

Probabilities of Counterfactuals

Richard Bradley

London School of Economics

Abstract

The subjective probability of a subjunctive conditional is argued to be equal to the expected conditional credence in its consequent, given the truth of its antecedent, of an ‘expert’: someone who reasons faultlessly and who, at each point in time, is as fully informed about the state of the world as it is possible to be at that time.

Keywords: Conditionals, Suppositions, Probability, Subjunctives, Counterfactuals, Chances.

1. The Suppositional Theory of Conditionals

The Suppositional Theory says that one should believe a conditional to the degree that one believes its consequent to be true on the supposition that its antecedent is or were true, a claim known as the Ramsey Test hypothesis.¹ The theory is widely reckoned to do a good job in explaining the patterns in our attitudes to conditionals and the linguistic behaviour that manifest them: our assertions and denials of different conditionals, the inferences we make with them, and so on. There has been much controversy around the consistency of the Suppositional theory with truth conditional semantics, but I shall set this issue aside.² My concern will instead be to give the theory enough additional content to deal with the challenge of explaining how our attitudes to conditionals vary with their morphology; in particular, with mood and tense.

To get us started, let’s look at a set of conditionals with common component sentences, but varying morphology, concerning Jim, a canny investor who very rarely loses money, and a potential investment in ostrich farming futures, a fashionable financial instrument, the market for which many consider to be a bubble. Suppose that at time t_0 Jim must decide whether or not to invest in ostrich futures. Jim almost always makes money from his investments so I am inclined to believe that:

¹ Versions of the Suppositional theory have been proposed by, amongst other, Adams (1965, 1975), Bradley (2017), Edgington (1995, 2004, 2008), Skyrms (1981), Stalnaker (1984) and McGee (1987). The origins of the Ramsey Test lie in Ramsey 1929.

² For the record I believe this controversy to be resolved: there is no inconsistency.

(1) If Jim does invest in ostrich futures, then he will make a packet.

On the other hand, from my less well-informed perspective, speculation on ostrich futures looks like a good way to lose a lot of money. So, I am inclined to believe that he won't make the investment but that:

(2) If Jim were to invest in ostrich futures, he would make a loss.

Suppose that by time t_1 Jim has made his decision but I have not learnt what it is or how ostrich futures have performed. Given this, I will continue to regard it as probable that he didn't make the investment but that:

(3) If Jim did invest in ostrich futures, he made a packet.

(4) If Jim had invested in ostrich futures, he would have lost a packet.

I subsequently learn (at time t_2) how ostrich futures have performed. Suppose that they have performed badly. Then I would continue to believe (4) and that Jim didn't make the investment, but that, contrary to (3), he lost a packet if he did. Suppose instead that by t_2 I know that ostrich futures have performed well. Then I would continue to believe (3) and that, contrary to (4), he would have made a packet if he had invested in ostrich futures.

The challenge is to explain this pattern of beliefs with the Suppositional theory: in particular why our attitudes to the indicative conditional and the corresponding subjunctive differ at t_0 and t_1 but not at t_2 . These sentences by no means exhibit all morphological variation in conditional sentences of potential interest, but they serve to illustrate four important classes: the forward-looking or future-oriented indicatives and subjunctives ((1) and (2) respectively) and the backward-looking or past-oriented indicatives and subjunctives ((3) and (4) respectively). So getting them right is an important step in filling out the theory.

A terminological note. Subjunctive conditionals are often termed 'counterfactuals' in the philosophical literature, on the grounds that the subjunctive mood is typically used to convey the belief that the antecedent is false. But, as many authors have observed, this implication holds only for the backward-looking subjunctives. Moreover, believing the antecedent of an indicative false doesn't preclude asserting it, let alone rendering it senseless (as in 'If he is in the room, then I must be blind'). On the other hand, the indicative-subjunctive distinction is also problematic: von Fintel (2011) calls it 'linguistically inept'. But getting the linguistic nuances right doesn't, for present purposes, justify the increase in complexity that comes with refining the distinction. So I will continue to use it and reserve the term 'counterfactual' for backward-looking subjunctives.

To explain how and why our readiness to believe and assert conditional sentences, at different times (or informational contexts), varies with mood and tense, the Suppositional theory draws on two key facts: firstly, that there are different ways of supposing something true and, secondly, that different modes of supposition are appropriate for the evaluation of indicative and subjunctive conditionals. We might suppose that as a matter of fact something is true, such as when I suppose, to help with my financial planning, that I won't have enough money at the end of the month to pay the rent. Suppositions of this kind I will call *evidential*, since the proposition supposed true is treated like a new piece of evidence that we have acquired. Evidential supposition should not lead one to give up any beliefs about what is in fact true or to adopt ones one knows to be false. I should not, for instance, adopt the belief that I will secure a large inheritance to cover the rent

when reasoning evidentially from the supposition that I will be short at the end of the month.

