CONDITIONALS AND KK

JOHN HAWTHORNE AND YOAAV ISAACS

Introduction

In this paper we make vivid an interesting tension between two ideas. The first is a principle about knowledge that, while somewhat unpopular, has a number of recent adherents. It is the KK principle:

KK: If you're in a position to know p, then you're in a position to know that you're in a position to know p.[1]

The second is an attractive principle concerning indicative conditionals:

Felicity Heuristic: If one knows that one knows q and knows that p is epistemically possible, then one will be epistemically well-placed to felicitously assert 'If p, q'.[2]

Section 1 provides a *prima facie* case for the Felicity Heuristic. Second 2 introduces the infelicity data that will be key to the paper, and presents a natural account of that data that depends upon a rejection of KK. Section 3 explains why, in the presence of KK, the relevant infelicity data can only be reasonably accommodated by rejecting the Felicity Heuristic. Section 4 explores a defensive strategy that the KK proponent might employ, and concludes that it is untenable. In an appendix, we compare the Felicity Heuristic to some principles in Rothschild and Spectre (2018).

1. The Felicity Heuristic

The reader may be struck already by the presence of iterated epistemic operators in the statement of the Felicity Heuristic. What's wrong with the following, more straightforward principle?

Simplified Felicity Heuristic: If one knows q and it is epistemically possible that p, one is in a position to felicitously assert 'If p, q'.[3]

[1] A related formulation is that if one knows one is in a position to know that one knows; Goodman and Salow (2018) employ this variant.

[2] Of course, there may be myriad other reasons why a conditional is infelicitous to assert: it may be impolite, it may irrelevant, it may have already been established, and so on. The idea of this principle is that when the relevant conditions are met, a conditional will not be infelicitous due to the inadequacy of one's epistemic position.

[3] In some cases, the conditionals we consider are made smoother by the addition of an "even". We'll indicate such improvements with a parenthetical "even". We don't think anything crucial depends on the presence or absence of these "even"s. For more on "even if" conditionals, see Gomes (2020), Barker (1994), and REDACTED.

In this connection, it is important to notice that there are cases where one knows (and perhaps even knows that that one knows) q, one does not know whether or not p, but where one will still not be comfortable asserting 'if p, q'. Suppose that one knows—and indeed knows that one knows—that one's mother is alive. One further thinks that one knows and thinks that one knows that one knows that one's friend Jones is alive. Moreover, one does not think of one's knowledge that one's mother is alive as in any obvious way more secure than one's knowledge of Jones' well-being (as one would in a case where one is currently looking at one's mother but using more indirect methods to form a view of Jones' current well-being). But as it happens, Jones has died of an unexpected heart attack ten seconds ago and the news hasn't yet spread. Suppose, uninformed of the news about Jones, one considers the conditional

> If one out of Jones and my mother has just died, it is not my mother.

One will not be comfortable asserting this conditional even though (i) one knows that one knows the consequent, and (ii) one does not know whether or not the antecedent is true.

Notice that is crucial to this example that the antecedent is something one thinks one knows to be false. If one knows that one knows q and knows that p is something one cannot rule out, then the conditional "if p, q" will sound impeccable. Suppose for instance that one knows and knows that one knows that Jones is coming to the party but knows that for all one knows, Smith won't. Then the conditional

> (Even) if Smith doesn't come, Jones will.

is totally unproblematic. In summary, there are some straightforward counterexamples to the Simplified Felicity Heuristic that are are not counterexamples to the Felicity Heuristic. (Indeed, we know of no counterexamples to the Simplified Felicity Heuristic which apply to the Felicity Heuristic.) Moreover, a cursory look at cases where one knows that one knows a consequent but knows an antecedent is epistemically possible lend plausibility to the Felicity Heuristic.

Further support for the Felicity Heuristic comes from the interaction of indicative conditionals with epistemic modals. Consider the following inferences:[4]

> She might not be there but he can't but be there. So, (even) if she isn't there, he is.
>
> He might fail, but she must pass. So, (even) if he fails, she won't.
>
> Smith must be dead by now, though Jones might still be alive. So, (even) if Jones is alive, Smith is dead.

The reasoning sounds impeccable. It is thus natural to think that the following inference rule is valid:

> Might p
>
> Must q
>
> Therefore, (even) if p, q.

[4]Note to the reader: Please read instances of 'has to', 'must' and so on as epistemic, not deontic.

