

# EPISTEMIC SAFETY: RISK AND ITS NEGATIVITY IN THE GOOD CASE

Wai Lok CHEUNG

**ABSTRACT:** Practical encroachment may be understood with an influx of epistemic possibilities into the epistemic context through raising epistemic standard. The same piece of evidence no longer epistemically justifies the corresponding belief because of the extra alternative possibilities. I differentiate relevance into ethical relevance and epistemic relevance through practical interest, and note that, even though one lost confidence in some facts one used to know, one did not lose epistemic justification. I will illustrate the difference between making an epistemic justificatory difference to knowledge and its required precision, and making only an ethical justificatory difference about what one decides to do. If banks are not open on Saturday mornings, one would have to decide differently concerning when it is that one deposited one's pay cheque, but the salience of the alternative does not constitute its epistemic relevance: one retained the epistemic justification to know that the bank will be open tomorrow all the same. One's belief is thus still epistemically safe as before. Epistemic risk measures the likelihood of having committed an epistemic wrong, which is constituted by believing falsely. I propose a numerical relation between epistemic safety and epistemic risk as a result.

**KEYWORDS:** epistemological invariantism, epistemic safety, practical encroachment, risk, epistemic risk, margin of error

Epistemic safety obligates belief only with facts, and it entails that a belief constitutes knowledge if and only if it is with an epistemic immunity from error. Often in the scholarly literature, such immunity is described using proximity of worlds – a Lewisian notion – such that a belief is epistemically safe if and only if it is true of nearby worlds. The belief, therefore, is epistemically immune from error among them. In this paper, I first discuss how practical interests, using Jason Stanley's cases, make a difference to what nearby worlds are. It seems, with much more at stake, that the numerosity of the amount of nearby worlds increases, thus setting a higher standard of epistemic safety for a belief of the same fact for the constitution of knowledge. In the section that follows, I will illustrate what epistemic risk is through discussing, instead of nearby worlds of which the belief is true, those of which it is false, using Timothy Williamson's margin of error. The epistemic immunity from error a belief is supposed to have is qualified with a corresponding margin, such that, within its width, its immunity is understood with a given deviation from actuality. In the next section, how the risk of being wrong, epistemically, interacts with its

cost – the practical cost – is demonstrated using John Hawthorne’s salience explanation of subject sensitivity in knowledge attribution. Through making an inference from the falsity of the conclusion of an argument to the falsity of the premise, the agent’s anxiety caused by the worry about the fictitiousness of the fact that was once known brings to mind what ought not have been considered because epistemically irrelevant. In the last section, I will argue for the illusion of danger, through its epistemic irrelevance, in some cases in which interest destroyed knowledge, to put forward some epistemic justificatory difference that some alternatives made is, in fact, of negative relevance. The upshot is the differentiation of relevance with epistemic relevance, and the corresponding distinction between epistemic proximity and ethical proximity of worlds in their measurement of distality. A numerical relation between epistemic safety and epistemic risk will also be proposed.

### 1. Interest

The bank case has brought into play what is now known as practical encroachment, or pragmatic encroachment. Consider the following two bank cases modified with a more concrete consequence of failure. The decision to modify thus is justified by the need to make more vivid the scenario in which practicality allegedly makes a difference through psychology to knowledge.

#### *The low-stake case*

It is winter and they need the electricity to keep the heating, for otherwise they would have caught a serious cold at night and be very ill. It is Friday, and Hannah and her wife drove by the bank, observing a long queue. Since they have sufficient fund in their bank account for paying the electricity bill, they decided to leave till the morning after to deposit their recently received pay cheque. Although, on Saturdays, banks open not as regularly as during weekdays, it is open the day after in the morning.

#### *The high-stake case*

It is winter and they need the electricity to keep the heating, for otherwise they would have caught a serious cold at night and be very ill. It is Friday, and Hannah and her wife drove by the bank, observing a long queue. Since they do not have sufficient fund in their bank account for paying the electricity bill, they decided to queue up to deposit their recently received pay cheque. They could have waited till the morning after, when the bank is less busy, but they worry that the less regular hours on Saturdays would have ruined their plan, leaving them ill from the cold weather. Even so, the bank is open the day after in the morning.

Stanley states that, in the low-stake case, the protagonists know that the bank will be open the next day, whereas in the high-stake case, they do not. Instead of evaluating the truth of any corresponding knowledge ascription, I will illustrate the epistemic difference between the two cases using informativeness relativity. If they were to call someone on the phone to inquire because of their doubt, the difference in informativeness of the same utterance can be illustrated thus. Before they decided, in either case, to deposit the cheque on their respective day, they called Bill on the phone to ask whether the bank would be open the next day. Bill answers, 'I was there the Saturday morning before last week, and it was open. Banks open on alternate Saturday mornings, and yes, the bank will be open tomorrow'. Bill's utterance constituted an assertion that transferred knowledge to the protagonists in the low-stake case, and they are thus informed of something that they presumably doubted – that the bank will be open the next day. Therefore, in the low-stake case, Bill's utterance is informative with regard to whether the bank will be open the next day. However, Bill's utterance failed to transfer his knowledge to them in the high-stake case. Furthermore, as Stanley notes, they would even deny Bill knowledge that the bank opens the next day. (Stanley 2005, 5) Therefore, in the high-stake case, the same utterance is not informative with regard to whether the bank will be open the next day.