In contrast to evidential supposition, we might also suppose something to be true, contrary to, or independently of, the known facts, such as when I suppose that it rained yesterday (it did not) in order to consider what I would have done had this been the case. Suppositions of this kind, I will call *interventional*, since we treat the assumed proposition as if it had been made true by an intervention from outside the actual system of causal relationships. Interventional suppositions are often best accommodated by giving up some beliefs that one knows to be true (in actual fact), to allow retention of well-entrenched conceptions about the way that the world works. For example, when supposing that it rained yesterday, in order to think about what I would have done had this been the case, I might have to give up my belief that I went for a walk in the mountains that day, even if I did in fact do so (and have sore feet to prove it).

To harness these distinctions for the purpose of explaining our attitudes to conditionals the Ramsey Test hypothesis must be augmented by postulates about the form of supposition relevant to the evaluation of each kind of conditional. I will make two.

- (RT1) An indicative conditional is evaluated *evidentially*, by supposing that its antecedent is, as a matter-of-fact, true. This involves *adding* the antecedent to our current set of beliefs and then determining whether, or to what degree, the truth of the consequent follows.
- (RT2) A subjunctive conditional is evaluated *interventionally*, by supposing that, potentially contrary to the facts, its antecedent is true. This involves *accommodating* the truth of the antecedent by suspending our beliefs about its causal preconditions and then inferring the truth or falsity of its consequent, drawing on our beliefs about the causal relationships between the two.

To illustrate and assess these two claims, consider the first pair of (forward-looking) conditionals (1) and (2). Although they make ‘opposite’ claims, assertion of both is quite reasonable from the point of view of our two RT hypotheses. The first sentence, being an indicative conditional, is evaluated by supposing that Jim will as a matter of fact make the investment. Since he rarely makes a mistake, it is reasonable to infer from the evidence furnished by his action that we were wrong about the bubble and that the investment in ostrich futures will be profitable. On the other hand, although the corresponding subjunctive conditional is also evaluated by supposing that Jim makes the investment, his so-doing is not treated as evidence for the quality of the investment. This is because we accommodate the supposition of his investing by giving up our belief that he is well informed about the state of the market for ostrich futures in order to retain our belief about the causal consequences of investing in the current state of the market. So instead of inferring the state of the market from the fact that Jim has made an investment, we infer that he will lose money from his investment from what we believe about the relationship between the state of the market, investing and securing profits.

Much the same applies to our t_2 attitudes to the second pair of (backward-looking) conditionals. To evaluate indicative (3), we suppose that Jim did in fact invest in ostrich futures by adding the occurrence of him making the investment to our stock of beliefs and inferring from this that the investment will prove to be profitable. To evaluate (4) on the other hand, we suppose, potentially contrary to

the facts, that Jim made the investment but without making the adjustment to our beliefs about the quality of the investment that we would have to make if reasoning evidentially. Hence, we infer that his investment would have made a loss.

By t_2 however our informational situation has changed. Consider first the case where we have learnt that (as expected) the ostrich futures have performed badly. No doubt we will be all the more convinced that Jim would not have invested in them. Nonetheless, supposing evidentially that he has in fact done so, he must have lost money. Similarly, unless we think that Jim's investment would have *caused* the ostrich futures to perform well, we must conclude that on the contrary-to-fact supposition that he made the investment, he lost money. This explains why, by t_2 , the backward-looking indicative and subjunctive conditionals that we accept coincide in what they assert.

2. Probabilities of Conditionals as Conditional Probabilities

To make precise this informal explanation of the stylised facts represented by our attitudes to the morphologically different conditional sentences concerning Jim's investment, we need to be correspondingly precise about the differences between evidential and interventional supposition and how they should be embedded within the Suppositional Theory. Let us start by stating the Ramsey Test hypothesis in a general form and introducing the formal vocabulary necessary to do so.

Throughout I will assume a background set L of sentences (denoted by italicised capitals) closed under the sentential operations of conjunction \wedge , disjunction \vee and negation \neg , and subjective probability functions Pr_i on L that represent the degrees of belief of a rational agent in the L sentences at time t_i (more exactly the degree to which they believe at that time that what these sentences say is true). The sentence $A \wedge B$ will usually be written AB . Sentences expressing the occurrence of an event at a particular time t_i will be denoted by i -subscripted capital e.g. X_i is the sentence asserting that X occurred at t_1 . The subscript will be dropped when the time is fixed for, or is irrelevant to, the discussion. Indicative conditionals will be represented by expressions of the form $A_i \rightarrow C$; subjunctive conditionals by expressions of the form $A_i \rightharpoonup C$. Either type is called backward-looking if the time of its assertion is later than t_i and forward-looking otherwise.³ I will make no assumptions about the logic of these conditionals other than they should obey Modus Ponens.

Most of the literature on the Ramsey Test hypothesis has focused on the version appropriate to indicative conditionals, known as Adams' Thesis, and which asserts that the credibility (and hence assertability) of an indicative conditional is the conditional probability of its consequent given its antecedent. More formally, for any non-conditional sentences $A, C \in L$ such that $\text{Pr}(A) > 0$ and any time t_i :

$$\text{(Adams' Thesis)} \quad \text{Pr}_i(A \rightarrow C) = \text{Pr}_i(C|A)$$

Note that if Adams' Thesis is to apply to backward-looking conditionals even when the (current) probability of its antecedent is zero, then conditional

³ No implication that the indicative and subjunctive conditionals have different semantic content should be drawn from this choice of formalism. Indeed, I am inclined to believe the contrary: that they have the same content but that this content is evaluated differently depending on mood of the sentence asserting it. But nothing will depend here on whether this is true or not.

probabilities must be defined for zero-probability sentences. For this the usual ratio definition of conditional probability will not suffice and so we must draw on one of the alternative treatments, such as that of Renyi 1955, which allow conditional probability to be defined for conditions that have measure zero.