Once the validity of this inference pattern is embraced, there is a reasonably short path to the Felicity Heuristic. Assume now the popular idea that assertion is normed by knowledge[5], we can see that insofar as one is epistemically well placed to assert the two premises, one is epistemically well placed to assert the conclusion. It is also standard to think that when what you know is consistent with p, that makes for the truth (in a natural and standard context) of 'It might be that p' and when you know q, that makes for the truth of 'It must be that q'. Assuming these connections, someone who knows that for all they know p and knows that they know q will be epistemically well placed to assert the two premises. Putting this all together we get a persuasive case for the Felicity Heuristic.

2. Some Infelicity Data

There are 1,000 fair coins that are about to be flipped in sequence. Here are some schemas for conditionals that are quite obviously unassertable:

> (1) (Even) if the first n coins come up heads, the $n + 1^{\text{th}}$ coin will not come up heads.
>
> (2) (Even) if the first n coins all come up heads, the first $n + 1$ coins will not all come up heads.

Anyone who asserted such conditionals would seem to have fallen prey to the gambler's fallacy, erroneously taking a long sequence of coins landing heads as decisive evidence that the next coin will land tails.[6]

We will restrict our attention to conditionals about the 1,000 coins that are about to be tossed. It bears emphasis that our data here concern assertability. Whether certain of the conditionals, though unassertable, are nevertheless true, is a more delicate matter. Various authors, including recently Williamson (2020), have defended the thesis that simple indicative conditionals of English, of the form "if p, q", have the truth conditions of the material conditional. If they are right, then all such conditionals with a false antecedent, as well as those with both a true antecedent and consequent get to be true.

Relatedly, whether certain of these conditionals, though unassertible, nevertheless express knowable propositions is also a delicate matter. If the material conditional analysis is correct and one knows that there will not be 100 heads in a row then (assuming single premise closure) one will know the proposition expressed by '(Even) if the first 100 coins come up heads the first 101 coins will not come up heads'. Our inclination is to reject the material conditional analysis. But since we're restricting our attention to facts about assertability[7], that is an argument for another occasion.

---

[5]For more on this norm see Williamson (2000).

[6]For more on the gambler's fallacy, see Laplace and Solovine (1814) and Kahneman and Tversky (1973).

[7]Note that the material conditional view has to concede that the felicity data do not pattern with the truth conditions provided by the material conditional analysis. Williamson (2020), for example, is explicit that ordinary heuristics for assertion and denial are often poor guides to truth.

The conditionals conforming to the schemes above are uniformly unassertable. Since the Felicity Heuristic is good, we know that under certain conditions those conditionals would be assertible. So why is it impossible for those conditions to obtain? If, for the relevant *p*s and *q*s, it could not be that *q* is known and *p* is epistemically possible, then that would be sufficient answer. However, there is a convincing case that in some situations, *q* is known and *p* is epistemically possible. To see this, we assume two principles: Coin Modesty and Coin Anti-Skepticism.

Coin Modesty: For some *n*, one is not in a position to know that the first *n* coins will not all land heads.

Coin Anti-Skepticism: For some *n*, one is in a position to know that the first *n* coins will not all land heads.

Coin Modesty is clearly true. It is obviously impossible to know that the first coin will not land heads, and that alone secures the truth of Coin Modesty.[8] Coin Anti-Skepticism is also a very plausible principle. We thus assume *inter alia* that the agent is not in the bizarre "skeptical" scenario where all 1,000 coins do land heads. Coin Anti-Skepticism is not a necessary truth, but it holds in a wide variety of normal situations. Admittedly, some will balk at the idea that one can have any non-trivial knowledge about how fair coins will land. Fair coins may seem too "lottery-like" to conduce to knowledge. To this we have two responses, one specific and one general. First, we think it will prove quite difficult to distinguish "lottery-like" propositions from other propositions, and forbid knowledge about the former while allowing it for the latter. There are liable to be various known (or suspected) connections between the "lottery-like" propositions and the other propositions, and these will make it very hard to maintain the pattern of knowledge proposed.[9] Second, the larger philosophical argument we want to make about KK doesn't actually depend on the particularities of coins; they just make for conveniently tractable models. One could instead number one's meat-eating friends and consider whether they've eaten pepperoni pizza in the last week. The analogous principles of Pizza Modesty—that for some *n*, one is not in a position to know that the first *n* of one's friends (relative to a chosen ordering) have all not eaten pepperoni pizza in the last week—and Pizza Anti-Skepticism—that for some *n*, one is in a position to know that the first *n* of one's friends have all not eaten pepperoni pizza in the last week—would work just as well for our purposes. Indeed, we will shortly consider analogously unassertable conditionals that have nothing to do with "lottery-like" propositions (for example, propositions about which objects are red and about which people are alive). Denying all principles analogous to Coin Anti-Skepticism amounts to embracing skepticism itself. Since proponents of KK typically do not want it to be vacuously true (as it would be if we knew nothing) it is dialectically appropriate to think about a case in which we have some knowledge.