The difference in informativeness may be explained by a difference in the numerosity of the amount of epistemic possibilities in their respective epistemic context. Bill's utterance is informative for the protagonists in the low-stake case because there are fewer epistemic possibilities in their epistemic context, and his utterance is sufficient to rule out all epistemic possibilities in which the bank is not open the next day. His testimony, therefore, encodes sufficient information for Hannah and her wife to restore their knowledge after their putative moment of doubt. However, it is not informative for them in the high-stake case because there are more epistemic possibilities in their epistemic context. Bill's utterance is not sufficient to rule out all such alternatives in which the bank is not open the next day. Whatever information Bill's testimony encodes, it is insufficient to restore their knowledge given that they doubted that the bank will be open the next day. In the former case, given an epistemic infallibilism, they come to know, after having doubted, that the bank will be open the next day, because all relevant alternatives are ruled out. In the latter, they do not because not all relevant alternatives are ruled out. There are still some, in the epistemic context, epistemic possibilities in which the bank will not be open the next day. The important issue here is with the relevance of those alternatives, and how interest determines relevance. I will return to it by differentiating relevance into epistemic relevance and otherwise later on.

Before that, I will focus on my disagreement with Stanley with regard to the high-stake case.

Stanley notes that the difference of the high-stake case from the low-stake case may be explained with psychological difference, instead of epistemic difference.

‘[O]ne might attempt to explain away the force of the intuitions behind these scenarios, by arguing that, when someone recognizes that the costs of being wrong are particularly high, his or her confidence is shaken. The result of having one’s confidence shaken is either to reduce one’s degree of belief below the threshold required for knowledge or to defeat the evidence one has for one’s belief in some other manner’. (Stanley 2005, 6)

To respond to this, Stanley invites us to consider a case in which the confidence of the protagonists did not decrease. Given that, in the low-stake case, the protagonists knew that the bank would be open the next day, the problem is whether the same persons in a high-stake situation, without knowing it, would also know that the bank will be open the next day. Suppose, unbeknownst to Hannah and her wife of the impending electricity bill, they decided, without hesitation, to deposit the pay cheque the next day, having observed the long queue on that Friday. They, eventually, made the deposit on Saturday, without any disturbance to their daily routine, the next month. During their discussion, Hannah says to her wife, “I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning”.<sup>1</sup> (Stanley 2005:5) On that, Stanley comments,

‘[O]ur reaction is that Hannah’s utterance of “I know the bank will be open tomorrow” is false... After all, Hannah is more knowledgeable about her situation in the [high-stake case]... It does not seem correct that adding a little ignorance increases knowledge’. (Stanley 2005, 5-7)

Since they are in the same practical situation, Stanley insists that it is a wrong epistemology if it entailed that Hannah’s ignorance of their high-stake situation increases their knowledge through knowing something more than otherwise. The protagonists know that the bank will be open the next day because, in contrast to the high-stake case, they are ignorant of their practical situation. Contrastively, the protagonists in the high-stake case, in virtue of knowing their practical situation, failed to know that the bank will be open the next day. Knowledge ought not have destroyed knowledge. If knowing the stake would have contrasted with knowing less – losing the knowledge that the bank will be open the next day – in comparison

---

<sup>1</sup> Stanley calls this the ignorant high-stake case.

with the present contending case, then we are faced with a case of knowledge-destroying knowledge.

I approach the problem with the difference in restoration of knowledge through Bill's utterance among the low-stake and the high-stake cases. I accept the datum of the difference in informativeness, but without a difference in epistemic justification. The protagonists have the same epistemic justification in both cases for their knowledge that the bank will be open the next day, but their difference in practical situation differentiated them with a difference in relevant alternatives. The extra possibility brought in is the one in which they decided to deposit the pay cheque the next day and they got seriously ill – because of having no electricity to sustain the heating of their flat. Since this possibility can be eliminated if they had sufficient fund in their bank account, the alternative makes a justificatory difference with regard to when it is that they deposited the pay cheque – a decision whose deliberation factors in the fact of the bank's being open the next day. The alternative to actuality that the bank will not be open the next day is thus, allegedly, relevant, but not *epistemically relevant*.<sup>2</sup> Although the alternative made a justificatory difference to whether to decide to deposit the pay cheque the next day, it made no epistemic justificatory difference to whether one knew that the bank is open the next day because of the invariance of the epistemic justification across the two cases. Their evidence remains the same, and evidence of the high-stake situation does not constitute an epistemic defeater of their prior evidence of the bank's opening hours.

