

STALNAKER'S THESIS IN CONTEXT

ANDREW BACON

Abstract. In this paper I present a precise version of Stalnaker's thesis and show that it is both consistent and predicts our intuitive judgments about the probabilities of conditionals. The thesis states that someone whose total evidence is E should have the same credence in the proposition expressed by 'if A then B ' in a context where E is salient as they have conditional credence in the proposition B expresses given the proposition A expresses in that context. The thesis is formalised rigorously and two models are provided that demonstrate that the new thesis is indeed tenable within a standard possible world semantics based on selection functions. Unlike the Stalnaker–Lewis semantics the selection functions cannot be understood in terms of similarity. A probabilistic account of selection is defended in its place.

I end the paper by suggesting that this approach overcomes some of the objections often levelled at accounts of indicatives based on the notion of similarity.

A popular form of contextualism concerning indicative conditional statements embraces the following three theses:

CONTEXTUALISM: (i) Indicative conditionals semantically express (i.e., can be used to assert the truth of) propositions. (ii) Which proposition is asserted by an utterance of an indicative conditional sentence sometimes depends on the context in which it is uttered. Moreover, (iii) which proposition is asserted depends, absent other sources of context sensitivity, on some piece of evidence or knowledge that is salient in the context of utterance (perhaps the utterer's evidence or some pooled piece of evidence being assumed by the participants of the conversation.)

Contextualism holds a distinguished place in recent philosophy, and for good reason.¹ It promises to answer a number of rather puzzling issues in the philosophy of conditionals – the apparent validity of 'or-to-if' arguments and the so-called 'Gibbardian stand-offs' to name but a couple of examples (see Stalnaker, 1975; Van Rooij, 1999) While it also has its critics² it is interesting to note that, by contrast, the context sensitivity of conditionals is all but taken for granted by linguists working within the framework of Kratzer's (see Kratzer (1986)). This work draws on connections between modals and conditionals, and the context sensitivity of former, at least, appears to be quite pervasive.

My aim here, however, is not to defend contextualism or its applications but to show that contextualism can be put to work to shed light on another difficult issue, namely that of

Received: March 1, 2014.

¹ See, for example, Stalnaker (1975), Van Rooij (1999), Nolan (2003), and Santos (2008).

² There is a long tradition rejecting thesis (i) of the contextualist program: see Adams (1975), Edgington (1986), and Bennett (2003) for representative examples. More recently some theorists have attempted to accommodate the data keeping (i) but rejecting (ii) by adopting a form of *relativism* about the propositions expressed by conditional sentences; see Weatherson (2009) for discussion of this approach.

providing a theory that predicts our intuitive judgments about the probability of conditional statements. As early writers noted, contextualism provides a potential way around the triviality results. However, despite some limited initial optimism regarding this project (see in particular Harper (1976) and Van Fraassen (1976)), few philosophers still see this as a viable option. Most likely this is due to the fact that existing constructions either fall apart when one considers conditionals embedded within other conditionals or are not compatible with an orthodox possible worlds semantics such as a selection function semantics. Finally these constructions invariably require that the locus of context sensitivity be, not a salient piece of evidence in accordance with thesis (iii), but an entire credential state. It is therefore hard to integrate such approaches into an orthodox contextualist framework of the kind that is now popular in linguistics and philosophy. In fact the only construction I know of that overcomes the first limitation, that of accommodating embedded conditionals, is given in van Fraassen's (1976), and this falls afoul of the other two constraints.

In Section 1 of this paper I will present and defend a weakening of a principle known as 'Stalnaker's thesis'. I show that the thesis originally presented by Stalnaker, in addition to being fraught with difficulties arising from the triviality results, is unfriendly to contextualism. I argue that my weakened principle is both strong enough to predict our intuitive judgments about the probabilities of conditionals and compatible with contextualism while also being weak enough to avoid the triviality results. In Section 2, I give the principle a possible world semantics that is intended to integrate straightforwardly with contextualist accounts of indicatives. Unlike the Stalnaker–Lewis semantics the selection functions cannot be understood in terms of similarity. A probabilistic account of selection is defended in its place. The appendices contain two different tenability results establishing the consistency of the new principle with the semantics.