Assume that one is in the scenario outlined above and that both Coin Modesty and Coin Anti-Skepticism hold. This combination induces a cutoff: for some *n*, one

[8] One way to make this even more obvious would be to stipulate that the first coin will land heads, and invoke the factivity of knowledge.

[9] See Hawthorne (2004).

is in a position to know that the first n coins flipped will not all land heads, but for all one knows the first n - 1 coins flipped will all land heads. For concreteness, let's suppose that n = 9. Then one knows that the first nine coins will not all land heads, but for all one knows the first eight coins will all land heads. Under such circumstances, the following conditional has a known consequent and an epistemically possible antecedent:

> (Even) if there are eight heads in a row, there still won't be nine heads in a row.

So we have a p and q such that q is known and p is epistemically possible. As we have seen, the combination of a known consequent and an epistemically possible antecedent does not suffice for assertability: the Simplified Felicity Heuristic is wrong. But knowing that one is dealing with a combination of a known consequent and an epistemically possible antecedent does suffice for assertability: the Felicity Heuristic is right. So something must preclude knowing that such circumstances obtain even when they do. But what?

3. Our Analysis

To explain the unassertability of conditionals like (1) and (2), we introduce three ideas—all Williamsonian. The first is about knowledge, the second is about standard assessments of indicative conditionals, and the third is about a norm of assertion.

The first idea is that there is that knowledge requires a margin for error.

> **Margins for Coins:** If all of the first n coins land heads, then one is not in a position to know that not all of the first n + 1 coins land heads.

Suppose, for example, one is in a freaky world where the first 87 coins all land heads. Margins for Coins then dictates that one is in not in a position to know that the first 88 coins do not all land heads. One is just a single coin flip away from error, and thus too close to error to count as knowing.

The second is defended at length in Williamson (2020) and belongs to a long tradition dating back to Ramsey (1929). It is that the main heuristic that we use to assess indicative conditionals is to suppose the antecedent and then form attitudes to the consequent conditional on that supposition. Williamson articulates this idea in the form of an empirical conjecture that the following rule is our primary way of prospectively assessing conditionals (though it is only the indicative conditional that concerns us here):

> **Suppositional Rule:** Take an attitude unconditionally to 'If A, then C' just in case you take it conditionally to C on the supposition that A. (Williamson, 2020, p. 19)

The third idea is a natural extension of the idea that knowledge is normed by assertion: we deploy a knowledge norm of assertion when deciding what we are willing to assert, both unconditionally and conditionally. Thus, regarding our conditional willingness to assert—if we think conditional on p, that we do not know q, then, conditional on p we are not willing to assert q. Given the Suppositional Rule,

this means that in such a scenario we will not be willing to assert the indicative 'if p, then q'.

The application to the issue at hand should be clear enough. Take a conditional like (2). For concreteness, let $n = 102$. Let's see what happens when one supposes that the first 102 coins will all land heads. Assuming you are aware of Margins for Coins, then conditional on that supposition, you will think you do not know that the first 103 coins will not all land heads. So, conditional on that supposition, you will not be willing to assert that the first 103 coins will not all land heads. And given the relevant application of the Suppositional Rule, one will not be willing to assert the conditional

> (Even) if the first 102 coins all land heads, the first 103 coins will not all land heads.

The unassertability of conditionals like (1) is even more obvious—one knows that, whatever happens to the first 102 coins, one does not know what will happen to the 103$^{\text{rd}}$.