Now, at both t_0 and t_1 , the conditional probability that Jim will make, or has made, a packet, given him making the investment, is high because in almost every credible circumstance in which Jim makes an investment, the investment is profitable. But by t_2 I know how the ostrich futures have performed. And in case they have performed badly, the conditional probability of him having made money, given an investment in them, has to be at or near zero. So Adams' Thesis does a good job explaining my acceptance of the indicative conditionals (1) and (3) at t_0 and t_1 and the conditions under which I would accept or reject indicative (3) at t_2 . On the other hand, since my attitudes to the subjunctive conditionals (2) and (4) are different to my attitudes to the corresponding indicatives (1) and (3), they cannot also be explained by Adams' Thesis.

Adams himself recognised very early on (in Adams 1970) that backward-looking indicatives and subjunctives were evaluated differently, offering as evidence his (now famous) example of the difference in our attitude to the past indicative "If Oswald didn't kill Kennedy, someone else did", which we regard as almost certainly true (as Adams' Thesis predicts), and the counterfactual "If Oswald hadn't killed Kennedy, someone else would have", which we don't. He nonetheless speculated that some other conditional probability would explain our attitude to the counterfactual.

One idea, explored and then rejected by both him (Adams 1975) and Edgington (2004), is that the degree to which a counterfactual should be believed at the time of its utterance is the conditional credibility of its consequent, given its antecedent, at the earlier time at which the truth of its antecedent was resolved. Let's call this the Prior Conditional Credence view of the probability of conditionals. More formally, if the truth of the antecedent of the conditional $A_0 \rightarrow B$ was resolved at t_0 , then at time $t_i \geq t_0$:

$$\text{(ConCred)} \quad \Pr_i(A_0 \rightarrow B) = \Pr_0(B|A_0)$$

The Prior Conditional Credence theory offers no account of forward-looking conditionals and hence of our attitudes to the contrasting pair (1) and (2), but its prescriptions do fit with some usage of backward-looking subjunctives. Here is Edgington's example. I say "If I leave before noon, I will be on time for my appointment with the doctor". I am distracted and fail to leave. Later I say (regretfully) "If I had left before noon, I would have made my appointment". It also explains the difference in our attitudes to the two Oswald-Kennedy sentences. For while Adams' Thesis accords with our willingness to accept the indicative, ConCred accords with our unwillingness to accept the matching subjunctive, assuming that the subjective probability at the time of Kennedy's assassination of someone other than Oswald attempting it was very low.

On the other hand, the Prior Conditional Credence view doesn't do very well in our running example. For one thing, it doesn't provide any explanation of the difference in the attitude I take to the counterfactual (4) at t_1 and at t_2 , since it predicts that its probability equals my t_0 conditional probability of him making a packet, given that he invests in ostrich futures, at *all* times later than t_0 . Furthermore, it predicts that the probability of subjunctive (4) at t_1 will agree with that of indicative (1) at t_0 . But it does not. While the conditional probability is high of

Jim making money, given that he invests, the probability is low that he would have made a packet had he invested in what we believe to be a bad venture. In fact, the t_1 probability of counterfactual (4) agrees, not with the forward-looking indicative (1), but with that of forward-looking *subjunctive* (2), while its the backward-looking indicative (3) that agrees with the forward-looking indicative (1). The general message is thus not that the counterfactuals align with forward-looking indicatives uttered earlier but that, *so long as no new relevant information is obtained*, both backward-looking subjunctives and backward-looking indicatives align with their corresponding forward-looking subjunctives and indicatives.

When new information *is* obtained, in particular about the truth of the consequent, then this alignment breaks down. Edgington's example continued: I arrive late for my appointment having failed to leave before noon. Before I can make my excuses to the receptionist, he says "I am surprised you made it. I heard that most trains have been cancelled". I now say to myself "Even if I had left before noon, I would not have made my appointment". Similarly, in our running example, what I learn at t_2 breaks the alignment between the backward-looking conditionals and the corresponding earlier forward-looking one. If I learn that ostrich futures have performed badly then my earlier conviction that if Jim did make the investment then he made money is overturned. So, my earlier acceptance of (3) gives way to its rejection. On the other hand, if I learn that they performed well, then it's my earlier conviction that had Jim made the investment he would lose money that is overturned. So, my t_1 acceptance of (4) turns into rejection of it at t_2 .