§1. Contextualism and Stalnaker's thesis. Suppose that a card has been picked at random from a standard 52 card deck and placed face down in front of you. Assuming that you are not more confident that some card will be selected over any other, how confident should you be about asserting the following sentences?

1. The selected card is an ace if it's red.
2. It's spades if it's black.
3. It's diamonds if it's an eight.

To be clear, when I ask how confident you should be about asserting a sentence I mean: what degree of belief should you have in the proposition that you would assert by uttering that sentence. The obvious answers to these questions are, in order: $\frac{1}{13}$, $\frac{1}{2}$, and $\frac{1}{4}$. For example, to calculate 2 I just determine what proportion of black cards are spades. Since one in two black cards are spades the answer is $\frac{1}{2}$.

An initially attractive theory, known as 'Stalnaker's Thesis'³, gives us a general way to make these calculations. It states:

³ Not to be confused with Adams' Thesis which employs the notion of 'assertability'. The assertability of a conditional, according to Adams, need not be identified with the probability of a proposition.

STALNAKER'S THESIS: The degree of belief one should assign to a conditional sentence, 'if A then B ', should be identical to one's conditional degree of belief in B given A .⁴

In Stalnaker (1970) Stalnaker works with a theory of probability that assigns degrees of belief and conditional degrees of belief to *sentences* and not, as is normally done, to propositions. In what follows I shall read the thesis as saying that, if p , q , and r are the propositions that would be asserted by the sentences A , B and 'if A then B ' in a given context, then one's degree of belief in r must be identical to one's conditional degree of belief in q given p .

Stalnaker's thesis gets its appeal from its simplicity and its ability to straightforwardly explain the probability judgments reported in 1–3. However, enthusiasm for the thesis quickly died down as a slew of results trivialising this theory appeared (see Bennett, 2003, Chapter 5), for a good summary of the highlights) and Stalnaker himself quickly dropped the theory.

While I find this literature to be conclusive with regard to the thesis (at least as currently interpreted) there is an even more basic reason to be dissatisfied with Stalnaker's theory if you are a contextualist of the stripe described in the opening paragraph. According to the contextualist, there are lots of conditional propositions that one and the same conditional sentence, 'if A then B ', can be used to assert, but only one conditional probability for the probabilities of those propositions to be identical with (provided neither A nor B is themselves context sensitive.)

To spell the worry out in full, suppose that neither A nor B is context sensitive and express p and q respectively in every context. Assume also that I can assert 'if A then B ' in one context and thereby assert the proposition r , and in another context the proposition r' . Since, by assumption, r and r' are two different propositions, there is no general reason why one must assign them equal confidence. Yet according to Stalnaker's thesis, one's degree of belief in both r and r' must be identical to one's conditional degree of belief in p given q . In other words, Stalnaker's thesis entails that one must be equally confident in r and r' after all. For a contextualist this ought to be extremely puzzling. For after all, people are driven to contextualist theories by cases where one hears two utterances of the same sentence in which one seems to be true while the other seems to be false. If Stalnaker's thesis were true then these seemings would be utterly irrational—any two utterances of a conditional sentence must express equally probable propositions, so no rational person could be even somewhat confident in the truth of one and the falsity of the other.⁵

What to make of these problems? One radical response, often made in connection to the triviality results, is to take these highly theoretical arguments to undermine the original probability judgments to 1–3.⁶ To my mind this response is excessive: triviality results do not cast doubt on particular probability judgments such as those reported in 1–3. They merely refute a general theory that predicts those judgments—the judgments themselves

⁴ If Cr is a function representing your degrees of belief, then your conditional degree of belief in B given A , $Cr(B | A)$, is defined to be $\frac{Cr(A \wedge B)}{Cr(A)}$ when $Cr(A) > 0$.

⁵ Indeed the possibility of having unequal degrees of confidence in r and r' is essential if we are to account for the puzzles for which contextualism was introduced to explain (see, in particular, Gibbard's puzzle (Gibbard, 1981).

