One or Two Puzzles about Knowledge, Probability and Conditionals

Moritz Schulz

Technische Universität Dresden

Abstract

Rothschild and Spectre (2018b) present a puzzle about knowledge, probability and conditionals. This paper analyzes the puzzle and argues that it is essentially two puzzles in one: a puzzle about knowledge and probability and a puzzle about probability and conditionals. As these two puzzles share a crucial feature, this paper ends with a discussion of the prospects of solving them in a unified way.

Keywords: Conditionals, Probability, Knowledge, Adams thesis, Closure

1. Introduction

The topic of this paper is a puzzle by Rothschild and Spectre (2018b). This puzzle involves premises about knowledge and probability on the one hand and premises about knowledge, conditionals and probability on the other. If in this puzzle one were to stick to the premises about conditionals, one could learn that a certain form of skepticism is true. Conversely, denying this form of skepticism could lead one to deny certain principles about conditionals. In this way, it could seem that deep questions in epistemology show a surprising and hitherto unnoticed connection to deep questions about conditionals. This would be exciting news.

But there are also reasons to be somewhat skeptical. As I shall argue in this note, the assumptions about conditionals already constitute a puzzle on their own. Moreover, the remaining assumptions in epistemology are (an incomplete) part of a well-known difficulty of squaring seemingly intuitive assumptions about knowledge and probability with similarly intuitive assumptions about a closure-condition on knowledge. Thus, to some extent we have two puzzles in one: a puzzle about conditionals as well as a puzzle about knowledge and probability. But there is also an interesting connection: the puzzle about conditionals has an effect for single-premise inferences which can otherwise only be duplicated with multi-premise inferences. Given that the two puzzles seem to be generated in a similar way, one may conjecture that they might have similar solutions. I explore this possibility in a final section.

Argumenta 6, 2 (2021): 195-210
ISSN 2465-2334

DOI 10.14275/2465-2334/202112.sch

2. The Puzzle

Here is the puzzle presented by Rothschild and Spectre 2018b, which uses the basic setting of Dorr, Goodman and Hawthorne 2014.[1] You know that a thousand fair coins got tossed yesterday. As a matter of fact, not all coins came up heads. According to Rothschild and Spectre (2018b: 473f.), the following five assumptions are plausibly made about this case (I quote the principles directly from their paper):

(Anti-skepticism) You know that not all the coins landed heads.

(Knowledge & Probability) If you are in a position to know something then you cannot assign it a probability of one-half or less.

(Independence) You should treat each of the coin flips as probabilistically independent.

(Restricted Adams Thesis) Where A and B are non-conditional statements about coin-flips in the setup, you should assign a conditional statement of the form *If A then B* as its probability the conditional probability of B given A.

(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 that B is true, then you are in a position to know that *if not A then B*.

As Rothschild and Spectre (2018b) show, these five assumptions are jointly inconsistent. Their argument goes like this. Take m to be the least number so that one knows that not all coins came up heads. By **(Anti-skepticism)**, there is such a number and it is ≤ 1000. Given the minimal choice of m, one does not know that any of the first $m-1$ coins came up heads. And given **(Knowledge & Probability)**, one cannot know that m came up heads. But one does know that either one of the first $m-1$ coins came up heads or m came up heads. So one may apply **(Restricted or-to-if)** to conclude that one knows that if none of the first $m-1$ coins came up heads, m came up heads. What is the probability of this conditional? By **(Restricted Adams Thesis)**, it is the conditional probability and by **(Independence)**, this conditional probability is $1/2$. But this contradicts **(Knowledge & Probability)**: one would know a conditional which only has a probability of $1/2$.

Let me make a few initial comments about the five assumptions constituting this puzzle. Note that the first three assumptions are general epistemological assumptions which are not specifically concerned with conditionals. Call them the *epistemological assumptions*. The last two assumptions, in contrast, are specifically concerned with conditionals. Call those the *assumptions about conditionals*.

Almost all of these assumptions are individually controversial. The principle **(Anti-skepticism)** grants the possibility of knowledge which is purely based on

[1] For related discussion see Bacon 2014 and Rothschild and Spectre 2018a.

probabilistic grounds. Many would hold that just as one cannot know on probabilistic grounds alone that one's lottery ticket lost, so one cannot know on probabilistic grounds alone that in a sequence of coin tosses not all coins came up heads. The second assumption, **(Knowledge & Probability)**, is comparatively much more plausible. It is widely held that for a proposition to be known, it must be likely in the light of one's evidence. If one's rational credences are supposed to match one's evidential probabilities, then it seems to follow that one cannot (rationally) assign to a proposition one knows a probability of 1/2 or less. Even if there is no such requirement on rational credences, one may argue in favor of **(Knowledge & Probability)** by assuming that knowledge implies rational belief, which in turn may require a credence higher than 1/2. The last epistemological assumption could easily be mistaken to be true in virtue of the description of the case. It is built into this description that we are dealing with fair coins which all have the same chance of landing heads. That is, it is fair to assume that the *objective chances* are all probabilistically independent. But this is not what **(Independence)** says: it says that one's rational credences (or evidential probabilities) are independent. But this requires some kind of bridge principle (see Bacon 2014 for discussion).