3. Probabilities of Conditionals as Conditional Chances

The Prior Conditional Credence view is clearly inadequate and it has no current defenders. But a similar theory, that draws on objective conditional probabilities rather than subjective ones, is more promising. According to what I will call the Prior Conditional Chance view one should set one's degrees of belief in a subjunctive to what one takes to be the conditional objective probability or conditional chance, at the time of the resolution of the truth of the antecedent of the conditional (or immediately prior to it), of the truth of its consequent given the truth of its antecedent. More formally, let CH_0 be a random variable ranging over a set $\{\pi^j\}$ of probability functions, measuring the t_0 -chances. Then, on the Prior Conditional Chance view, at any time t_i your degree of belief in the counterfactual $A_0 \rightarrow C$ should go by your t_i expectation of the t_0 conditional chances of C given that A , i.e.:

$$(\text{ConCh}) \Pr_i(A_0 \rightarrow C) = \sum_j \Pr_i(CH_0 = \pi^j) \cdot \pi^j(C|A_0)$$

Different interpretations of objective probability or chance will yield different instances of this view. Perhaps the most prominent in the literature is that of Brian Skyrms (1980, 1981, 1988) who takes the relevant objective probabilities to be prior propensities. But Moss (2013), Williams (2008), Joyce (1999) and Pearl (2000) all endorse versions of it—more on the latter later on. (Here I gloss over the fact that Skyrms himself did not endorse this exact view because he was sufficiently convinced by the triviality results of Lewis and others to accept that conditionals do not have truth values. For this reason, he presented his claim as pertaining to what he called the Basic Assertability Value of counterfactuals rather than to their probabilities of truth.)

Most of the exponents of the Prior Conditional Chance view propose it only as theory of the credibility of counterfactuals, but my more general formulation allows that it applies to forward-looking subjunctives as well (i.e. that t_i be earlier than t_0). In this form it offers a general explanation, in terms of the difference between objective and subjective probability, for our contrasting attitudes to the indicative and corresponding subjunctive conditionals concerning Jim's investment. While acceptance of (1) follows from the fact that Jim's investment is evidence for ostrich futures being a good investment, our acceptance of (2) follows from our belief that the objective probability or chance of ostrich futures performing well is low. Similarly, for our acceptance of the contrasting conditionals (3) and (4).

The view also seems to capture the way in which our attitude to counterfactual (4) changes between t_1 and t_2 in response to the information acquired about the performance of ostrich futures. For definiteness suppose that you give non-zero credence to just three hypotheses concerning the conditional chances at t_1 of making a packet, given an investment in ostrich futures: that they are zero, that they are a half and that they are one. At t_1 your credence will be concentrated on the first of these, or perhaps the first and second. This explains why at t_1 you regard (4) as credible: the expected conditional chance at that time of making a packet given an investment is low in virtue of the high probability that it is zero. If you learn that, as expected, ostrich futures have performed badly, you will have all the more reason to concentrate your belief on the hypothesis that there was no chance of making a packet from an investment in ostrich futures. But if you learn that ostrich futures have performed well, you are likely to shift probability from the hypothesis that the t_1 conditional chance was zero to the hypothesis that it was one. This explains why you now (at t_2) reject (4).

All of this seems to offer confirmation of the Prior Conditional Chance view. But there is a problem. For the view underestimates the strength of my t_2 attitudes to counterfactual (4). When I learn how ostrich futures have performed, I become *certain* of one or the other of these counterfactuals (depending on whether they performed well or badly). But unless I attached no credibility *at all* at t_1 to the hypothesis that the conditional chance of making money given an investment in ostrich futures is 0.5, my t_2 estimation of this conditional chance will fall short of one. This is because whatever information I get about the performance of ostrich futures is perfectly consistent with this hypothesis. More exactly, on this hypothesis neither a good performance nor a bad one is more likely, given an investment, so the observation of its performance is uninformative regarding the truth of the hypothesis.

Explaining my newfound certainty about the truth of these counterfactuals is an instance of a famous old problem for accounts of conditionals: Morgenbesser's Coin.⁴ An indeterministic fair coin is to be tossed and you are invited to bet on it landing heads. You demur, the coin is tossed and lands heads. Your interlocutor says "If you had bet, you would have won". As you don't believe that your betting would have influenced the toss, you are forced to agree. But the fact that the coin has landed heads is no evidence that it is not fair. So your estimate of the prior conditional chances of landing heads stays at 0.5. It would seem to follow that on the Prior Conditional Chance view you should not regret your

⁴ This example was reported in Slote 1978, who attributed it to Sydney Morgenbesser.

failure to bet: what your interlocutor says is as likely to be false as it is to be true. But since what the interlocutor says seems true, this view must be rejected.

A variant. You in fact decide to bet and in due course win it. Someone mistakenly believes that you did not bet and says “If you had bet, you would have won”. You reply “Yes, that’s true. In fact, I did bet and I did win”. Your agreement here with your interlocutor is in line with the widely held semantic principle (known as Centring) that the truth of $A \wedge C$ implies that $A > C$. But it is inexplicable if ConCh is correct, for what they have said is, on this account, as likely to be false as it is to be true.

The explanatory problem presented by the Morgenbesser coin problem and its variant is rather different to that presented by the Oswald-killing-Kennedy one. Before the coin has been tossed we are inclined to accept neither the forward-looking indicative “If you bet, you will win”, nor the corresponding subjunctive “If you were to bet, you would win”; after the hands landing has been observed, we are inclined to accept both the backward-looking indicative “If you did bet, you won” and the corresponding counterfactual “If you had bet, you would have won”. So here our attitudes to the indicative and corresponding subjunctive conditional is the same at every moment of time and the challenge is to explain why both change over time in the same way.