One way in which the alternative made a justificatory difference epistemically is through, for example, knowing better about the bank's opening hours during a week for their decision next time around. The alleged epistemic relevance of the alternative, if at all, is thus explained with a sameness in epistemic justification for their belief that the bank will be open the next day among the two cases; how is an alternative supposed to make an epistemic justificatory difference among them if not through increasing it in one case? This relates to the reasonableness of the doubt that the bank will not be open the next day, which will be discussed in the last section as a case of negative relevance. Let me now turn to the risk of epistemic wrong, through epistemic safety, having set the stage with epistemic context.

---

<sup>2</sup> In the end, it was on supposition of the epistemic relevance of the alternative that it made an ethical justificatory difference. More on this later.

## 2. Error

Since, on an epistemic safety conception of knowledge, a piece of evidence ruled out all relevant epistemic possibilities only if the resultant belief is true of all nearby worlds, what belonged to the epistemic context, especially when interest makes a difference to relevance without it being an epistemic difference, can be considered through distinguishing proximity of worlds via epistemology versus via ethics. I shall now turn to Williamson's notion of margin of error to motivate what epistemic risk is, in contrast to risk that is not epistemic but ethical, having to do with practicality.

Suppose Hannah and her wife believed truly, of a typical week, that bank employees work for 5.5 days. Banks are open on Monday, Tuesday, Wednesday, Thursday, Friday, and Saturday morning, because that is the protagonists' own working schedule.<sup>3</sup> However, they believed it with a certain margin of error: that it is with a margin of 0.5 day. It might as well be that bank employees do not work on Saturday, or if they do, they do so including Saturday afternoon. The number of workdays in a week for bank employees relative to the protagonists ranges from 5 to 6, through their believing it to be 5.5 days, with the aforementioned margin of 0.5 day of error.<sup>4</sup>

Suppose it is indeed the case that banks open on the 5 weekdays plus Saturday morning, constituting a 5.5 days of work per week. The margin of error principle states that

- (1) 'Banks open for 5.5 days per week' is true of all cases similar to cases of which  
'It is known that banks open for 5.5 days per week' is true. (Williamson 1994, 227)

Similarity among cases determines epistemic proximity of worlds. If the protagonists' belief in banks' 5.5 days of work per week constituted knowledge, then the notion of similarity here at work may be interpreted in such a way that the margin of error is 0.5 day; their evidence rules out worlds in which it is less than 5 days and more than 6 days because, with such evidence, only those worlds are distinguishable from the actual world, and it epistemically justifies their belief with the aforementioned margin of error. It was insufficient evidence for fully epistemically justifying the belief, leaving a positive width of epistemic margin, which indicates a less than maximum credence, but it still epistemically justifies the corresponding belief nonetheless.<sup>5</sup> Their belief has an epistemic immunity from

---

<sup>3</sup> This modifies the fact that banks open on alternate Saturday mornings for simplicity.

<sup>4</sup> Supposing both the upper and lower bounds to be inclusive.

<sup>5</sup> Note however, alternatively, that it is metaphysically possible to believe outright knowing the margin of error. One would have thus formed a belief with certainty whose truth is within a given

error among worlds in which banks open on Monday, Tuesday, Wednesday, Thursday, and Friday only, and also those in which banks open on Monday, Tuesday, Wednesday, Thursday, Friday, and Saturday. This is an epistemological datum, had an epistemic error of 0.5 days been permitted of their epistemic decision to believe that banks open for 5.5 days per week.. If they had confidence, or degree of belief, greater than the credence, then they would have believed more than are epistemically justified by their evidence. Confidence is psychological, whereas credence is epistemic.

However, I will illustrate with a case in which the epistemic width makes much more sense than when the piece of knowledge is specified with the numericity of 5.5. If banks are open not on Saturday, such that it is open only for 5 days per week, then 'Banks open for 5.5 days per week' is not true of some cases similar to the case of which 'It is known that banks open for 5.5 days per week' is true – on a more commonsensical understanding of what it is for Hannah and her wife to know that banks open for 5.5 days per week. We need something with which their belief, with such an epistemic margin, could have constituted a piece of knowledge. Consider the following.

(2) Banks provide sufficient financial services.

In evaluating the truth of (2), Hannah and her wife check banks' working days against the working days of other businesses, and conclude that, whether they open on 5.5 days, as they believe it, or their amount of working days per week lies on the lower margin with only 5 days, (2) is still true. Therefore, given their evidence from their own number of work days per week, (2) is true in all cases similar to the one in which they know (2). Here, their knowledge of (2) is understood with their belief's epistemic immunity from error among nearby worlds, even if their evidence is insufficient to have ruled out worlds in which banks open only 5 days a week.

Now contrast someone to whom it matters much if banks are not open on Saturdays. Suppose Jill works 5 days a week, but her job does not permit her to take any leave. She occasionally needs to use the bank for financial services, such as depositing her pay cheque, and it thus matters much to her whether banks open only 5 days a week, lying on the lower margin of Hannah and her wife's belief, or indeed 5.5, if not even 6 days, a week. If Jill were to evaluate the truth of (2), she might have a different epistemic standard, such that she would not affirm (2) except with a belief in the banks' opening for 5.5 days a week without such a width for epistemic error, but much less. This is how, I think, practicality determines a different epistemic standard for an allegedly identical piece of knowledge.

