely that some person or other has suffered this fate. In something like the party example, however, there *is* an identifiable belief – the belief that Peter and Summer and Shelley and… will attend – that is likely to be false. Isn't that a crucial difference?

Here's how I see the analogy holding up: The judge/doctor's actions are likely, given their evidence, to have led to a wrongful conviction/misdiagnosis. Similarly, my conjunctive belief is likely, given my evidence, to be false. The judge/doctor's evidence doesn't, however, support any particular narrative as to *how* their actions could have led to this consequence – every conviction/diagnosis was backed by strong incriminating evidence/medical tests. Rather, the likelihood of this consequence comes from the sheer number of different ways in which it could have occurred – the first person could have been wrongfully convicted/misdiagnosed, the second person could have been wrongfully convicted/misdiagnosed etc. Similarly, my evidence in the party example doesn't support any particular narrative as to how my conjunctive belief could be false – everyone gave a positive reply. Rather, the likelihood of falsity comes, once again, from the number of different ways in which this could happen – Peter could fail to show up, Shannon could fail to show up etc. The analogy, then, is between my conjunctive belief and the judge/doctor's 'conjunctive' action. And the fact that there is no identifiable person who is likely to have been wrongfully convicted/misdiagnosed corresponds with the fact that there is no identifiable conjunct that is likely to be false.

References

Alston, W. (1985) 'Concepts of epistemic justification' *The Monist* v68(1), pp57-89

Alston, W. (1988) 'The deontological conception of epistemic justification' *Philosophical Perspectives* v2, pp257-299

Armstrong, D. (1973) *Belief, Truth and Knowledge* (Cambridge: Cambridge University Press)

Beddor, B. (2020) 'New work for certainty' *Philosophers' Imprint* v20(8), pp1-25

Blome-Tillmann, M. (2015) 'Solving the Moorean puzzle' *Philosophical Studies* v172(2), pp493-514

Christensen, D. (2004) *Putting Logic in its Place: Formal Constraints on Rational Belief* (Oxford: Oxford University Press)

Coliva, A. (2012) 'Moore's proof, liberals and conservatives: Is there a (Wittgensteinian) third way?' in Coliva, A. ed. *Mind, Meaning and Knowledge: Themes From the Philosophy of Crispin Wright* (Oxford: Oxford University Press)

Douven, I. (2002) 'A new solution to the paradoxes of rational acceptability' *British Journal for the Philosophy of Science* v53(3), pp391-410

Dretske, F. (1970) 'Epistemic operators' *Journal of Philosophy* v67(24), pp1007-1023

Ebert, P., Smith, M. and Durbach, I. (2020) 'Varieties of risk' *Philosophy and Phenomenological Research* v101(2), pp432-455

Evnine, S. (1999) 'Believing conjunctions' *Synthese* v118(2), pp201-227

Foley, R. (1993) *Working Without a Net: A Study of Egocentric Epistemology* (Oxford: Oxford University Press)

Foley, R. (2009) 'Beliefs, degrees of belief and the Lockean thesis' in Huber, F. and Schmidt-Petri, C. eds. *Degrees of Belief* (Dordrecht: Springer)

Gardiner, G. (2021) 'Relevance and risk: How the relevant alternatives framework models the epistemology of risk' *Synthese* v199(1-2), pp481-511

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

Hawthorne, J., Rothschild, D. and Spectre, L. (2016) 'Belief is weak' *Philosophical Studies* v173(5), pp1393-1404

Hintikka, J. (1962) *Knowledge and Belief: An Introduction to the Logic of the Two Notions* (Ithaca, NY: Cornell University Press)

Jackson, E. (forthcoming) 'On the independence of belief and credence' *Philosophical Issues*

Kaplan, M. (1995) 'Believing the improbable' *Philosophical Studies* v77(1), pp117-146

Kaplan, M. (1996) *Decision Theory as Philosophy* (New York: Cambridge University Press)

Klein, P. (1995) 'Scepticism and closure: Why the evil genius argument fails' *Philosophical Topics* v23, pp213-236

Kraus, S., Lehmann, D. and Magidor, M. (1990) 'Nonmonotonic reasoning, preferential models and cumulative logics' *Artificial Intelligence* v44(1-2), pp167-207

Kyburg, H. (1961) *Probability and the Logic of Rational Belief* (Middletown: Wesleyan University Press)

Lasonen-Aarnio, M. (2008) 'Single premise deduction and risk' *Philosophical Studies* v141(2), pp157-173

Leuenberger, S. and Smith, M. (2021) 'Epistemic logic without closure' *Synthese* v198(5), pp4751-4774

Lewis, D. (1982) 'Logic for equivocators' *Noûs* v16(3), pp431-441

Makinson, D. (1965) 'The paradox of the preface' *Analysis* v25(6), pp25-27

Mandelkern, M. and Dorst, K. (forthcoming) 'Assertion is weak' *Philosophers' Imprint*

Nelkin, D. (2000) 'The lottery paradox, knowledge and rationality' *Philosophical Review* v109(3), pp379-409

Pedersen, N. (2009) 'Entitlement, value and rationality' *Synthese* v171(3), pp443–57

Routley, R. and Routley, V. (1975) 'The role of inconsistent and incomplete theories in the logic of belief' *Communication and Cognition* v8(2/4), pp185-235

Ryan, S. (1996) 'The epistemic virtues of consistency' *Synthese* v109(2), pp121-141

Schotch, P. and Jennings, R. (1980) 'Inference and necessity' *Journal of Philosophical Logic* v9, pp327-340

Smith, M. (2013) 'Two notions of epistemic risk' *Erkenntnis* v78(5), pp1069-1079

Smith, M. (2016) *Between Probability and Certainty: What Justifies Belief* (Oxford: Oxford University Press)

Smith, M. (2018) 'The logic of epistemic justification' *Synthese* v195(9), pp3857-3875

Smith, M. (2021) 'Civil liability and the 50%+ standard of proof' *International Journal of Evidence and Proof* v25(3), pp183-199

Smith, M. (forthcoming) 'The hardest paradox for closure' *Erkenntnis*

Spectre, L. (2019) 'Compartmentalized knowledge' *Philosophical Studies* v176(10), pp2785-2805

Spohn, W. (2012) *The Laws of Belief: Ranking Theory and its Philosophical Applications* (Oxford: Oxford University Press)

Stalnaker, R. (1984) *Inquiry* (Cambridge, MA: MIT Press)

Tang, W. (2018) 'In defense of single premise closure' *Philosophical Studies* v175(8), pp1887-1900

Thorn, P. (2017) 'Against deductive closure' *Theoria* v83(2), pp103-119

Weiss, Y. (2019) 'Are contradictions believable?' *Thought* v8(1), pp42-49

Williamson, T. (2000) *Knowledge and its Limits* (Oxford: Oxford University Press)

Wright, C. (2004) 'Warrant for nothing (and foundations for free)?' *Aristotelian Society Supplementary Volume* v78(1), pp167–212.

Wright, C. (2014) 'On epistemic entitlement (II): Welfare state epistemology' in Dodd, D. and Zardini, E. eds. *Scepticism and Perceptual Justification* (Oxford: Oxford University Press