We are aware that there is another style of explanation for the unassertability of such conditionals, one that appeals to conditional credences. The key idea would be that the epistemic probability of the consequent on the antecedent is only .5, and this underlies our unwillingness to assert the conditional. We prefer to give an account of assertabilty and conditional assertability that is heavily knowledge-oriented rather than oriented towards a notion of epistemic probability that is not fundamentally tied to knowledge.[10] We feel that if the heavy lifting gets done by a notion of epistemic probability that is not tied to knowledge, then knowledge will be something of a side-show (and a potentially unsystematic side-show at that).

Moreover, we're skeptical that a notion of epistemic probability that is not tied to knowledge can handle the relevant heavy-lifting. There are plenty of felicitous indicative conditionals for which the probability of the consequent on the antecedent—relative to epistemic probabilities not tied to knowledge—is low. For example

> If he wins this race, I'll eat my hat.

is a normal thing to say, but it does not commit one to having high credence in hat-eating conditional on the guy winning. The statement does not commit one to endorsing anything like

> On the off-chance he wins the race, I'm highly confident I will eat my hat.

[10]There are large issues about the relation of epistemic probability to knowledge that we cannot enter into fully. If (like Williamson) you have a notion of evidential probability (and perhaps even idealized credences) where knowledge confers probability 1, then it will be problematic to assume that the epistemic probability of the consequent on the antecedent is .5. (The reader will notice that this conception of epistemic probability does not conform to the Principal Principle, at least for non-sceptics about the future and notions of chance drawn from physics, since one can know things about the future that have non-zero chance of being false. See Dorr et al. (2014) and Bacon (2014) for more.)

But if epistemic probabilities not tied to knowledge and the conditional probabilities thereof are what govern the assertability of indicatives, then that's what would happen.[11] The probability that he wins the race—though low—is liable to be non-zero, and the conditional probability of hat-eating on winning is liable to be low. What's really going on is much simpler: you know the antecedent is false, so the conditional is guaranteed to work out.[12]

Here's another case. Suppose that there are identical twin brothers racing against each other in a marathon. They're evenly matched (as one would expect), and so you know that their marathon times will be within an hour of one another. The following utterance seems fine

> They're not going to be more than an hour apart—they're identical twins! If either they finish more than an hour apart or within a millisecond, it's within a millisecond.

But for a notion of epistemic probability not tied to knowledge, the probability that the brothers finish more than an hour apart is liable to be greater than the probability that they finish within a millisecond of one another (or microsecond, etc.). Again, we think that knowledge has a role to play here.[13] Epistemic probabilities not tied to knowledge do not govern the assertability of conditionals.

So far so good. However, as we shall soon see, matters become much more delicate if we introduce KK. In that setting, a clash with the Felicity Heuristic is more or less inevitable.

4. KK Complications

We have presented what seems to us a plausible explanation of why conditionals like (1) and (2) are unassertable. But, interestingly, it is an explanation that is unavailable to proponents of KK (at least to proponents of KK who accept Coin Modesty and Coin Anti-Scepticism). This is because proponents of KK who accept those two principles cannot also accept Margins for Coins. The reasons why are already familiar. As with standard presentations of this kind of argument (by

[11] Again, nothing much depends on this case employing a plain conditional rather than an 'even if' conditional. Consider the following dialogue:

> A: If he wins this race, I'll eat my hat.
>
> B: He might only win by a little bit.
>
> A: Even if he only wins by a little bit, I'll eat my hat.

[12] We won't present a full theory of how knowledge of the antecedent's falsity guarantees that the conditional works out, as we're remaining neutral on the relevant truth-conditions in this paper. Our point is just that the governing architecture is clearly not conditional probabilities untied to knowledge.

[13] For example, consider the following reasoning:

> They couldn't finish more than an hour apart.
>
> They might finish within a millisecond of one another.
>
> Therefore, if they either finish more than an hour apart or within a millisecond of one another, they finish within a millisecond of one another.

Williamson[14] and others) we will assume that the relevant applications of closure are unproblematic, that Margins for Coins is knowable if true, and that the agent in our scenario has taken the trouble to know what she is in a position to know.

First, given Coin Anti-Skepticism, there is some n such that it is known that the first n tosses will not all land heads. Given KK, it is known that it is known that the first n tosses will not all land heads. But then given knowledge of Margins for Coins, it is also known that (the first n 1 coins all land heads $\supset$ it is not known that the first n coins do not all land heads). Then given closure, the agent is in a position to deduce and come to know that the first $n - 1$ coins will not all land heads. And once more, given KK this is something the agent will know that she knows. The machinery of this reasoning can work again and again, with the agent deducing that the first $n - 2$ coins will not all land heads, that the first $n - 3$ coins will not all land heads, all the way down to that the first coin will not land heads. But, by Coin Modesty, the agent cannot know that the first coin will not land heads.