---

range with some probability distribution, instead of with maximum precision.

Stanley notes that, with regard to the difference in knowledge attribution in different practical situations, one may ‘deny that the proposition expressed by [the] utterance in [the low-stake case] is really the denial of the proposition expressed by [the] utterance in [the high-stake case]’. (Stanley 2005, 12) Hannah, with regard to (2), is in a low-stake situation, whereas Jill a high-stake situation. If, without having checked further the working days of the bank, Jill, doubting that banks are not open on Saturdays, deny (2), one may understand the epistemic difference between Hannah the low-stake case and Jill the high-stake case, given the same evidence and thus the same epistemic justification, with a difference in the proposition being entertained.<sup>6</sup>

With regard to (2), different epistemic agents in different practical situations thus have a different margin of error for a related fact – a fact that made an epistemic justificatory difference to one’s affirmation of (2). The notion of similarity Williamson uses brings out the question of what it is for a world to be nearby; when Jill knows (2), are worlds in which banks open only 5 days a week nearby worlds? Presumably, those worlds had once been in the epistemic context, but had been ruled out by Jill’s further, subsequent, evidence; once she checked with the banks that they are open also on Saturday mornings. Her belief, therefore, is epistemically immune from error within her epistemic context – of all worlds epistemically possible of the actual world, her belief that banks open 5.5 days a week is true. This explains her knowledge of banks’ workday, through her knowledge, via her judgment, of (2). Given a margin of error, her belief, within the epistemic context, is epistemically immune from error within the width. Her practical situation determines the ethical relevance of the alternative to the actual world in which banks open only 5 days a week, and thereby raised the epistemic standard for knowing (2), and correspondingly, for knowing banks’ workday to be 5.5 days a week.

There is a difference between people with different standard, and a same person, holding standard constant, in different situation. Suppose Hannah and her

---

<sup>6</sup> This invites the epistemic content understanding. Notably, to borrow from the previous section, the epistemic content of (2) relative to someone in the high-stake situation is understood with the epistemic dynamic of having ruled out worlds in which she got seriously ill from a cold because of her failure to pay the electricity bill, whereas, relative to someone in the low-stake situation, the epistemic content does not include such worlds. Someone who knows that Hesperus is Phosphorus needs no further evidence other than Hesperus is Venus to know that Phosphorus is Venus, and thus the epistemic content of ‘Phosphorus is Venus’ is different relative to someone ignorant of the identity between Hesperus and Phosphorus in contrast to someone who knows the identity. See Cheung (2017) for a theory of epistemic content in response to the epistemological problem of necessary *a posteriori* truths.



wife had the same personal standard as Jill. When is it that a nearby world no longer so, and when is it what was not a nearby world became one? Notice the difference between epistemic proximity through similarity, and ethical proximity. The high-stake case tempts us to take worlds in which one decided to deposit the pay cheque the next day and got seriously ill as a nearby world, and I shall now turn to why we ought not have done so. Are worlds in which one got seriously ill very similar to one in which one stayed healthy? Clearly not. However, shall there be a slight modification of the fact – that the bank is not open the next day – carrying through with the plan to deposit the pay cheque the next day constitutes an ethical proximity of a world in which one got seriously ill because of the high potentiality of catching a cold through not having the electricity to sustain the heating in one's flat during winter. If the bank were not to open the next day, one could have *very easily* caught a cold, and it is extremely bad that one got so seriously ill. Epistemic proximity of a fact with regard to a belief based on a piece of evidence is measured with the proportion of possibilities in which the fact obtained among all possibilities given the decision to believe in that fact with the given evidence. Ethical proximity of an outcome with regard to an act in a situation is measured with the proportion of possibilities in which the outcome obtained among all possibilities given the decision to act towards that outcome in the given situation. Epistemic proximity is understood with similarity because of an epistemic accessibility of indistinguishability. (Cheung, forthcoming)

Although such a practical situation makes an epistemic justificatory difference to (2), it ought not have done so with regard to the belief in banks' 5.5 workdays, and thus the bank being open the next day, at least in the morning. Even if it made a justificatory difference to whether their present situation suffices for their sustenance, and whether any change was needed, analogous to whether they are to carry through with their plan to wait out the long queue on that Friday by depositing the next day, it is not an epistemic justificatory difference that it made to their belief in the bank's being open the next day. Deliberation about what to do in the actual world considers what it might have been given one's act, and since the alternative of serious illness does not belong to the epistemic context, it also does not belong to the deontic context.<sup>7</sup>

---

<sup>7</sup> The probability of outcome given act, which measures the ethical proximity of the outcome, belongs to the epistemic state of the deontic agent through being encoded in the probability function. The utility function, on the other hand, encodes weak preference among outcomes. Decision to act operates with what one knows, and thus deontic context is restricted by epistemic context.