The final two assumptions about conditionals are well-known for being problematic, particularly when they are considered jointly. **(Restricted Adams Thesis)** is intuitively very plausible but also very hard to satisfy. That said, there are various ways of satisfying it, either in innovative truth-conditional settings or in a non-truth-conditional framework. Lastly, **(Restricted or-to-if)** is known for lending support to the material analysis, for it effectively states that knowledge of the material conditional (the disjunction *not-A or B*) puts one in a position to know the natural language conditional. But this would suggest that the natural language conditional cannot be stronger than the material conditional. As it is arguably not weaker, the material analysis would be vindicated.

This should suffice as an initial overview of the issues one faces when addressing the puzzle by Rothschild and Spectre (2018b). Let us now turn to some of the details.

3. The Puzzle and Closure

To begin with, I would like to highlight the perhaps obvious fact that the first three assumptions—which are not about conditionals at all—are already in tension with a well entrenched principle about knowledge: multi-premise closure under conjunction introduction.[2]

I shall focus on the following closure condition:

(Closure) If you know statements $A_1, \ldots, A_n$, then you are in a position to know their conjunction $A_1 \wedge \ldots \wedge A_n$.

Note that this does not mean that one's knowledge is actually closed under conjunction. **(Closure)** merely states that conjunctions of what one already knows

[2] This is acknowledged by Rothschild and Spectre (2018a) in the context of a related discussion. Cf. also the exchange between Hawthorne and Lasonen-Aarnio (2009) and Williamson (2009).

are within one's *epistemic reach* (I take the latter notion from Egan 2007, 8). By way of conjoining bits of what one already knows, one can come to know a conjunction. This is extremely plausible: competent application of the rule of conjunction introduction seems to be one of the most basic ways of extending one's knowledge by way of inference.[3]

The tension between **(Closure)** and the epistemological assumptions in the puzzle primarily arises because **Independence** implies

> **(Less-than-1)** The probability you assign to not all the coins landing heads is less than 1.[4]

The implication is not hard to recognize (see Bacon 2014 for further discussion). If the coin tosses are independent, the probability of all the coins landing heads is computed by multiplying the probability of the individual coin tosses. If the probability of not all coins landing heads were 1, the probability of all coins landing heads would be 0. But for a product to be 0, one of the factors would have to be 0. This, however, would mean that one could exclude with certainty that a certain individual coin landed heads, which, we may plausibly assume, one cannot.

The three assumptions **Anti-skepticism**, **Knowledge & Probability** and **Independence** are not in *direct* conflict with **Closure**, for **Anti-skepticism** merely provides a single proposition which is known despite not being assigned probability 1.

Yet the setup of the example can easily be duplicated. Simply think of m such sequences of coin tosses being realized simultaneously but independent of each other. For every such sequence i, this will provide a statement A_i expressing that not all coins reported in this sequence landed heads. We may assume that all statements A_i are true and true for the same reason (they may even show the same pattern of heads and tails). As we are evidentially positioned in exactly the same way vis-à-vis the statements A_i, one may generalize **Anti-skepticism** to:

> **(Multiple Anti-skepticism)** You know $A_1, \ldots, A_m$.

If **Closure** and **Multiple Anti-skepticism** are added to **Knowledge & Probability** and **Independence** from the original puzzle, we face again an inconsistency. With **Closure**, one can infer from **Multiple Anti-skepticism** that one is in a position to know the conjunction $A_1 \wedge \ldots \wedge A_n$. By **Less-than-1**, each A_i should be assigned a probability less than 1. Given that all the coin tosses are causally independent of one another, they should be treated as independent (by **Independence**). Hence, we can choose an m large enough so that the probability of the conjunction $A_1 \wedge \ldots \wedge A_m$ is below 1/2, which contradicts **Knowledge & Probability**. This is because the probability of the conjunction is going to equal $P(A_1) \times \ldots \times P(A_m)$

[3] For ways of fine-tuning a closure condition on knowledge, see Hawthorne 2005.

[4] Rothschild and Spectre (2018b) are aware of this consequence but treat it as another reason to deny that knowledge requires probability 1.

which is going to drop further and further the more terms with a value less than 1 are added.[5]

Neither **Closure** nor **Multiple Anti-skepticism** are uncontested principles about knowledge. My point here is merely that one can set up a similar puzzle by substituting the assumptions about multi-premise closure for the two assumptions about conditionals in the original puzzle proposed by Rothschild and Spectre (2018b). What one gains is, roughly speaking, a puzzle about knowledge and probability only. This gives us some reason to think that the puzzle is not essentially about knowledge of conditionals, but is more concerned with the intricate issues surrounding knowledge of facts which are the result of probabilistic processes.