It is no surprise then that proponents of KK present models that deny Margins for Coins.[15] But it is important to notice that in the absence of Margins for Coins, it is much more challenging to explain the unassertability of conditionals like (1) and (2). This is particularly clear with conditionals like (2): Suppose the strongest thing that one knows is that the first 9 coins flipped will not all come up heads. Assuming KK one knows that one knows this. By hypothesis, one does not know whether or not the first 8 coins will all come up heads. Moreover, given the margins-free model that KK needs, the truth of the proposition that the first 8 coins all land heads is perfectly compatible with one's knowing, and indeed, knowing that one knows that the first 9 coins do not all land heads. So to avoid a clash with the Felicity Heuristic, the KK defender must maintain that in any such setting, while it is epistemically possible that the first 8 coins land heads, it is not known to be epistemically possible that the first 8 coins land heads. Can the KK proponent maintain this?[16]

This approach requires that there be three ranges of coins: First, a range of coins such that you know you don't know that not all the coins up to them come up heads. Second, a range of coins such that you don't know whether or not you know that not all of the coins up to them come up heads. Third, a range of coins such that you know that you know that not all the coins up to them come up heads. Here's why: Given KK, knowledge of Coin Modesty requires the existence

[14] See Williamson (2000), chapter 2.

[15] For example, the KK friendly model of coin flipping knowledge presented in Goodman and Salow (2018) is very explicit about denying Margins for Coins.

[16] When discussing KK, authors often focus on somewhat idealized agents who, for the relevant propositions, know them if and only if they are in a position to know them. (See, for example, the discussion of Mr. Magoo in Williamson (2000), and the discussion of agents who have "thought through a matter" and know everything that they are in a position to know, in Dorst (2019), p. 1230.) As part of our idealization, we shall assume that the agents not merely know KK as stated but, at least, for each relevant proposition, know that if they know them, they know that they know them. As with many of the arguments about KK in the literature, it is open to the KK proponent to accept KK but resist these idealizations. This strategy has not been much pursued by those looking to defend KK, and we will not pursue it here.

of the first range and the truth of Coin Anti-Skepticism requires the existence of the third range. And the defensive strategy crucially relies on having the second range act as a buffer between the first and third; without it this approach would be irrelevant to the case.

Interestingly, it is very hard for proponents of KK to coherently maintain the existence of all three ranges. Indeed, given some apparently modest suppositions, proponents of KK are committed to the first range extending through all 1,000 coins, thereby ruling out two of the three needed ranges. These suppositions are as follows: knowledge of KK, knowledge of the infelicity data[17], knowledge of Coin Modesty, knowledge of the Felicity Heuristic, and closure for knowledge.

Here's the issue: Given knowledge of Coin Modesty, you know that it's epistemically possible that the first coin will land heads. Given KK, you know that you know it's epistemically possible that the first coin will land heads. You also know (thanks to knowing the infelicity data) that you are not epistemically well-placed to assert the conditional

> (A) (Even) if the first coin lands heads, the first 2 coins will not both land heads.

But given knowledge of the Felicity Heuristic, you know the following material conditional

> (You know that it is epistemically possible that the first coin will land heads and you know that you know that the first two coins will not all land heads) ⊃ you are epistemically well-placed to assert (A).

But you know that you know that it is epistemically possible that the first coin lands heads, and you know that you are not epistemically well-placed to assert (A). You can thus deduce (and come to know, given closure) that you do not know that you know that the first 2 coins will not both land heads. Given knowledge of KK you can now deduce and come to know that you do not know that the first 2 coins will not both land heads, which is tantamount to knowing that it is epistemically possible that the first 2 coins will both land heads. The known premise that it is epistemically possible that the first 2 coins will both land heads can then serve as input into a new round of reasoning of exactly the same form in order to secure knowledge of the epistemic possibility of the first 3 coins all landing heads. Repeated application of this pattern of reasoning leads to the conclusion that it is epistemically possible that all 1,000 coins land heads, thus violating Coin Anti-Scepticism.