So troublesome has Morgenbesser’s Coin been for theories of conditionals in general (and not just for those under consideration here) that it is worth considering an error theory for the intuition that drives it. Consider a similar case involving a deterministic coin that lands heads if and only if some set of initial conditions C hold. The observation that the coin has landed heads licenses the inference that C is the case. From which it does follow that had you bet on heads you would have won, since the coin always lands heads when C holds. Now the error theory I have in mind says that we mistakenly believe this to be true in Morgenbesser’s case as well because we treat it as a deterministic case with epistemic uncertainty about the determining conditions, rather than a truly indeterministic one. But we are wrong: the observed actual outcome of an indeterministic process is completely uninformative with regard to the truth of counterfactual claims about what the outcome would have been had it occurred under different conditions (including those causally independent of the process).

The problem with error theories of this kind, as Edgington (2004) points out, is that they license what can only be regarded as wishful or magical thinking. I turn out to have the winning ticket in a lottery with 10,000 tickets. I say “I am glad that I rubbed my rabbit’s foot. For had I not I would have lost”. On the view that the error theory is designed to uphold, this sentence is very probably true. I don’t think that this is an implication that we should accept. But if we don’t, then we must accept that the Prior Conditional Chance view is false.

It is not difficult to identify where things have gone wrong for the Prior Conditional Chance view. In the Morgenbesser’s Coin case and others like it, what we learn about particular outcomes of chancy processes trumps what we know or believe about the prior chances of these outcomes. Edgington (2004) suggests a simple fix: look not to the prior conditional chances of the outcomes but to the prior conditional chances *updated* by any relevant information subsequently received and attach a degree of belief to the corresponding counterfactual equal to your expectation of these updated chances. More formally, let S be a proposition expressing all relevant events occurring between (and including) the occurrence

of the antecedent and the occurrence (but not including) the consequent. Then, what we can call the Updated Conditional Chances view says that at any time t_i :

$$(\text{UpdConCh}) \Pr_i(A_0 \rightarrow C) = \sum_j \Pr_i(CH_0 = \pi^j) \cdot \pi^j(C|A_0 \wedge S)$$

The crucial question for the Updated Conditional Chances view is what information is relevant, i.e. what event S we should conditionalize on. According to Edgington:

The objectively correct value to assign to such a counterfactual [that if A had been the case, then C would have been] is not (or not always) the conditional chance of C given A at the time of the fork; but the conditional chance, at that time, of C given $A \& S$ where S is a conjunction of those facts concerning the time between antecedent and consequent which are (a) causally independent of the antecedent, and (b) affect the chance of the consequent (Edgington 2004: 21).

Edgington's proposal nicely explains why our attitudes to the counterfactual conditionals change over time. To see this, let us continue to assume that the performance of ostrich futures is independent of whether or not Jim makes an investment in them. Then, while at t_1 the expected chance is low of Jim making money, given the prospective investment in ostrich futures, at t_2 the expected updated chances of him making money will be one or zero, depending on how in fact the ostrich futures performed. This explains my t_2 -acceptance of (4). And unlike the Prior Conditional Chance view it correctly predicts not just which counterfactuals we would assert at each time, but also how strongly we would believe them.

What about the troublesome Morgenbesser Coin case? Here, presumably, we want to update the chances by the information that the coin was tossed and did in fact land heads, but not the information that I did not bet on how it would land. And, indeed, the conditional chance of it landing heads, given that it was tossed and landed heads, is one. This explains why we accept the counterfactual "If you had bet, you would have won". The problem is that we should also be willing to assert the counterfactual "If the coin had been tossed, it would have landed heads"; since its truth is the reason why you would have won had you bet. But Edgington's condition (a) does not allow us to update on the coin landing heads since the heads-landing of the coin is not causally independent of it being tossed. So, her proposal doesn't allow us to predict our attitude to this second counterfactual. And, in general, it fails to ensure the high credibility of "If A had been the case, B would have been" in cases in which A and B are both true, because it doesn't prescribe updating the conditional chances of B given that A , by the truth of B whenever B depends causally on A .

It is tempting to conclude that the Updated Conditional Chances view should dispense with restriction (a), for there will often be facts that are *not* causally independent of the antecedent but which nonetheless affect the chance of the consequent in a manner relevant to the evaluation of the counterfactual itself. But Edgington has a good reason for not allowing such information. Suppose that at t_1 Jim invests, not in ostrich futures, but in a housing development which makes him a large amount of money. So, at t_2 , Jim has made a packet. This fact should not by itself ensure the falsity of counterfactual (4), for the performance of the housing market is of no relevance to that of ostrich futures. But the prior conditional chance of Jim making a packet, given that he invests in ostrich futures, updated by the fact that he makes a packet (through his investment in housing) is

one. So without Edgington's clause (a), the Updated Conditional Chances view will prescribe disbelief in subjunctive (4), even in cases in which ostrich futures turn out to perform badly. (Clause (a) blocks this inference because Jim's making a packet from housing is *not* causally independent of his (not) investing in ostrich futures.)