To further substantiate this diagnosis, let me highlight a crucial property of probabilities of conjunctions:

(Fact 1: Probability & Conjunction)

Certainty Preservation. Always: If $P(A) = 1$ and $P(B) = 1$, then $P(A \wedge B) = 1$.

Probability Drop. Often (for $x < 1$): $P(A) \geq x$ and $P(B) \geq x$, but $P(A \wedge B) < x$.

Probability 1 is preserved under conjunction introduction. If you are certain about the truth of A and B, you can be certain about their conjunction. However, this does not hold for any threshold less than 1. For instance, as in the type of case which currently interests us, if A and B have equal probability x with $x < 1$ and are probabilistically independent, then their conjunction will be below the threshold set by x. So unless one is willing to give up on any fixed positive threshold for knowledge, be it probability 1/2 or less, one runs into a problem with **Closure** if knowledge can be accumulated while being assigned a probability of less than 1.

4. The Puzzle and Conditional Probability

In a key respect, conditional probabilities relate to disjunctions similarly to how unconditional probabilities relate to conjunctions. The following is the relevant fact I have in mind (it plays an important role in the debate about conditionals and probability; see Adams 1975; Adams 1998, Edgington 1986; Edgington 1995):

(Fact 2: Probability & or-to-if)

Certainty Preservation. Always: If $P(A \vee B) = 1$, then $P(B|\neg A) = 1$ as long as the conditional probability is defined.

Probability Drop. Often (for $x < 1$): $P(A \vee B) \geq x$, but $P(B|\neg A) < x$.

The two observations concern the relation between conditional probabilities and probabilities of disjunctions. A disjunction which is certain forces the corresponding conditional probability to be 1, while a disjunction which is uncertain may go with a comparatively low conditional probability.

[5] Effectively, we have reached a problem that is structurally very similar to the lottery paradox. An obvious difference is that the present puzzle is concerned with knowledge rather than rational belief. Another difference is that the starting assumptions, and hence their conjunction, are all true.

Here is a way of illustrating the two properties. Suppose it is quite likely that either the gardener or the cook was responsible for a murder. Assume further that this is so because it is fairly likely that it was the gardener, while it is almost certain that the cook did not do it. However, let's suppose that the disjunction is uncertain because it could also have been the butler. And compared to the possibility of the cook having committed the crime, it is much more likely that the butler did it. How likely is it that the cook did it given that the gardener did not do it? Rather unlikely, for it is much more probable that on the assumption that the gardener did not do it, it was the butler. So, despite the disjunction being probable, the corresponding conditional probability is rather low, illustrating the probability drop reported in **Fact 2**.

Note how the situation changes if we make the disjunction certain. In our example, this means that we eliminate any chance of the butler being the murderer. Once it is certain that it was either the gardener or the cook, it becomes certain that it was the cook given that the gardener did not do it. This illustrates the preservation of certainty reported in **Fact 2**.

The similarity between **Fact 1** and **Fact 2** is obvious. It would therefore not be surprising if assumptions about conditionals, disjunctions and conditional probability could generate a problem similar to that generated in terms of knowledge, probability and closure.

5. The Puzzle and Conditionals

As it stands, **Fact 2** is unconcerned with (indicative) conditionals. However, it becomes immediately relevant once we combine it with the first assumption about conditionals Rothschild and Spectre (2018b) introduce, namely **Restricted Adams' Thesis** stating that the probability of a conditional just is (at least in a certain class of cases) the corresponding conditional probability. What we then find are two corresponding theses: **(a)** if a disjunction is certain, then the conditional is guaranteed to be certain, but **(b)** if a disjunction is merely probable, the probability of the conditional will often be improbable.

Now, Rothschild and Spectre (2018b) put forward a second assumption about conditionals—**Restricted or-to-if**—concerning the knowability of conditionals: knowing a disjunction $A \vee B$ puts one in a position to know the conditional *if not A, B* (provided one does not know either disjunct). This is effectively an assumption about *single-premise closure*: knowledge of a disjunction enables one to knowingly infer the corresponding conditional. Given the probability drop reported in **Fact 2**, one can already identify the outlines of why this leads to trouble. As knowledge is not required to receive probability 1 (as a joint consequence of **Anti-Skepticism** and **Independence**) but is required to respect a certain minimal threshold (as expressed by **Knowledge & Probability**), the probability drop possible for probabilities less than 1 is incompatible with the or-to-if inference being knowledge preserving. Thus, the two assumptions about conditionals, **Restricted Adams' Thesis** and **Restricted or-to-if**, generate a structurally analogous problem to that of closure under conjunction while requiring only a certain instance of single-premise closure, not multi-premise closure. In both cases, the relevant inferences are certainty preserving but not probability preserving. Once knowledge with probability less than 1 is granted but a certain threshold is still maintained, we run into parallel problems. Still, that we can in this way mirror the problems with multi-premise closure in terms of a possible case of single-premise closure is very surprising. How is this even possible?