⁶ I do not mean include those, such as Adams, who rejects probabilities in favour of talk about 'assertibilities'. The people who make the radical response disagree about the numerical values.

do not imply the refuted theory. Furthermore, if the answers I listed to 1–3 are not correct then those who make the radical response owe us an answer to the question: what are the correct answers to these particular questions? If they are not respectively $\frac{1}{13}$, $\frac{1}{2}$, and $\frac{1}{4}$ then what on earth are they?

The contextualist, in my view, has a better response; one that predicts the intuitive probability judgments in 1–3, but does not commit us to Stalnaker's thesis. According to the contextualist the judgments we make about the probabilities of conditional sentences are determined by two pieces of evidence. One piece of evidence determines which proposition is asserted by the utterance of the conditional being evaluated, the other is the evidence you actually possess, which determines what your degrees of belief are if you're rational (i.e. determines the probability function with which we make the actual judgments of probability.) In other words, the former determines which proposition is to be evaluated, and the latter determines how probable that proposition is. I propose that when these two pieces of evidence are identical, the probability of the conditional and the conditional probability coincide – the probability of the proposition you assert with a conditional when E is salient is the same as the conditional probability when your total evidence is E .

The revised thesis entails, for example, that when the utterer's evidence is identical to the contextually salient evidence she will assert a proposition using a conditional sentence that she takes to be exactly as probable as the conditional probability of the proposition expressed by the consequent on the proposition expressed by the antecedent. This is plausibly what is going on when we make the judgments reported in 1–3.

Let us write $A \rightarrow_E B$ for the conditional expressed when evidence E is salient. The informal version of our revised thesis says:

CP $Cr(A \rightarrow_E B) = Cr(B | A)$ provided Cr is a rational credence for an agent whose total evidence is E at world x .

The above instance of the thesis is silent about the credences of agents whose total evidence is not E . That said, whatever your evidence is there will be another conditional corresponding to that evidence and another instance of the thesis which does apply to you.

In order to make this thesis precise two questions must be addressed. Firstly, we must say how the contextually salient piece of evidence determines which proposition is expressed by a conditional sentence. Secondly, we must say what it is for a credence to be rational given a total body of evidence E .

It turns out that the first question can be treated in very different ways; the models constructed in appendices §4 and §5 provide two such treatments. For the time being let us just use the notation $A \rightarrow_E B$ to represent the proposition that would be expressed by the conditional, in a context in which evidence E (representing an accessibility relation) is salient, whose antecedent expresses A and consequent expresses B . In order to address the second question I shall adopt a relatively standard Bayesian picture according to which the epistemic state of a rational agent at a time and world w is represented by a pair consisting of a probability function Pr and an accessibility relation E . Pr represents the agents initial probability function, sometimes called a 'prior' or an 'ur-prior'.⁷ E represents their evidence at t at each possible world by mapping each world x to the agents total evidence at x , $E(x) = \{y \mid Exy\}$. In order to determine what that agent's informed

⁷ Sometimes philosophers use the word 'prior' to represent an agent's credences before they have undergone some episode but in which they are still informed about some matters. This is not how I am using it—by an 'ur-prior' I mean the credences of a completely uninformed agent.

credences are at a time, t and world w , assuming she is rational, we condition the agent's ur-prior on her total evidence at t and w : if Pr is her ur-prior and $E(w)$ her total evidence at w and time t then her credence at t is $Pr(\cdot \mid E(w))$ if she is rational. My conditional credence of B on A at t is therefore just $Pr(B \mid A \cap E(w))$. The revised thesis says that this is just identical to my credence in $A \rightarrow_E B$ at t : $Pr(A \rightarrow_E B \mid E(w))$. Thus the precise statement of the thesis is

CP $Pr(A \rightarrow_E B \mid E(w)) = Pr(B \mid A \cap E(w))$ for every rational ur-prior Pr , evidence E and world w .