### 3. Salience

Besides identifying, through epistemic error, over which fact practicality determines epistemic status, a way to explain the seeming destruction of knowledge with more knowledge is through salience. In discussing the lottery paradox, Hawthorne proposes an explanation using the following psychological phenomenon. Suppose that, contrary to Stanley, Hannah and her wife, even in the high-stake case, knew that the bank would be open the next day. They drove past the long queue and formed the intention to deposit the pay cheque the next day. However, once apprised of the fact that they have an impending electricity bill coming due, without sufficient funds in their bank account, they are overcome with anxiety. They wanted another uneventful month, but the alternative in which they got seriously ill from a cold they caught during the few days without the electricity to sustain the heating in their flat became salient. Salience, however, does not constitute relevance. (Hawthorne 2004, 160-2)

Anxiety differs from worry, and worry differs from doubt. Sometimes, one could have been anxious without having anything specific in mind, and, with such anxiety, pays too much attention to what was not, in fact, a danger and starts to worry. One could have also worried about something without doubting it, knowing such worry to be unfounded, and thus irrational, if not also unreasonable. One tries to contain one's such psychology and restrict the episode to as short a duration as possible. Lastly, one might start to doubt what began only as a worry.

Therefore, there is a difference between doubting what one ought to have doubted and doubting what one ought not to have doubted. Consider

- (3) For any  $p$ , one ought to have doubted that  $p$  only if  $p$  is an epistemically relevant alternative.

Although getting seriously ill makes a great justificatory difference to one's daily routine, such an alternative to actuality is relevant not epistemically, and thus one ought not have doubted that one would get seriously ill. If doubting the bank to open the next day psychologically necessitated the former doubt, we are also to evaluate whether one is to doubt it as well. Is the bank not opening the next day an epistemically relevant alternative?

If, as any piece of ordinary knowledge, banks' 5.5 days of work per week, or at least that banks open on alternate Saturday mornings, is a piece of knowledge that any citizen is to have, then, if any epistemic difference between the low-stake case and the high-stake case, it shall have to be established independently from a prior commitment to the epistemic effect of practicality. The lottery paradox interacts cognition with epistemology through something much more mundane than getting

seriously ill from not having sufficient funds in one's bank account. Now consider Hawthorne's construal. Ordinary citizens like you and I shall not have sufficient funds to travel around the world thrice in one go. However, like most ordinary citizens, we purchase lottery tickets from time to time, and if I won, I would have sufficient funds. Therefore, it seems that knowledge of

- (4) I will not have sufficient funds to travel around the world thrice in one go in the coming month.

is incompatible with the ignorance of

- (5) I will not win with the lottery ticket I recently purchased.

since (5), given auxiliary pieces of knowledge, is logically entailed by (4).

To assimilate Hawthorne's salience theory to the bank cases, I will now reformulate the high-stake case analogously. It seems, from Stanley's assessment, that knowledge of

- (6) I will not get seriously ill from catching a cold through not having the electricity to sustain the heating in my flat in the coming week.

is incompatible with the ignorance of

- (7) The bank will be open the next day.

since (7), given knowledge of the supplementary act one decided to perform, is logically entailed by (6).

If (7) is false, then (6) is false, and it is indeed anxiety-provoking. One might get seriously bedridden, suffering from various body aches, let alone the potential of the cold being fatal. One would also have to give up some recent work projects that would have been a pity that one invested so much in. The anxiety drives one to worry about various things, and if it has developed into an actual doubt of (6) – that one considers the alternative in which one got seriously ill thus, it does not, Hawthorne insists, thereby constitute the epistemic relevance of the serious illness. One ought not have doubted because one knew that the bank would be open the next day, and thus depositing the pay cheque on the Saturday ought not have been paired, epistemically, with getting seriously ill from catching a cold, thus through not having the electricity to sustain the heating in one's flat.

Equivalently, such an alternative should not have been in the epistemic context. Worlds including such facts are not nearby worlds, and one's evidence, prior to the doubt, was sufficient to have those worlds ruled out from the epistemic context. Worlds in which one failed to deposit the pay cheque on time and got seriously ill from having no electricity for heating, given the decision to deposit the

next day, are not live options at all because one had sufficient evidence that the bank would be open the next day.

The issue now is whether, with a change in salience, the margin of error changed. Presumably, since epistemically irrelevant, worlds in which the bank is not open the next day are not cases similar to the case in which the protagonists know that the bank is open the next day, and thus, in all similar cases, the bank is open the next day; they are not supposed to make an epistemic justificatory difference to whether the protagonists know because one's prior evidence is sufficient for such alternative to have been ruled out from the epistemic context. Psychologically, one cannot help but think of the serious illness one would have caught, had the bank not opened, and it seems to raise an epistemological point about how this worry ought not have constituted a doubt. Only with a change in epistemic standard is the margin of error narrowed, leaving the corresponding alternatives not ruled out by prior evidence, and the question becomes, thus, whether the possibility, albeit illusory, of getting seriously ill changed it. If it is only an ethical standard that is changed, it leaves the notion of epistemic safety with the illustration of ease of being wrong being the suspect of the origin of the present confusion of the ethical with the epistemic.<sup>8</sup> A slight modification of facts – that the bank is not open the next day – results in the serious illness, constituting its ethical proximity given one's decision to deposit the next day.<sup>9</sup> Even if one could have easily gotten seriously ill with such a decision, it does not constitute the similarity of such an alternative to actuality because one's prior evidence suffices for its distinguishability from actuality.