What makes this surprising is the fact that probability is always preserved under logically valid inferences. If A implies B and $P(A) \geq x$, then $P(B) \geq x$ by the laws of probability. The reason is, very roughly, that if the set of worlds in which B is true includes the set of worlds in which A is true, then B cannot be less probable than A. However, if **Restricted Adams' Thesis** is assumed, the or-to-if inference does not preserve probability (**Fact 2**). As a matter of fact, the probability drop for the or-to-if inference is without any lower bound: to any $x < 1$ and any $\epsilon > 0$, we can find a case with $P(A \vee B) \geq x$ but $P(B|\neg A) \leq \epsilon$ (the murder example above can easily be precisified to fit this structure; the present observation is extensively discussed in Adams 1975; Adams 1998). Such a drop in probability is not even matched by multi-premise closure: although conjunction introduction does not preserve a threshold for probability, the probability of a conjunction always has a lower bound determined by the sum of uncertainties of each conjunct.

Given that valid inferences are probability preserving while the or-to-if inference is not if Adams' Thesis is assumed, a natural conclusion to draw is that the or-to-if inference is simply not valid. Further evidence for this conclusion can be gathered by observing that none of the existing semantics for conditionals which can accommodate a version of Adams' Thesis validates the or-to-if inference (Bacon 2015, McGee 1989, Van Fraassen 1976). As a final point, recall that the or-to-if inference is (equivalent to) the inference from the material conditional to the indicative conditional of natural language. As the material conditional is defined by $\neg A \vee B$, the or-to-if inference gives us $if \neg\neg A, B$. Eliminating the double negation allows us then to derive the indicative conditional from the material conditional. But when indicative conditionals are construed along the lines of Adams' Thesis, the point is usually that they appear to be stronger than the material conditional (the material conditional is easily shown not to satisfy Adams' Thesis).

If the or-to-if inference is invalid by the lights of Adams' Thesis, it is of course a question how this can be squared with this inference being knowledge preserving, as **Restricted or-to-if** would have it. To see the problem, take a case which invalidates this inference, that is, a case in which the disjunction is true but the conditional false. If such cases exist, one should be able to gain knowledge about the disjunction without knowledge of the conditional. This possibility would be backed by the following principle:

> **(Knowledge & Inference)** If A does not imply B, then if A can be known at all, A can be known without being in a position to know B.

With this principle in place, **Restricted Adams' Thesis** and **Restricted or-to-if** can be shown to be incompatible. If **Restricted Adams' Thesis** is true, the or-to-if inference is not valid. By **(Knowledge & Inference)**, the disjunction can be known without being in a position to know the indicative conditional. But this contradicts **Restricted or-to-if** (it is easily verified that the side assumptions necessary for these principles to apply—ignorance about the truth of antecedent and consequent, knowability of the antecedent—pose no problem).[6]

The puzzle about indicative conditionals created by **Restricted Adams' Thesis**, **Restricted or-to-if** and **(Knowledge & Inference)** is a real one. It is not easy to reject any of these assumptions. Thus, the assumptions about conditionals in the original puzzle by Rothschild and Spectre (2018b) already form a puzzle on their own as they are inconsistent with an at least *prima facie* harmless principle such as **(Knowledge & Inference)** (more on this principle below, though).

6. An Intermediate Conclusion

Putting all this together, let me try the following diagnosis. The puzzle about knowing conditionals put forward by Rothschild and Spectre (2018b) is actually two puzzles in one. On the one hand, there is the puzzle about knowledge and probability we already face when we aim at a very modest closure constraint. On the other hand, there is the puzzle of how a logically invalid inference (the or-to-if inference) can still be certainty/knowledge preserving (if indeed it has these two features).

It is definitely interesting that we can substitute the second puzzle for the closure constraint in the first puzzle and still get an inconsistency. This is, as we have seen, because the invalidity of the or-to-if inference generates a similar kind of probability drop we otherwise only find with (valid) multi-premise inferences (valid single-premise inferences are always probability preserving).

Yet in my mind this weakens rather than strengthens the puzzle at hand. It is much more puzzling that a modest closure constraint leads to trouble given the three background principles **Anti-skepticism**, **Knowledge & Probability** and **Independence** than that we gain a puzzle by adding two assumptions about conditionals which already form a puzzle on their own.

Nevertheless, given that the two puzzles embedded in the larger puzzle by Rothschild and Spectre (2018b) are structurally fairly similar—both caused by probability drops in inference—one may wonder whether they also have structurally similar solutions.

7. Towards a Solution

It is beyond the scope of this paper to discuss in any detail the deep questions concerning knowledge, probability and conditionals relevant to the above puzzle(s). What I can do, however, is to sketch one particular solution. This solution is the mirror image of a solution I have given to a similar puzzle involving counterfactuals (Schulz 2017, Chap. 4). By way of conclusion, I will briefly consider what this would mean for the general epistemological assumptions about knowledge and probability.