Here $E(w)$ is the agents total evidence at possible world w . Simplifying by writing Cr_E for my credences at w with evidence E (i.e. $Cr_E(\cdot) = Pr(\cdot \mid E(w))$) we get the informal thesis mentioned above.

Of course the appeal to ur-priors is controversial and questions about their status are important, however, I shall not delve into those issues here. I expect the thesis to be formulable without them, however the presentation of the principle **CP** in terms of them is particularly simple and will be easy to use in what follows.⁸

1.1. Evidence. The evidence in favour of **CP**, I claim, is exactly the evidence usually adduced in favour of Stalnaker's thesis. Stalnaker's thesis, at least as I have precisified it, is more general: it implies, for example, that if I express a proposition, p , with an utterance of an indicative conditional relative to my evidence, and you were to evaluate this proposition by your evidence you would assign it the conditional probability of the consequent on the antecedent. It is here that the two theories diverge.

To prize apart the two theses, then, we have to consider a case where the evaluator's evidence and the contextually salient evidence are distinct. Thus we want a probability judgment associated with a conditional utterance made in a context other than your own. The problem is this: when you hear an assertive utterance of a conditional your evidence usually changes in such a way that your evidence matches, in the relevant matters, the person who is making the utterance. In these cases the evaluator's evidence and the contextually salient evidence are not so different after all. Thus getting concrete judgments of probability about the propositions expressed by people who have different evidence than you do is a hard task; we shall have to be a bit more indirect than that. At the same time this point is one of the principal virtues of the theory I am proposing—Stalnaker's thesis only seems to be generally true because the cases that appear to confirm it are the special cases in which it is true. The contextually salient evidence and the evaluators evidence are almost always the same, so the cases that disconfirm Stalnaker's thesis are hard to come by.

I'll start by showing that the judgments that motivate Stalnaker's thesis are really instances of **CP**. Then I will try and show, indirectly, that probabilities of conditionals and conditional probabilities come apart when the evaluator's evidence and the contextually salient evidence are different.

Let us begin with a typical example of a probability judgment involving an indicative conditional. Suppose that Alice and Bob know that there was a small chance a given fair coin was flipped earlier today. Alice asks Bob how probable he thinks it is that it landed heads if it was flipped, and he answers that it is a half since all they know is that the coin is fair. Their evidence here is incomplete: neither of them know whether the coin was flipped

⁸ For example, if you can make sense of a credence function being 'rational to have when your total evidence is E ' you can just stick to the informal version of **CP** that does not invoke ur-priors.

or whether it landed heads or tails if it did. On the other hand, they both know that there is a coin, that it's fair and so on. Thus it is this knowledge that determines both which question is asked when Alice utters the question 'how probable do you think it is that the coin landed heads if it was flipped?', and also how likely Bob will find the answer to be. Since the contextually salient evidence and the utterer's evidence are the same, CP delivers the verdict we predicted.

Matters change somewhat when the questioner knows something the questionee doesn't. Let's suppose that both Alice and Bob know that Alice will be informed if the coin is flipped and it lands tails and that otherwise she will not be informed. Then there are two cases to consider. If she is not informed Alice can reason that, although she doesn't know whether the coin was flipped, it's not the case that it was flipped and landed tails. So she truthfully concludes that if the coin was flipped it landed heads. Bob can see this and surmises that:

Case 1: The proposition expressed in Alice's mouth by 'if the coin was flipped it landed heads' is true in cases where she's not informed.

On the other hand, if she is informed that the coin was flipped and landed tails the sentence 'if the coin was flipped it landed heads' is obviously false in her context.

Case 2: The proposition expressed in Alice's mouth by 'if the coin was flipped it landed heads' is false when she is informed.