Given our understanding of (2), with the narrowed epistemic margin through what was brought in via salience, there might as well have been nearby worlds that are in the epistemic context, not being ruled out by one's prior evidence, that now constitute the failure of the belief's epistemic immunity from error. What was not an epistemic error – that banks open not on Saturdays while still having constituted the sufficiency of their financial service – is now an epistemic error with regard to one's belief of (2), but this shall not have modified the epistemic context. The epistemic possibility of banks' being open on Saturday remained the same, but the epistemic possibility of banks' insufficient financial service is modified because of one's having been updated about the practical situation. Even if the alternative of serious illness made a justificatory difference to one's decision to deposit the pay cheque the next day – that it is ethically relevant –, it did not make an epistemic

---

<sup>8</sup> See, for example, Pritchard (2005), for the formulation of epistemic safety using ease of being wrong.

<sup>9</sup> Again, this is on supposition of the corresponding epistemic relevance.

justificatory difference to one's belief that the bank will be open on the Saturday morning. The alternative of serious illness is epistemically irrelevant. In no epistemically proximal worlds is one getting seriously ill. It was ethically relevant only on supposition that it were epistemically relevant, and this illusion of danger, brought in via one's anxiety, motivates, wrongly, practical encroachment.

One should note that although evidence epistemically justifies one's knowledge of (6), given the nontransitivity of epistemic accessibility, with which knowledge does not entail knowledge of knowledge, one might not yet know that one knows (6).<sup>10</sup> Consider first perceptual knowledge. Even when one knows, through perception, that there is a hand, one may still doubt that one knows that there is a hand. This is when one brings into salience the epistemic possibilities of one's perceptual belief being false, seeking, with such doubt, to know that one knows via ruling out such epistemic possibilities. Similarly, without knowing that one knows, one doubts that one knows (6), and such doubt is driven by what motivated its salience, especially the badness of the serious illness. Here, the importance of anxiety is demonstrated: had it only been the lottery case in which, through epistemically competent deduction, one fancies the possibility in which one has sufficient fund to travel around the world thrice, it might not have raised doubt as the alternative might not have been thus similarly salient through provoking anxiety. If one shall have presented oneself with some sufficient evidence to fully justify one's knowledge of (7) epistemically, through knowing what one's evidence for one's knowledge of (7) is, one may know that one knows (7), and via epistemically competent deduction in another logical direction, know that one knows (6). One thus eliminated the doubt, via epistemically eliminating the epistemic counterpart identical to oneself who does not know (6) – an epistemic possibility of oneself – and came to know that one knew (6).

#### 4. Risk

Our inquiry is sometimes guided by our practical interests. However, it is also guided by theoretical interests. With regard to a given topic, one formulates questions, with attempts to answer them. Sometimes, one questions what has already been epistemically established – it is sometimes good practice to call a well-established fact into question. In a community, it might be common practice that what is with much at stake is doubted, for a better assessment of risk. However, as noted in previous sections, I sharply distinguish between epistemic risk and otherwise. I shall

---

<sup>10</sup> See Cheung (forthcoming) for an epistemic accessibility that is reflexive, nontransitive, and nonsymmetric.

now turn to my conception of epistemic risk, in relation to doubt, especially with regard to its permissibility on reasonableness, before discussing briefly the interaction between the epistemic risk of an epistemic act with its practical cost. I will end with a proposal of the numerical relation between epistemic safety and epistemic risk.

When something seems relevant, we may, reasonably, call that into question by bringing it into the epistemic context if it is not already in it. It might have been a fact that one has called into question, such that, if any alternative, it is not immediately epistemically relevant because not actual. However, before we concluded the inquiry with an epistemic decision in any way, what seems relevant are either positively relevant or negatively relevant – it made some justificatory difference epistemically, but we do not yet know whether it epistemically justified a belief in it, or its opposite. Doubting a fact, especially an epistemically well-established fact, calls one's knowledge into question; even if one knew, one wants to know, through such inquiry via doubt, whether one knew. In establishing epistemically that the evidence that fully justified one's knowledge before is indeed such evidence, one is better justified to believe that one is fully justified to know the corresponding facts. The corresponding doubt, therefore, even of some facts, made some epistemic justificatory difference through justifying more, higher-order, knowledge.

In evaluating a piece of evidence, one asks what ought to have stayed in the epistemic context. It used to rule out a given set of epistemic possibilities, but ought the set have been smaller? Ought less epistemic possibilities be ruled out? Was it with a correct epistemic standard, with a correspondingly correct epistemic margin for epistemic error? What are the cases similar to the case in which one knew in which what is known ought to still have been a fact concerning the this epistemic margin? This question over the margin of error interacts with the following question: if, given my evidence, I restored my belief, what is the risk of having believed wrongly?