[6] The structure of this problem is known from the debate about conditionals (Edgington 1986; Edgington 1995, Stalnaker 1975; I discuss it in relation to counterfactuals in Schulz 2017, Chap. 4).

It will be helpful to explicitly consider the relation between indicative conditionals and the material analysis. Very roughly, Adams' Thesis implies that the indicative conditionals is not implied by the material conditional, while the or-to-if inference suggests the opposite. So we can put the puzzle in the following form (see Schulz 2017: 104 for a counterfactual version of it):

The Puzzle (indicative version) For most indicative conditionals *c*:

(P) It is possible that the probability of the corresponding material conditional is high while the probability of *c* is low.

(K) It is not possible that one knows the corresponding material conditional without being in a position to know *c*.

The two problematic assumptions (P) and (K) derive from Adams' Thesis and the or-to-if inference respectively. If something like Adams' Thesis holds, then the probability of the material conditional and the probability of the indicative conditional will not align. The probability of the material conditional will often be high, while the probability of the indicative conditional is low (recall the murder example above). As this indicates how Adams' Thesis supports (P), the second assumption, (K), is almost a mere reformulation of **Restricted or-to-if**. As observed earlier, the or-to-inference could be equivalently stated as the inference from the material conditional to the indicative conditional. But then it states that knowing the material conditional puts one into a position to know the corresponding indicative conditional.

Similar to the discussion above, it is easily seen that there is a tension between (P) and (K). In the light of (P), it seems the material conditional cannot imply the indicative conditional. For if it did, how can the probability of the material conditional oftentimes be higher than the probability of the indicative conditional? After all, the laws of probability guarantee that if A implies B, then the probability of B is never lower than the probability of A. On the other hand, (K) suggests that the material conditional must imply the indicative conditional. For if it did not, there are cases in which the material conditional is true while the indicative conditional is false. If one comes to know the material conditional in such a case, one will not be in a position to know the indicative conditional simply because the indicative conditional is false. This is brought out by **Knowledge & Inference**. In sum, the puzzle is that the material seems to imply and not to imply the indicative conditional. Is there a way out?

Of course, one may simply reject one of the horns of the present dilemma. Adams' Thesis is a highly controversial principle, so there are clearly good theoretical reasons to reject it despite its strong intuitive support. Note, however, that although Adams' Thesis gives rise to (P), (P) itself is a much weaker assumption. On almost all accounts which have the indicative conditional stronger than the material conditional, (P) will still be true. Thus, giving up (P) comes very close to defending the material analysis of the indicative conditional. Instead, one could also try to give up (K), perhaps by finding some kind of explanation of why it holds in many circumstances without being true in full generality.

There is one further option which could easily be overlooked. One may also challenge the principle **Knowledge & Inference**. It is fairly clear that this principle is at best a good approximation to a valid principle. To see that it is false as it stands, let p be the proposition 'snow is white' and consider q to be defined as the conjunction of p with 'I exist', a contingent a priori truth (for discussion, see Schulz 2017, 4.3). It is clear that snow being white does not imply that I exist, that is, p does not imply q. But it seems that any time I am in a position to know that snow is white, I am also in a position to know that snow is white and I exist. This is because I am always in a position to know that I exist. Hence, the principle **Knowledge & Inference** fails. Sometimes, a knowable proposition p can be logically weaker than a stronger proposition q, yet any time one is in a position to know p, one is also in a position to know q.

This might provide a means of escape. Could it be that the indicative conditional is logically stronger than the material conditional (as (P) would have it), but any time one is in a position to know that the material conditional is true, one is also in a position to know that the indicative conditional is true (as (K) would have it)? What would the truth conditions of the indicative conditional have to look like for this to be the case?

In Schulz (2017, 9.3) I consider the following semantics for indicative conditionals, modeled after a similar semantics for counterfactuals. The basic idea is to slightly twist Stalnaker's (1975) semantics for the indicative conditional. According to Stalnaker's semantics, an indicative conditional is true at a world if the consequent is true at the closest epistemically possible world at which the antecedent is true (it is vacuously true if the antecedent is false at all epistemically possible worlds). To make this a bit more precise, let w be a salient possible world and C be a set of worlds epistemically possible at w. On Stalnaker's semantics, we would select a closest antecedent-world from C to fix the truth conditions of a conditional at w. But suppose we do not select a closest antecedent-world from C. Suppose instead that any world in C could be selected and that we leave the selection up to chance. The chance of a consequent-world being selected from the antecedent-worlds in C would be given by the probability of a salient probability function. This probability function assigns a conditional probability to the consequent given the antecedent.