These are, at least, the truth values these two utterances would have in the respective scenarios if they had truth values at all.⁹ Furthermore, according to the contextualist framework outlined earlier, the proposition expressed by Alice will be the same in both scenarios. Even though the proposition that constitutes her evidence at *t* is different in the two contexts, the salient accessibility relation, which tells us what her evidence at *t* is in each case, is the same. So according to a popular contextualist theory, which relativises the proposition expressed to an accessibility relation (sometimes called a modal base), rather than a salient proposition, Alice will have said the same thing in both cases. This point is crucial because it means that Bob can know what Alice has said, even if he does not know what she actually knows – he knows that in the first type of scenario Alice's evidence at *t* will include the fact that she hasn't been informed, and that in the second type of scenario the evidence that the coin was flipped and landed tails.¹⁰ Call the proposition expressed in both cases *p*.

⁹ There seems to be a straightforward analogy with the Gibbard cases that are sometimes taken to motivate contextualism. The crucial difference is that these two utterances are made in different worlds, relative to the same agent's evidence whereas the Gibbard cases the utterances are made in the same world relative to different agent's evidence. Thus unlike Gibbard cases we have no trouble accounting for the different truth values of these utterances (they are made in different worlds where the facts are different) and moreover, since it is Alice's evidence at the time of utterance, *t*, that is salient in both scenarios, it is natural to think that the two utterances express the same proposition.

¹⁰ It is a common misconception about contextualism that it involves pervasive ignorance about what the speaker is saying (see, for example, Gibbard, 1981, p. 232–234; Stalnaker, 1987, p. 110–111). There is an alternative contextualist view that is susceptible to this charge. The alternative view uses a contextually salient proposition, rather than a contextually salient accessibility relation, to determine what is said. On this view what Alice said in the two scenarios was distinct. Moreover, a natural contextualist account (a variant of HARPER'S CONDITION below) implies that when the contextually supplied evidence entails $A \supset B$ the corresponding indicative expresses a necessary proposition, and when it entails $A \wedge \neg B$ the indicative expresses an impossible proposition.

Here it seems as though when Alice outright asserts 'if the coin was flipped it landed heads', she is reporting something about her knowledge, not about what is common knowledge between them: since Bob doesn't have this information, she communicates something useful by uttering this conditional.¹¹ Unfortunately, for this very reason, we cannot get clean intuitions about what Bob's prior credence is in the proposition Alice asserted, because what he will report is his credence after he's heard the utterance and updated on what's been said. However we can still ask him to consider, before he's heard anything, how probable it is that an utterance of that sentence would be true, were it to be made by Alice.

Here is how he reasons. Either Alice was informed that the coin was flipped and landed tails or she wasn't. If she's not informed (i.e., Case 1 above) then, p , the proposition she would have asserted if she'd uttered the conditional, is true. If she was informed that the coin was flipped and landed tails (i.e., Case 2) p is false. Thus he knows that p is true if and only if she was not informed, and this happens if and only if the coin was not flipped or it was flipped and landed tails. Thus the probability of what the conditional says in Alice's context is just the probability that either it wasn't flipped or it was flipped and landed tails (i.e. the material conditional.) Since there was only a small chance that it was flipped, Bob's credence in this disjunction is high, and not identical to one half. Stalnaker's thesis predicts that Bob's credence in p should be a half, CP does not.

The above is an attempt to get a direct probability judgment regarding a proposition expressed by a conditional relative to a context that has different information to that of the judge. Such examples are hard to come by, and are certainly harder to evaluate. Let me now try and show that Stalnaker's thesis fails by calculating these probabilities indirectly.

Suppose that both Alice and Bob know that a car, with some unknown amount of gas, is to be driven in a straight line until it runs out of gas.¹² While they do not know how much gas is in the car, they know that it will run anywhere between 0 and 100 miles and then stop. Both Alice and Bob begin with evenly distributed credences regarding how far the car travelled – suppose they in fact have exactly the same credences, represented by the function Pr . Alice then goes out and checks the last 30 miles of the road. The car is not there and she concludes that the car didn't travel more than 70 miles – her credences are now represented by the function $Pr(\cdot \mid E_{\leq 70})$, where $E_{\leq 70}$ is the proposition that the car travelled at most 70 miles. Bob does nothing and his credences remain the same at Pr .

Thus we get the puzzling consequence that in the first case Alice said something necessarily true, and in the second case she said something necessarily false.