These considerations do not decisively refute the Updated Conditional Chances view since it is possible that some other specification of what information is relevant will deliver the goods. But instead of pursuing it further I want to make a different suggestion which, I will argue, avoids the difficulties of all the accounts considered thus far. The proposal is a very simple one: that the credibility of a subjunctive conditional, counterfactual or otherwise, is the expected *posterior* objective conditional probability of *C* given that *A*, and not the prior one. But to defend it I must first make a detour.

4. Expert Probabilities

We saw that the credibility of 'matching' indicative and subjunctive conditionals can differ because when evaluating the former, but not the latter, we treat the truth of the antecedent as evidence about the state of the (actual) world. To better understand what difference this makes let's look at how conditionals would be evaluated by a perfect Bayesian reasoner (hereafter called Expert) who, at each point in time, is as informed about the state of the world as it is possible to be at that time. By a perfect Bayesian reasoner I mean someone that makes no mistakes in their probabilistic reasoning, draws all the inferences that they should from what they know and none that they should not. By as informed as it is possible to be, I mean that they are apprised of any truths that it is physically possible to learn. This will include most facts about the past, but not *a posteriori* truths about the future. Nor will it include laws of nature or counterfactuals, these being things about which Expert must form beliefs by inference from what they do know.

This characterisation of Expert leaves open difficult questions about what inferences she should draw from what she knows and exactly what it is possible to know at any point in time. But all that is important for present purposes is that the Expert's degrees of belief at any particular time are as good as they can be at that time. They cannot be improved by more information, because none is available. And they cannot be improved by elimination of errors of reasoning, for they make none. It follows that our own degrees of belief will be as good they can be when we have correctly aligned them with those of Expert. In this sense Expert's credences constitute ideal or objective probabilities: not because they are mind or judgment independent, but because they are the probabilities that our credences should aim at.

In the light of these observations, let us consider how Expert would evaluate our four conditionals concerning Jim and his investments. Suppose for the sake of the exercise that Expert is fully apprised at t_0 of the state of the market and that she infers from this a probability for the profitability of various possible investments and for Jim investing in any one of them (drawing also on any accessible facts about Jim's preferences and beliefs). Suppose that by t_1 , she has learnt of Jim's decision and by t_2 of how the various assets have performed.

We saw earlier that my willingness to accept both the indicative (1) and its 'contrary' subjunctive (2) stemmed from the fact that at t_0 Jim's investment was evidentially relevant for me to the question of the state of the market, and hence

whether ostrich futures would provide good returns, even though it was causally irrelevant to it. In contrast, Jim investing or otherwise in the ostrich futures is completely *uninformative* for Expert at t_0 regarding the state of the market. For at t_0 , Expert knows what the state of the market is and so Jim's decision cannot contain information about it that she does not already hold. It follows that, for Expert, the performance of the various possible investments do *not* depend probabilistically on whether Jim invests in them or not. This is of course trivially true when Expert's knowledge of the state of the market suffices for her to predict with certainty how investments will perform. But the independence of the two holds even when it does not.

Suppose, as seems reasonable, that it is necessary and sufficient for Jim to make a packet from ostrich futures that he invests in them and that they perform well. Then it follows from the probabilistic independence of the two that Expert's t_0 degree of belief in Jim making a packet, conditional on him investing in ostrich futures, equals her t_0 unconditional degree of belief in this investment performing well. But given that the performance of this investment is causally independent of Jim's decision, this just the same as her degree of belief in him making a packet on the interventional supposition of him investing. So, by application of RT1 and RT2, Expert's degree of belief in the indicative "If Jim does invest in ostrich futures, then he will make a packet" will equal her degree of belief in the subjunctive "If Jim were to invest in ostrich futures, then he would make a packet", i.e. she will accept indicative (1) iff she *denies* its subjunctive contrary (2).

This observation holds at later times as well. Learning at t_1 whether or not Jim made the investment in ostrich futures will, for the same reason as before, make no difference to Expert's evaluation of the returns on it. And so at t_1 and t_2 she will accept indicative (3) iff she denies its counterfactual contrary (4).⁵ On the other hand, Expert's evaluation of the returns to the investment will, of course, be sensitive to any information she gains about how ostrich futures have performed. If she knows at t_2 that they have performed badly she will deny (3) and accept (4); if she knows that they have performed well, it will be just the other way around. But in each case her evaluation of the indicative conditional and the corresponding subjunctive will be the same. And in each case her evaluation of both will be independent of whether or not Jim made the investment at t_0 or the degree to which she believes he did.