By availing ourselves of an operator, the so-called *epsilon-operator*, we may designate an arbitrarily selected world.[7] More precisely, this world is always an antecedent-world among the epistemically possible worlds (if such exist; otherwise the conditional is vacuously true). The chance of an antecedent-world being a consequent-world is the conditional probability of the consequent given the antecedent. An indicative conditional is then said to be true at a world according to this semantics iff the consequent is true at the arbitrarily selected antecedent-world.

A semantics based on arbitrary selection shares with a Stalnakerian semantics the feature that the truth of the conditional is tied to the truth of the consequent at a single antecedent-world. It differs from a Stalnakerian semantics by leaving the selection of this world in a certain sense up to chance.

[7] For more on the epsilon-operator, see Schulz 2017: Chap. 6.1.

There is clearly a lot to say about a semantics of this kind.[8] However, here I would like to focus on its potential to solve the puzzle about indicative conditionals. The crucial thing to note is that thinking about arbitrary antecedent-worlds is very similar to thinking about conditional probabilities. This is no surprise. An arbitrarily selected antecedent-world will have a certain feature depending on how probable it is that a world has that feature given that it is an antecedent-world.

To see this more clearly, consider the example from above: 'If the gardener did not do it, the cook did'. Now consider a world at which the gardener did not do it, which has been arbitrarily selected from the 'the gardener did not do it'-worlds. How likely is it that this world is a world at which the cook did it? The answer is that this likelihood corresponds to the conditional probability of the cook being the murderer given that the gardener did not do it. In this way, the present semantics puts one in a position to secure a version of Adams' Thesis.

What about knowledge of conditionals on this semantics? This turns on the question of when one can know that an arbitrary object has a certain feature. Consider the case of mathematics, where one frequently introduces names for arbitrary numbers. 'Let n be a prime number not greater than 100', we may stipulate.[9] Can we know that n is an odd number? It seems we cannot, for we cannot exclude that $n = 2$. This suggests that a necessary condition on knowledge about arbitrary selection is that we can only know that an arbitrary F is G if all Fs are Gs. This corresponds to the mathematical practice of concluding that all Fs are Gs after having shown than an arbitrary F is G.

If we apply this observation to our semantics of conditionals, we find that in order to know that an arbitrary antecedent-world is a consequent-world, all antecedent-worlds must be consequent-worlds. In others words, if some antecedent-world is not a consequent-world, one cannot know the corresponding conditional. A consequence of this is that knowledge of an indicative conditional requires the conditional probability to be 1.

The present semantics offers a solution to the puzzle about indicative conditionals. The easy part is to show that (P) can be satisfied. Recall that (P) states that oftentimes, the probability of the material conditional is high, while the probability of the indicative conditional is low. As the indicative conditional is stronger than the material conditional on the present semantics, just like Stalnaker's, there is no obstacle to how (P) can be true. If the material conditional is likely in large parts because the antecedent is likely to be false, the indicative conditional can still be unlikely if most of the antecedent-worlds are worlds at which the consequent is false.

The harder part is to show how (K) can be true. How can it be that any time one is in a position to know the material conditional, one is also in a position to know that the indicative conditional is true? This part is harder because the material conditional does not imply the indicative conditional on the present semantics. But it is not too hard either, for knowing that the material truth conditions are satisfied means that among the epistemically possible worlds, any world is ei-

[8] I discuss it primarily in Schulz 2014 and Schulz 2017. See also Andreas 2018, Cross 2019 and Khoo 2020.

[9] For further discussion, see Breckenridge and Magidor 2012.

ther not an antecedent-world or it is a consequent-world. This means that among the antecedent-worlds, all worlds are consequent-worlds. A fortiori, an arbitrary antecedent-world is a consequent-world. Thus, when one reviews the epistemically possible worlds and finds that all worlds are either not antecedent-worlds or worlds at which the consequent is true, one is in a position to see that an arbitrary antecedent-world is a consequent-world. This can explain how knowing the material conditional puts one in a position to know that the indicative conditional is true. The crucial assumption in this argument is that gaining knowledge of the material truth conditions has an effect on the set of epistemically possible worlds, relative to which the truth conditions of the indicative conditional are defined.[10]

By way of conclusion, let us step back and ask: If this is the right solution to the puzzle about indicative conditionals, what does this mean for the larger puzzle introduced by Rothschild and Spectre (2018b)?

The conditional referred to in the puzzle is this:

> If none of the first $(m-1)$ coins came up tails, m came up tails.

Recall that m was chosen to be the minimal number such that one (allegedly) knows that among the m coins at least one came up tails.

This conditional is established by way of an or-to-if inference from the following disjunction:

> Either one of the first $(m-1)$ coins came up tails or m came up tails.

On the present semantics, knowledge of this disjunction puts one in a position to know the indicative conditional, even though the inference itself is not logically valid. For this reason, one would take the following stance on the epistemic status of this disjunction and the corresponding conditional: either they are both known or if the conditional is not known, then the disjunction is not known either. Let the first option be *option 1* and the second option be *option 2*. I shall consider them in turn.