This alternative view is described, for example, in Van Rooij (1999), although van Rooij finds an alternative way to make sense of communication in this setting. The idea, by analogy, is that if you hear Fred utter the sentence 'I am hungry', but you do not see who made the utterance, you do not know what has been said (that Fred is hungry) but you may still update your beliefs on a proposition determined by the Kaplanian character and conclude that the speaker of the context is hungry. Thus even when there is ignorance about what has been literally said communication is still possible.

¹¹ That we use the speaker's knowledge, and not the mutual knowledge, in these cases is also crucial for solving the Sly Pete cases in Gibbard's (1981). Note that there may still be some cases where two speakers have unequal evidence, yet it is only the mutual knowledge of the participants that determines what is said. This kind of flexibility regarding what evidence to use is widely acknowledged (see, for example, the discussion in Chapter 4 of Kratzer (2012).) This phenomenon is also present with epistemic modals.

¹² I take this example from Edgington (2014), although she uses it for different purposes.

Now Alice and Bob both consider the following conditional in their respective contexts:

If the car went at least 50 miles it went exactly 60 miles.

Since there is more evidence available in Alice's context than in Bob's it follows, given CONTEXTUALISM, that the propositions expressed by Alice and Bob are (at least potentially) different. Call the propositions expressed by Alice and Bob A and B respectively. If we were to ask Alice and Bob to report their credences in the above conditional it seems clear that Alice would report her credence in A and Bob would report his credence in B , and not vice versa. CP predicts the intuitively correct result that these two verbal reports will match Alice and Bob's respective conditional credences: when Alice evaluates A the evaluator's evidence and the contextually salient evidence are the same, and similarly when Bob evaluates B .

Let us run through this idea explicitly for Bob's judgment. Bob doesn't know how far the car went: for him there are 50 possibilities where the car went at least 50 miles, and in only one of them did the car go 60 miles. Intuitively he judges the probability of the above conditional to be $\frac{1}{50}$. Since ' B ' was the name introduced for whatever proposition Bob is judging in his context, Bob's judgment corresponds to:

$$1. Pr(B) = \frac{1}{50}$$

Not coincidentally $\frac{1}{50}$ is also his conditional credence, and this is exactly what CP predicts. Analogous things can be said about Alice's credence in A .

Stalnaker's thesis also predicts this. However Stalnaker's thesis, as I've precisified it, predicts much more—it also predicts that Alice's credence in B (not just in A) should match her conditional credences, and that Bob's credence in A (not just in B) should match his conditional credences. This is not predicted by CP since when Alice evaluates B the contextually salient evidence is Bob's evidence not Alice's, and similarly for when Bob evaluates A . Neither of these extra predictions of Stalnaker's thesis are substantiated by a verbal report, and moreover, I shall now argue, the predictions are incorrect. I shall argue that if Bob's credence in B is his conditional credence of $\frac{1}{50}$, as suggested by his verbal reports, then Alice's credence in B isn't her conditional credence. (A parallel problem could be raised for the prediction that both Alice and Bob's credence in A must be their conditional credences.)

Since Alice knows the car went at most 70 miles, her conditional credence that the car went 60 given it went at least 50 is $\frac{1}{20}$. Thus if Alice's credence in B were her conditional credence then:

$$2. Pr(B | E_{\leq 70}) = \frac{1}{20} \text{ (Alice's credence in } B \text{ is her conditional credence.)}$$

Note that 2 is *not* motivated by an intuitive judgment in the same way that 1 is, since it is A , not B , that Alice would evaluate if she were to consider the above conditional.

Here is the problem: 2 must be false if 1 is true, for they are jointly inconsistent. Unluckily, Stalnaker's thesis predicts both. Note that by 2 $Pr(B | E_{\leq 70}) = \frac{1}{20}$. Also since Bob's credences were uniform his credence that the car went at most 70 miles is $\frac{7}{10}$, so $Pr(E_{\leq 70}) = \frac{7}{10}$. Now by probability theory we have $Pr(B) = Pr(B | E_{\leq 70})Pr(E_{\leq 70}) + Pr(B | \neg E_{\leq 70})Pr(\neg E_{\leq 70})$. Since the second summand is positive, and the first we have calculated we know that $Pr(B) > \frac{1}{20} \cdot \frac{7}{10}$. Yet $\frac{1}{20} \cdot \frac{7}{10} > \frac{1}{50}$ contradicting 1.