In sum, knowledge of the Felicity Heuristic, of the infelicity data, and of Coin Modesty, together with the truth of closure and Coin Anti-Scepticism, leave no room for the knowledge of KK. This puts the defender of KK in an uncomfortable

[17]The infelicity data is that you are not epistemically well-placed to assert instances of conditionals (1) or (2). We assume that the infelicity data is not an artifact of not having thought things through; even if you knew everything you are in a position to know, you would still not be epistemically well-placed to assert those conditionals.

position. Absent a Sorensonian blindspot[18], the natural conclusion for those persuaded of the Infelicity Heuristic is that one cannot know KK simply because KK is false.

5. A Defensive Strategy

It is important to acknowledge that sensible KK models (which will deny Margins for Coins) will embrace the phenomenon of defeat with open arms. After all, it is pretty much non-negotiable that if one knows that a fair coin is about to be flipped then one doesn't know whether it will land heads or tails. (This is what Coin Modesty amounts to.) Suppose that the strongest thing one knows at the outset is that the first n coins will not all land heads, and that the first n 1 coins do in fact land heads. (Assuming the failure of Margins for Coins, this is a perfectly coherent supposition). For concreteness, let $n = 9$. Suppose one learns that the first 8 coins all land heads. If one retained one's knowledge that the first nine coins do not all land heads, one could then deduce and come to know in advance that the 9th coin will be flipped but not come up heads, a violation of the non-negotiable constraint above. So if you favor KK you will think that while the first 8 coins all landing heads is compatible with knowing the first 9 coins will not all land heads, learning that the first 8 coins all land heads will make knowledge that the first 9 coins do not all land heads go away.

This all suggests one strategy for explaining the unassertability of conditionals like (1) and (2). Here's the idea: when we assess a conditional like

> (Even) if the first 8 coins all land heads, the first 9 coins won't all land heads.

we suppose not merely that the antedecent is true, but that we have learned (that is, come to know) the antecedent. But learning the antedecedent is incompatible with knowing the consequent. If we indeed suppose not merely that the antecedent is true but that we have learned it, this gives the proponent of KK wiggle room to explain our unwillingness to assert the consequent conditional on the antecedent, and thus (given the Suppositional Rule) our unwillingness to assert the conditional itself. This strategy essentially involves rejecting the Felicity Heuristic in favour of a rather different heuristic, namely that a conditional is assertable only if, *on the supposition that one has learned the antecedent*, one is willing to assert the consequent.

This defensive strategy is, it turns out, rather unpromising. As Williamson is careful to emphasize, "when one supposes something, one does not simulate knowing or believing to the point of supposing that one knows or believes it."[19] After all, if one did then conditionals like the following would be assertable in a situation where one has no view whether or not Jones is a thief:

> If Jones is a thief, I believe that Jones is a thief.

[18]See Sorensen (1988). The key idea would have to invoke the idea that one of the key principles is unknowable. Either The Felicity Heuristc, though true, is unknowable by people when asserting conditionals, or else the KK principle itself will have this status.

[19]Williamson (2020), p. 27

If Jones is a thief, I have learned that Jones is a thief.

But those conditionals are quite clearly unassertable. And it would predict the following to be unassertable:

If Jones is a thief, I don't know it.

But that quite clearly is assertable. In sum, we are quite careful in our suppositional reasoning not to suppose we have learned p upon supposing p. And given that, this defensive strategy breaks down.

6. Conclusion

Conditionals like (1) and (2) are uniformly unassertable. But why? We've offered an explanation of their infelicity. Some of the principles employed in this explanation, however, are unavailable to proponents of KK. Insofar as the unassertability of conditionals like (1) and (2) would otherwise be mysterious, they make a positive case for these principles and against KK.

By our lights, the case against KK was already exceedingly strong. But since when do philosophers shrink from piling on?

Appendix: The Felicity Heuristic and some Principles from Rothschild and Spectre

It may be worth juxtaposing the Felicity Heuristic with an assertability-theoretic variant of some principles articulated in Rothschild and Spectre (2018). Here is their original version[20]:

> **Restricted or-to-if:** If you know a statement of the form *A or B* but you do not know that *A* is true or false or that *B* is true, then you are in a position to know that *if not A then B*.