We reach the same conclusion by looking at the issue from the other direction. Because Jim's investment decision is evidentially relevant for me to the state of the market and the latter is the determinant of how well the various investments perform, for me the probability of whether Jim would make a packet, were he to make an investment in ostrich futures, is sensitive to whether or not Jim will in fact make the investment. In contrast, for Expert, since Jim's decision is evidentially irrelevant to the state of the world, whether he would make a packet were he to invest is probabilistically independent of whether he invests. Now such independence of the Expert's degree of belief in the subjunctive conditional from its antecedent implies that $\Pr(A \rightarrow B) = \Pr(A \rightarrow B|A) = \Pr(B|A)$. But by Modus Ponens, $\Pr(A(A \rightarrow B)) = \Pr(AB)$ and hence $\Pr(A \rightarrow B|A) = \Pr(B|A)$. So Expert's probability for the counterfactual equals the conditional probability of its

⁵ It is true that since she knows whether or not John has made the investment, it is inappropriate, because misleading, for her to utter the indicative. But this is not to say that she does not have an attitude to it.

consequent given the truth of its antecedent, i.e. *Adams' Thesis holds for Expert's degrees of belief in subjunctive conditionals*, as well as for indicative conditionals.

I am now in a position to state my proposal regarding the probabilities of subjunctive conditionals. I claimed earlier that since Expert's beliefs are as good as they can be, we should align our degrees of belief, both conditional and unconditional, with hers. But since we don't know what these are, the best we can do is adopt the degrees of belief that we expect Expert to have. We have now seen that Expert will adopt as her degree of belief in a subjunctive conditional, her conditional degree of belief in its consequent given its antecedent. So we in turn should set our degrees of belief in the conditional to our expectation of the conditional credence of Expert. More formally, let Ex be a random variable taking values from a set of possible probability functions on L measuring the degrees of belief of Expert at time t . Then according to what I will call the Expert Conditional Credence view, the probability of a subjunctive conditional at time t_i is given as follows:

$$(\text{ExConCred}) \Pr_i(A \rightarrow C) = \mathbb{E}_i(Ex(C|A)) = \sum_j \Pr_i(Ex = \pi^j) \cdot \pi^j(C|A)$$

Let's test this proposal against the postulated attitudes to the subjunctive conditionals in our running example. If, at t_0 , the market is such that returns on an investment in ostrich futures will be positive, Expert's conditional probability for Jim making a loss given that he makes an investment will be high; if it such that returns will be negative, it will be low. Since we believe the latter to be true, ExConCred prescribes that we believe counterfactual (2) to a high degree. The same applies at time t_1 : our evidence regarding Expert's conditional probabilities has not changed and so a high degree of belief in counterfactual (4) is required. But by t_2 we know that Expert either knows that Jim will have made a packet, conditional on his having invested in ostrich futures, or that he will have made a loss, either undermining our degree of belief in counterfactual (4) or confirming it. So our proposal correctly predicts the evolution in our attitudes to the subjunctive conditionals.

What about the Morgenbesser's Coin case? Prior to it being tossed, Expert does not know how it will land (because it's an indeterministic process) and so, plausibly, will adopt a degree of belief of one half on a bet on heads winning. Consequently, the Expert Conditional Credence view prescribes degree of belief of one half in the forward-looking subjunctive 'if you were to bet on heads, you would win' (just what Adams' Thesis prescribes for the forward-looking indicative 'If you bet heads, you will win'). Once the coin has landed heads, Expert will update on this information and so will set her conditional degrees of belief in winning, given a bet on heads, to her prior conditional degrees of belief in winning, given a bet on heads and the coin landing heads, which of course equals one. (It matters not that she knows at this point that no such bet has been made.) Consequently, ExConCred prescribes full posterior belief in the subjunctive 'If you had bet on heads, you would have won' (just as the Adams' Thesis prescribes full posterior belief in the corresponding backward-looking indicative 'If you did bet on heads, you won'). So this account gets the Morgenbesser Coin case right as well.

Earlier we saw that the Updated Conditional Chance view proposed by Edgington adequately explained our attitudes to these (Morgenbesser) sentences, but not our acceptance of the subjunctive "If the coin had been tossed it would have landed heads" because the coin landing heads is not independent of it being

tossed. In contrast the Expert Conditional Credence view gets our attitude to this sentence right as well. For Expert's posterior conditional probability for the coin landing heads, conditional on it being tossed, is of course one, since at this point in time she knows how the coin has landed. And in general, the proposed view will correctly prescribe full belief to any subjunctive of the form "If A had been the case, B would have been" in cases in which A and B are both known to be true.

That the Expert Conditional Credence view correctly handles these cases is strong evidence in its favour. Let us now turn to some challenges to it. A first worry that might arise at this point concerns whether the conditional probabilities for the performance of ostrich futures given Jim's decision to invest in them are defined in case Jim does not in fact make the investment. And if they are, in virtue adoption of a suitable definition of conditional probability, whether these probabilities can be determined given that they concern counterfactual possibilities. These worries are misplaced. As we noted earlier, frameworks for conditional probability are available which allow that they be defined for conditions of probability zero. And the relevant conditional probabilities can, in this example at least, be determined without difficulty, even though they are not implied by the unconditional probabilities. The performance of the investment, being causally independent of Jim's decision and evidentially irrelevant to the state of the market, will be inferred by Expert from the state of the market alone. So her conditional expectation at any time for the performance of the investments, given Jim's decision, will equal her unconditional expectation at that time for their performance. Similarly, in the Morgenbesser Coin case, her conditional probability, at any time, for the bet on heads winning, in the event of it being made, will equal her unconditional probability for the coin landing heads, irrespective of whether the bet is or was made.