Unfortunately for Stalnaker's thesis it predicts both 1 and 2 which we have shown to be jointly inconsistent. Yet it is clear that it is only 1 that corresponds to our intuitive judgments about probabilities; CP predicts 1 but not 2 and thus fits the bill perfectly.

1.2. Other approaches. It is worth noting that despite the impossibility results there has been a number of attempts to resurrect Stalnaker's thesis in some limited form (see the two constructions in Van Fraassen, 1976, McGee, 1989, Jeffrey & Edgington, 1991, Jeffrey & Stalnaker, 1994, Kaufmann, 2009, Bradley, 2011.)

To simplify discussion, we can divide these attempts into two classes: those that place some restriction on the kinds of sentences for which a variant of Stalnaker's thesis holds, and those that don't. In fact all but one of the listed approaches falls into the former class, leaving only the first construction given in van Fraassen's (1976) providing us with an unrestricted version of Stalnaker's thesis.

One thing to highlight about the present results, that distinguishes them from the results in the former class, is that CP applies to all propositions *A* and *B*, without any restriction on what kinds of iterations of conditionals are allowed. A common sticking point for the former proposals is to account for the probability of nested conditionals, especially conditionals with conditionals in the antecedent place in a way that is consistent with the connection to conditional probability.¹³ While iterated conditionals are not as commonplace in ordinary discourse, it is just as important to account for them. For one thing, it is not clear that a syntactic restriction can rule out the instances of CP that these theorists find problematic. Consider (1) and (2):

If it breaks without significant deformation if it is subjected to stress, it (1)
is not a suitable material

If it is brittle it is not a suitable material (2)

Firstly note that (1), while an iterated conditional, is a perfectly reasonable thing to say – iterated conditionals are not a mere curiosity but a proper part of English. Secondly, even if (1) were improper (it is certainly harder to parse) the antecedents of (1) and (2) plausibly express the same proposition and (2) is certainly not improper; indeed (2) is a simple conditional. It therefore doesn't seem plausible that a purely syntactic distinction, such as the distinction between nested and simple conditionals, could carve out a significant epistemological distinction; after all (1) and (2) fall on different sides of the distinction yet plausibly they are semantically, and presumably epistemologically, of the same kind.

A more direct argument can also be given for including iterated conditionals within the scope of CP. Consider the following scenario:

Suppose you have ten numbered vases, three are shatter-proof and the remaining seven are fragile enough to break if dropped. You also know that two of the fragile vases are priceless, however, you don't know which of the vases are priceless or fragile. Suppose also that there has recently been an earthquake and there is a chance that some of the vases have fallen from their shelves onto the floor.

¹³ McGee, Jeffrey, and Stalnaker allow compounding in the consequent but not in the antecedent (e.g. Jeffrey writes 'Like McGee's treatment, this one allows compounding in the consequent ('If A, then if B then C') but not in the antecedent', see also Stalnaker and Jeffrey's 'Generalized Adam's Thesis' which restricts attention to conditionals with categorical antecedents.) Van Fraassen's second construction (the 'Stalnaker–Bernoulli model' in §4) allows for conditional antecedents and consequents provided these conditionals do not themselves contain conditional antecedents and consequents; this issue is inherited in Kaufman's approach. Bradley expresses optimism regarding the prospects of extending his approach to iterated conditionals, although whether this is possible remains an open question. Such approaches typically argue that such conditionals aren't evaluable anyway, or at least, that they shouldn't be evaluated by conditional probabilities.

How confident should you be that vase number eight is priceless if it was one of the vases that broke if it was dropped? The intuitive answer is calculated as follows: there are seven vases that will break if dropped. Furthermore, we know that out of those only two are priceless, so the proportion of priceless vases out of those that broke if they were dropped is intuitively $\frac{2}{7}$.