Option 1. On this option, the disjunction is known and the conditional is known on the basis of the disjunction. This option is challenged by the puzzle because Adams' Thesis implies that the conditional's probability is the corresponding conditional probability which is supposed to be 1/2. But the constraint on knowledge and probability excludes knowledge of a proposition whose probability is only 1/2 or less.

If we hold fixed—as the present option requires—that the disjunction is known, it is instructive to see what the present semantics says about the probability of the conditional. If the disjunction is known, then all epistemically possible worlds will be such that either the antecedent of the conditional is false or where it is true, the consequent will be true as well. But this means that an arbitrarily selected antecedent-world is guaranteed to be a consequent-world, for there are no antecedent-worlds which are not also consequent-worlds. This in turn means that on the present semantics, the indicative conditional is guaranteed to be true

[10] This story is essentially the same as the one given by Stalnaker (1975) in his account of why the material conditional *pragmatically but not semantically* implies the indicative conditional.

and will therefore receive probability 1. If this is so, there is no clash with the constraint that a known proposition should have a probability greater than 1/2.

Of course, one thereby denies that the conditional has a probability of 1/2 (bear in mind, however, that one needs to deny this only if one sticks with option 1 which takes the disjunction to be known). But doesn't this clash with Adams' Thesis? Not necessarily. Adams' Thesis just says that the probability of the conditional is the corresponding conditional probability. And this claim is validated on the present semantics. It is just that the conditional probability is taken to be 1 rather than 1/2, because all epistemically accessible antecedent-worlds are consequent-worlds.

Thus, on option 1, one ultimately takes issue with the intermediate consequence of the puzzle which has it that the conditional probability of m coming up tails given none of $(m-1)$ came up tails is 1/2. But isn't this consequence extremely plausible? That the probability of a coin coming up tails is 1/2 given the outcome of a number of independent tosses just seems to be a fact, for which a lot of empirical support could be mounted.

Now, what clearly is a fact is that the *objective chance* in this case is 1/2. This can be granted, even on the present option. However, what one will disagree with is that the kind of probability which makes the constraint on knowledge and probability plausible is objective chance. What one will hold instead is that for a proposition to be known, it's *epistemic probability* should not be 1/2 or less. So, one can grant that the conditional objective chance is low but simultaneously say that the conditional epistemic probability is high.[11]

But is this a plausible move? I think it is if one accepts the assumptions behind option 1. For on option 1, one grants that propositions can be known which have an objective chance of being false. So one is already prepared to accept that at least in one sense one's epistemic probabilities can depart from the (known) objective chances. So if one grants that knowledge that either one of the first $(m-1)$ coins came up tails or m came up tails, it is not too much of a further commitment to say that the epistemic probability of m coming up tails given that none of the first $(m-1)$ coins came up tails is 1.[12]

There is still a second option, though. Given the somewhat unintuitive consequence of a known conditionals with an objective chance of only 1/2, one may start to doubt that the relevant disjunction from which the conditional was inferred is actually known. This brings us to option 2.

Option 2. Recall that the disjunction states that either one of the first $(m-1)$ coins came up heads or m came up heads. That this disjunction is known is a fairly immediate consequence of **Anti-Skepticism**. The latter principle assumed that one can know that not all 1000 coins came up heads. The number m was then chosen to be minimal with the property that one knows (or is in a position to know) that not all m coins came up heads. Thus, knowledge of the disjunction is established without any detour through assumptions about conditionals.

So, basically, on option 2 one would have to challenge **Anti-Skepticism**. Just to get this out of the way, note that the label is a bit of a misnomer. By denying

[11] See Williamson 2009 for a defense of such a stance.

[12] Yet see Bacon 2014 for critical discussion.

Anti-Skepticism, one is not committed to classical skepticism. What one would be committed to is that knowledge cannot be gained on purely statistical or probabilistic grounds. One could still know that one has hands. But one could not know, based on probabilistic information alone, that a coin is not going to come up heads a thousand (or a million) times in a row. Call this position, for want of a better name, *probabilistic skepticism*.

Probabilistic skepticism is a fairly common position. There are more theoretical as well as more intuitive reasons for it. Lottery cases generate a strong intuition that one cannot know that one's ticket is a loser if the only information one has is the probabilistic set-up of the lottery. This invites the question of how the present case is epistemologically different from a lottery case.[13] There is also a worry that **Anti-Skepticism** is incompatible with the idea that what one is a position to know is a function of the evidence one has. It seems that the evidence one has in a case in which "One of the m coins came up tails" is true is the same as in a case in which all the coins came up tails, for in both cases one would possess the same probabilistic information. Thus, going with option 2 seems to be a fairly easy way out.