An epistemic risk is understood thus with epistemic wrongness. An instance of an epistemic wrong is believing in something fictitious. It is not a fact that Christmas celebrates the birthday of Santa Claus, and, thus, one commits an epistemic wrong in believing in it. With the evidence that one has, one is to assess the risk of having committed an epistemic wrong; how likely is it that one would have believed something false? What belonged to the epistemic context thus makes an epistemic justificatory difference because, had some alternatives not been in it, one could have believed in a fact without the belief having constituted any epistemic wrong. It would have been the case when one is in an epistemic situation in which

one is perfectly epistemically safe with regard to the fact. The possibilities that belonged to the epistemic context would have thus been only negatively relevant, should such alternatives have turned out to be actual.

The difference between bringing in alternatives into the epistemic context and keeping it constant while modifying the margin of error is crucial. Holding prior evidence constant, if one were evaluating a piece of incoming evidence, what epistemic possibilities it rules out depends on the margin of error. On the other hand, shall the margin of error of some prior knowledge be modified, what had been ruled out by prior evidence would be brought into the epistemic context, and the epistemic agent reassesses one's prior evidence, perhaps in light of newly acquired evidence. In the high-stake case, the protagonists are not to bring in the alternative to actuality because they already knew; practicality does not directly change the epistemic standard, but only the ethical standard. The question, again, is whether, holding the epistemic context constant, their margin of error for their belief that the bank will be open the next day changed. Epistemic context interacts evidence with epistemic possibility through understanding. With a better epistemic competence, the same evidence rules out more epistemic possibilities, and, on reasonableness, relaxes such restriction when in doubt, lowering one's corresponding confidence. Credence is the probability on evidence, whereas, with better epistemic competence, one's confidence would have better conformed to one's credence.

If, going back to the case of doubt, as in the high-stake case, one failed to resolve the doubt positively, via knowing that the bank is indeed open the next day, one's doubt would have been one that failed to, in the corresponding epistemic situation, made any epistemic justificatory difference – if it possibly made some epistemic justificatory difference. The difference – ethical but not epistemic – it made was in leaving the protagonists choosing a less justified option – the decision to wait in the long queue on the Friday to deposit the pay cheque. Therefore, the doubt is at best only ethically relevant, on supposition of its epistemic relevance, without being epistemically relevant, and thus, by (3) one ought not have doubted it. If it made any epistemic justificatory difference, one would have come to know better, perhaps knowing that one knew, of the fact, and thus the alternative had been negatively relevant epistemically.

Notably, the epistemic risk, like relevance, could have also been negative. What one doubts, albeit only suppositionally, to be a fact, and evaluates of its belief with some epistemic risk, might have turned out to pose no epistemic risk because indeed a fact. This negative epistemic risk interacts, through the high-stake case, with the practical cost of the epistemic wrong. Suppose the bank is indeed not open the next day. It would have been an epistemic wrong to have believed that it would

be open, bearing some practical cost of having thereby made a corresponding decision to deposit the pay cheque the next day. The epistemic risk of having believed falsely differs from its practical cost; the former can be reduced via better evidence acquisition and better understanding, but the latter cannot. Through projecting the possibility in which one decided to deposit the next day and got seriously ill from having caught a cold through having no electricity to sustain the heating in one's flat, the practical cost of the epistemic wrong is illustrated thus.<sup>11</sup>

Interestingly, holding one's ethical profile constant, if one were to calculate the product of the practical cost and the epistemic risk of the putative epistemic wrong, one would have obtained something that covaries, ignoring the product of utility of the outcome and its probability given act, with the expected utility, albeit in everyday utility, of the series of decisions based on the epistemic wrong. The constancy of the ethical profile entails the whole series of decisions to have been fixed, such that, in believing in the bank being open the next day, one would have decided to deposit the pay cheque the next day. However low the epistemic risk, given the high practical cost, and thus its great disutility, the expected utility, putatively, shall be of a large negative number. This explains the pull towards evaluating, not depositing on the Saturday as a rational, if not only reasonable, decision.

However, it was only a putative epistemic wrong, but not in fact so. The epistemic risk correlates with the credence of the opposite of the alleged fact to be believed; for it to be an epistemic wrong, the epistemic risk would have been the opposite being what one ought to have believed. The putative epistemic wrong is committed through believing that the bank will be open, and the credence of this fact thus correlates with the putative epistemic risk. Shall it, in fact, have been an epistemic right, instead of an epistemic wrong, the negativity of the epistemic risk of the epistemic decision would have interacted with the large negative number on the practical side, yielding a large positive. It is interesting to note that, here, the expected utility is through crossing a great disutility with a great negative probability through a negative epistemic risk – shall the correlation between epistemic risk and credence cash out theoretically favourably.<sup>12</sup>

---

<sup>11</sup> Without going into David Lewis' Principal Principle, I assume here that epistemic proximity of a fact relative to an epistemic agent is to approximate its objective chance. However, in this paper, I assume an epistemic actualism with which the actual world, which is the object of one's knowledge, is the way it is but could have been otherwise.