A second worry. The consequences of the proposed view for the examples we have been looking at are no different from those of Edgington's Updated Conditional Chances view when clause (a) restricting updates to information about events causally independent of the antecedent is removed. So how does the Expert Conditional Credence view handle the case that we used to show why this clause is required? Suppose, as before, that at t_0 Jim can invest in either ostrich futures or housing, but not both. Suppose also that Jim chooses to invest in housing in t_1 , that housing performs well but that ostrich futures do not, and that Jim duly makes a packet at t_2 . Intuitively, in view of the poor performance of ostrich futures, we should deem highly improbable the counterfactual 'If Jim had invested in ostrich futures, he would have made a packet'. Now the Expert Conditional Credence view equates the probability of a counterfactual with (the expectation of) Expert's conditional credence in the consequent, given the truth of its antecedent. But by this point in time Expert knows that in fact Jim has made a packet, albeit from his housing investment not from an investment in ostrich futures. Does this not entail that at t_2 Expert believes to degree one that Jim made a packet, conditional on investing in ostrich futures?

It does not. It is true that Expert's t_0 conditional probabilities for Jim making a packet, given that he invests in ostrich futures, updated by the fact that he makes a packet, must equal one. (This, recall, is why clause (a) is required by the Updated Conditional Chances view: to block the updating by the fact that Jim makes a packet.) But by t_2 , Expert knows that Jim did *not* invest in ostrich futures and so her t_2 conditional credence for Jim making a packet, given that he invests in

ostrich futures, is not required to equal one, despite the fact that she knows that he did make a packet. On the contrary, since Jim's investment decision is evidentially irrelevant for her to its performance and since Jim makes a packet from ostrich futures (or from housing) only if he invests in it and it performs well, her conditional credence at any time for making a packet, conditional on an investment, simply equals her credence in the investment performing well. But since it is known at t_2 that they have performed badly, the latter equals zero. So the counterfactual 'If Jim had invested in ostrich futures, he would have made a packet' too must have probability zero (for us) since we know that Expert knows that ostrich futures have performed badly.

References

- Adams, E.W. 1965, "On the Logic of Conditionals", *Inquiry*, 8, 166-97.
- Adams, E.W. 1970, "Subjunctive and Indicative Conditionals", *Foundations of Language*, 6, 1, 89-94.
- Adams, E.W. 1975, *The Logic of Conditionals: An Application of Probability to Deductive Logic*, Dordrecht: Reidel.
- Adams, E.W. 1976, "Prior Probabilities and Counterfactual Conditionals", in Harper and Hooker 1976, 1-21.
- Bradley, R. 2017, *Decision Theory with a Human Face*, Cambridge: Cambridge University Press.
- Dowe, P. and Noordhof, P. (eds.) 2004, *Cause and Chance: Causation in an Indeterministic World*, London: Routledge.
- Edgington, D. 1995, "On Conditionals", *Mind*, 104, 235-329.
- Edgington, D. 2004, "Counterfactuals and the Benefit of Hindsight", in Dowe and Noordhof 2004, 12-27.
- Edgington, D. 2008, "Counterfactuals", *Proceedings of the Aristotelian Society*, 108, 1, 1-21.
- Fetzer, J.H. (ed.) 1988, *Probability and Causality*, Dordrecht: Springer.
- Harper, W.L. and Hooker, C.A. (eds.) 1976, *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, Vol. I, Dordrecht: Reidel.
- Harper, W.L., Stalnaker, R. and Pearce, G. (eds.) 1981, *Ifs. Conditionals, Belief, Decision, Chance, and Time*, Dordrecht: Reidel.
- Joyce, J.M. 1999, *The Foundations of Causal Decision Theory*, Oxford: Oxford University Press.
- Moss, S. 2013, "Subjective Credences and Semantic Humility", *Philosophy and Phenomenological Research*, 87, 2, 251-78.
- Pearl, J. 2009, *Causality*, New York: Cambridge University Press.
- Ramsey, F.P. 1929, "General Propositions and Causality", in *Philosophical Papers*, ed. by D.H. Mellor, Cambridge: Cambridge University Press, 1990.
- Renyi, A. 1955, "On a New Axiomatic Theory of Probability", *Acta Mathematica Hungarica*, 6, 3, 285-335.
- Slote, M. 1978, "Time in Counterfactuals", *The Philosophical Review*, 87, 3-27.
- Skyrms, B. 1980, *Causal Necessity*, New Haven: Yale University Press.

- Skyrms, B. 1981, "The Prior Propensity Account of Subjunctive Conditionals", in Harper et al. 1981, 259-65.
- Skyrms, B. 1988, "Conditional Chance", in Fetzer 1988, 161-78.
- von Fintel, K. 2011, "Conditionals", in von Heusinger et al. 2011, 1515-38.
- von Heusinger, K., Maienborn C. and Portner, P. (eds.) 2011, *Semantics: An International Handbook of Meaning*, Vol. 2, Berlin: de Gruyter Mouton.
- Williams, J.R. 2008, "Chances, Counterfactuals, and Similarity", *Philosophy and Phenomenological Research*, 77, 2, 385-420.