In conclusion, there are two ways in which one can respond to the puzzle by Rothschild and Spectre (2018b). On both of these options, one can hold on to the two assumptions about conditionals. This means that one has to take issue with at least one of the three remaining epistemological assumptions. On option 2, one would deny **Anti-Skepticism** by holding that purely probabilistic information is never sufficient for knowledge. On option 1, the situation is somewhat more complicated. The crucial assumption is that knowledge is incompatible with an evidential probability of less than 1 (as defended by Williamson 2000, Chap. 10), while it is compatible with an objective/statistical chance of less than 1. Thus, if the constraint labelled **Knowledge & Probability** is understood by invoking a sense of 'probability' which goes with, or is tied to, objective/statistical chance, then one would deny this constraint. In contrast, if this constraint is concerned with evidential probability, or a kind of rational credence which aims at evidential probability, then one can accept this constraint. This is because a much stronger constraint would hold: not only is knowledge incompatible with an evidential probability of 1/2 or less, it is actually incompatible with any probability less than 1. So, on the present evidential interpretation of **Knowledge & Probability**, one would take issue with the fact that a joint consequence of **Anti-Skepticism** and **Independence** is that knowledge can have a probability of less than 1. One may blame **Anti-Skepticism** for this consequence, in which case the present rebuttal effectively converges with option 2. But one may also blame **Independence**: it may be that the coin tosses are statistically independent, but not evidentially. The reason would be that if one knows that at least one coin came up tails, then this assumption gets evidential probability 1. From this it follows that not all coin tosses are evidentially independent, as long as one assumes that none of the individual coins is evidentially certain to come up tails (see p. 198 for the reasoning behind this).

[13] Rothschild and Spectre (2018b) feel that they are not forced to assimilate their case with lottery cases.

In sum, then, one should either be a probabilistic skeptic or require knowledge to have evidential probability 1. As a matter of fact, one may well adopt both of these claims, for they seem to mutually support each other.

References

Adams, E. 1975, *The Logic of Conditionals*, Dordrecht: Reidel.

Adams, E. 1998, *A Primer of Probability Logic*, Stanford: CSLI Publications.

Andreas, H. 2018, "Moritz Schulz, *Counterfactuals and Probability*", *Notre Dame Philosophical Reviews*, URL: https://ndpr.nd.edu/reviews/counterfactuals-and-probability/.

Bacon, A. 2014, "Giving your Knowledge Half a Chance", *Philosophical Studies*, 171, 373-97.

Bacon, A. 2015, "Stalnaker's Thesis in Context", *The Review of Symbolic Logic*, 8, 131-63.

Breckenridge, W. and Magidor, O. 2012, "Arbitrary Reference", *Philosophical Studies*, 158, 377-400.

Cross, C. 2019, "Moritz Schulz, *Counterfactuals and Probability*", *Australasian Journal of Philosophy*, 97, 414-17.

Dorr, C., Goodman, J. and Hawthorne, J. 2014, "Knowing against the Odds", *Philosophical Studies*, 170, 277-87.

Edgington, D. 1986, "Do Conditionals Have Truth Conditions?", *Critica: Revista Hispanoamericana de Filosofía*, 18, 3-30.

Edgington, D. 1995, "On Conditionals", *Mind*, 104, 235-329.

Egan, A. 2007, "Epistemic Modals, Relativism and Assertion", *Philosophical Studies*, 133, 1-22.

Greenough, P. and Pritchard, D. (eds.) 2009, *Williamson on Knowledge*, Oxford: Oxford University Press.

Hawthorne, J. 2005, "The Case for Closure", in Steup, M. and Sosa, E. (eds.), *Contemporary Debates in Epistemology*, Oxford: Blackwell, 22-43.

Hawthorne, J. and Lasonen-Aarnio, M. 2009, "Knowledge and Objective Chance", in Greenough and Pritchard 2009, 92-108.

Jackson, F. (ed.) 1991, *Conditionals*, Oxford: Clarendon Press.

Khoo, J. 2020, "Moritz Schulz, *Counterfactuals and Probability*", *The Philosophical Review*, 129, 489-95.

McGee, V. 1989, "Conditional Probabilities and Compound of Conditionals", *Philosophical Review*, 98, 485-541.

Rothschild, D. and Spectre, L. 2018a, "At the Threshold of Knowledge", *Philosophical Studies*, 175, 449-60.

Rothschild, D. and Spectre, L. 2018b, "A Puzzle about Knowing Conditionals", *Noûs*, 52, 473-78.

Schulz, M. 2014, "Counterfactuals and Arbitrariness", *Mind*, 123, 1021-55.

Schulz, M. 2017, *Counterfactuals and Probability*, Oxford: Oxford University Press.

Stalnaker, R. 1975, "Indicative Conditionals", *Philosophia*, 5, Reprinted in Jackson 1991, 269-86.

Van Fraassen, B. 1976, "Probability of Conditionals", in Harper, W. and Hooker, C. (eds.), *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, Vol. I, Dordrecht: Reidel, 261-301.

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

Williamson, T. 2009, "Replies to Critics", in Greenough and Pritchard 2009, 279-384.