<sup>12</sup> As with the continuity of expected utility, Steele and Stefánsson (2015/2020) note in the following passage. 'Is there any probability  $p$  such that you would be willing to accept a gamble that has that probability of you losing your life and probability  $(1-p)$  of you gaining \$10? Many people think there is not.' Suppose that one is to receive \$10 to go solo skydiving. There shall have



Lastly, with regard to epistemic safety, I propose the following formula:

*Epistemic safety*

$$\text{epistemic safety} = 1 - \text{epistemic risk}$$

When a fact is more epistemically proximal, it is epistemically safer to believe in it. The proportion of possibilities in which the fact obtained among all possibilities given the decision to believe with a piece of evidence constitutes, through being the probability of that fact having obtained given evidence when evaluated against background evidence, epistemic safety of the belief. This is the upshot of epistemic proximity proposed in section 2. Given it was a belief through rejecting the null hypothesis, on supposition of the truth of the null hypothesis, epistemic risk correlates with the  $p$ -value in statistics because it is constituted by the likelihood of having committed an epistemic wrong through believing falsely given evidence. Shall the null hypothesis, being an alternative to the truth of the belief, be actually false, the epistemic risk is negative, yielding an epistemic safety greater than 1.0. This violates the probability axioms, if epistemic safety is just probability on evidence, and I shall attempt the following explanation. Considering the role of doubt in forming higher-order knowledge, I conjecture the phenomenon of probability greater than 1.0, and correspondingly negative probability, to contribute to knowledge of knowledge. On supposition of the falsity of one's belief, the opposite epistemic safety shall be 1.0, constituting an epistemic risk of -1.0 in the original belief. If this supposition is just the supposition of the truth of the null hypothesis, then, when it is rejected at  $p$ -value 0.0, such as when the sample exhausts the whole population, this is when one removes the doubt of one knowing, and comes to know that one knew. The high-stake case, therefore, may be understood as one in which one's failure to know that one knew drives one to give up the original belief. The anxiety that brought about the doubt, if epistemically reasonable, produced an act that is, even if reasonable, irrational because one did worse in depositing having to wait in a long queue. The epistemic decision to give up the original belief, because one in fact did worse than otherwise epistemically, is also epistemically irrational. Reasonable doubt invites epistemic agents to consider the truth of a hypothesis contrary to what one believed, and this, as proposed here, may be understood using

---

been no assignment of probability of death that one would have decided to do so. The negative epistemic risk of believing in the safety of solo skydiving, having crossed the great disutility of death, shall have bypassed the counterexample, restoring continuity. The problem is with what it is when one assigned a negative probability. I gesture at credence as what confidence must conform to, with negative credence simply indicating the obligation to have a positive degree of belief in the opposite.

epistemic risk. Reasonableness, via the phenomenon of doubt, is crucial in statistics with regard to considering alternatives through questioning what belongs to the reference class. The epistemic context in epistemology shall inform the discussion – if the former does not simply constitute the latter.

## 5. Conclusion

When taking an epistemic risk, one evaluates the putatively potential epistemic wrong. When the epistemic wrong is epistemically polluted by some practical cost through what is at stake, salience confuses what is only ethical with the epistemic. What had only been ethically relevant would have seemed epistemically relevant because of our psychology of doubt. Epistemic margin for epistemic error sets some epistemic standard, and, indeed, people with a different ethical standard for safety would have a different standard, given their situation, of epistemic safety. A belief based on an incoming piece of evidence that had been epistemically immune from error could have lost its immunity without any change in the epistemic context because the epistemic standard was modified. Although practicality sometimes modifies epistemic standard, I differentiated worlds of ethical proximity through ease of modification of facts from worlds of epistemic proximity through similarity. The difference between epistemic proximity and ethical proximity is to be stringently observed to avoid confusing ethical relevance with epistemic relevance. The latter determines reasonableness of our doubt, but the former does not – not immediately anyway.

## References

- Cheung, Wai Lok. forthcoming. “Epistemic possibility: Kripke versus Soames”. *Australasian Journal of Philosophy*. In press.
- Cheung, Wai Lok. 2017. *Referential presuppositions, epistemic possibilities, and epistemic content*. The Chinese University of Hong Kong. M.Phil. thesis. <<https://repository.lib.cuhk.edu.hk/en/item/cuhk-2188004/>>
- Cheung, Wai Lok. 2024. “Epistemic responsibility: an agent’s sensitivity towards the world”. *Logos & Episteme*. 15(4):389-403. <<https://doi.org/10.5840/logos-episteme202415431>>
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Clarendon Press.
- Pritchard, Duncan. 2005. *Epistemic Luck*. Oxford: Clarendon Press.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford: Clarendon Press.
- Steele, Katie, and Stefánsson, H. Orri. 2015/2020. “Decision Theory”. *Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/entries/decision-theory/>. Accessed 2024.1.23.
- Williamson, Timothy. 1994. *Vagueness*. New York: Routledge.