Let's consider an assertion-theoretic version of this principle, one whose consequent is about being in a position to felicitously assert rather than about being in a position to know. That is,

> **Assertoric Restricted or-to-if:** If you know a statement of the form *A or B* but you do not know that *A* is true or false or that *B* is true, then you are in a position to assert that *if not A then B*.

The above discussion shows that we cannot accept this principle. As Rothschild and Spectre are aware, if the strongest thing one knows about the initial number of heads in a row is that there won't be 9 heads in a row, then one will know the disjunction 'either there will not be initially 8 heads in a row or there will be initially 8 heads in a row but not 9 heads in a row'.[21] And the conditions for Assertoric Restricted or-to-if apply here. One does not know that it is true that there will be initially 8 heads in a row but not 9, and one does not know whether it is true that there will not be initially 8 heads in a row. So by the assertion-theoretic version of their principle, the conditional

[20] They consider plain indicative conditionals, but one can easily consider analogous principles about 'even if' conditionals.

[21] This disjunction being *a priori* equivalent to there not being 9 heads in a row.

> If there are 8 heads in a row there will not be 9 heads in a row.

will come out as assertable. And that's a problem, because it isn't. Rothschild and Spectre are aware of this sort of problem, and say it militates for a related principle with additional restrictions:

> **Further Restricted or-to-if:** If you know you know a statement of the form *A or B* but you do not know that *A* is true or that *B* is true, then you are in a position to know that *if not A then B*.

We worry that this restriction doesn't quite get the job done. Let's look at an assertability-theoretic version of this principle too, namely:

> **Assertoric Further Restricted or-to-if:** If you know you know a statement of the form *A or B* but you do not know that *A* is true or that *B* is true, then you are in a position to assert that *if not A then B*.

Consider again the case where one thinks one knows that one's mother is alive and that one's friend Jones is alive. One does not think of either piece of knowledge as being more secure than the other. However, one does know—and indeed, knows that one knows— that one's mother is well. But Jones has just died of an expected heart attack ten seconds ago and the news hasn't yet spread. Here one knows that knows (or at least one can easily know that one knows) the following disjunction (since it is logically equivalent to something that one knows that one knows):

> Either my mother and Jones are both alive, or Jones is not alive and my mother is alive.

In this scenario, one does not know the left disjunct (since it is false). And one will not know the right disjunct (something that one doesn't even believe). But in this setting, one cannot felicitously assert

> If it is not the case that my mother and Jones are both alive, Jones is not alive and my mother is alive.

Here a key point to observe is that while one does not know the antecedent, one does not know that one does not know the antecedent. In short, we think neither of the principles that figure in Rothschild and Spectre are good guides to assertability. Whether the originals are good guides to knowability depends, inter alia, one whether a material conditional analysis of the truth conditions of the indicative conditional is correct, an issue we are aiming to be officially neutral on for the purposes of this paper.

References

Bacon, A. (2014). Giving your knowledge half a chance. *Philosophical Studies* (2), 1–25.

Barker, S. J. (1994). The consequent-entailment problem foreven if. *Linguistics and Philosophy 17*(3), 249–260.

Dorr, C., J. Goodman, and J. Hawthorne (2014). Knowing against the odds. *Philosophical Studies 170*(2), 277–287.

Dorst, K. (2019). Abominable kk failures. *Mind 128*(512), 1227–1259.

Gomes, G. (2020). Concessive conditionals without even if and nonconcessive conditionals with even if. *Acta Analytica 35*(1), 1–21.

Goodman, J. and B. Salow (2018). Taking a chance on kk. *Philosophical Studies 175*(1), 183–196.

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

Kahneman, D. and A. Tversky (1973). On the psychology of prediction. *Psychological Review 80*(4), 237–251.

Laplace, P. and M. Solovine (1814). Essai philosophique sur les probabilités. *Revue de Métaphysique et de Morale 30*(1), 1–2.

Ramsey, F. P. (1929). General propositions and causality. In F. P. Ramsey (Ed.), *The Foundations of Mathematics and other Logical Essays*, pp. 237–255. Kegan Paul, Trench, Trübner.

Rothschild, D. and L. Spectre (2018). A puzzle about knowing conditionals. *Noûs 52*(2), 473–478.

Sorensen, R. A. (1988). *Blindspots*. New York: Oxford University Press.

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

Williamson, T. (2020). *Suppose and Tell: The Semantics and Heuristics of Conditionals*. Oxford, England: Oxford University Press.