The proposal that comes closest to this one, then, is the first construction found in van Fraassen's paper (see §3.) Given a probability measure with certain nice properties, this construction will produce a model in which Stalnaker's thesis holds for arbitrary conditionals, including those with conditionals embedded arbitrarily deep in the consequent and antecedent. Unlike CP, however, this construction does not explicitly provide an account of the dynamics of conditional belief: if you give it your credence function at time t , the construction will output a model of the conditional that satisfies Stalnaker's thesis relative to this single probability function, but it doesn't explicitly account for what happens as you learn new things and your credences change. That said, there is a natural way to incorporate van Fraassen's theory into a contextualist theory resembling CP, and that is to simply rerun the construction on the updated credence to provide a new interpretation of the conditional. On this construal the connective that is expressed by an indicative sentence in a context depends on a probability function supplied by the context (presumably the speaker's credences.)

It is worth comparing this idea to the now prominent version of contextualism in linguistics according to which the proposition expressed by a conditional sentence depends on a contextually supplied 'modal base'.¹⁴ For our purposes this can simply be modeled by an accessibility relation, which maps each world to a set of accessible worlds representing (something like) the total evidence available in the utterance situation at that world. In one regard, the modal base contains far less information than a probability function—given a world, the modal base merely tells you which possibilities are left open by the evidence at that world, and says nothing about how probable those possibilities are. There does not appear to be an independently motivated reason to think that two conditional utterances, made when the same epistemic possibilities are open, could express different propositions due to a small difference in how probable these possibilities are. This seems like a fairly radical form of context sensitivity, whereas the dependence on a modal base is much more modest and independently evidenced.

While CP integrates neatly with this kind of theory, van Fraassen's construction doesn't. Firstly, the modal base has a far reaching conversational role in those theories that extends well beyond the contribution they make to conditional assertions, and this role is perfectly well captured using a modal base and not a probability function. Secondly, the modal base is sometimes supposed to represent the pooled evidence of the conversational participants; to apply this to credences would require solving the problem of credence aggregation (a difficult problem – see Russell et al. (Forthcoming).)

There is another way in which van Fraassen's construction involves accepting more context sensitivity than contemporary contextualism does. Consider again the example in which Alice is informed at t if a coin was flipped and landed tails, but is otherwise not informed. Her credences at t in the two cases will be different, so according to van Fraassen

¹⁴ There are, of course, lots of variants and different terminologies, but the basic point I am making remains unchanged in these variants.

she will end up saying distinct things by the conditional 'if the coin was flipped it landed heads' depending on her credences. Thus in van Fraassen's model Bob will not know what Alice has said with the conditional since he does not know whether she was informed or not. This point holds even if we assume, implausibly, that Bob knows exactly what Alice's credences would have been in two different scenarios. A better model would treat the locus of context sensitivity as a function from worlds to probability functions, intuitively mapping each world to Alice's credences at that world. On this model Bob would know what Alice had said even if he didn't know whether she'd been informed (assuming he knows exactly what her credences would be like in the two scenarios).

The formalism I have adopted is more like the latter model, except that I am using something less fine grained than a function from worlds to probability functions. A modal base is effectively a function mapping each scenario to a proposition—Alice's evidence at t in that scenario. In both scenarios the contextually salient modal base will be the same and Alice will assert the same proposition by uttering a conditional sentence at t in either scenario – informally the context sensitivity isn't due to *which* proposition actually constitutes her evidence at t , it's due to her evidence at t whatever it might be. For Bob to work out what has been asserted by a conditional utterance at time t all he needs is the function that maps each world to the utterer's evidence at t at that world. Thus, presumably, all he needs to know is who is speaking at what time in that context – he does not need to know what the speaker's evidence is.

Other issues are more specific to van Fraassen's construction. The interpretation of a conditional, on his approach, is generated by assigning conditionals subsets of the unit interval (or any space that is 'full' – see Van Fraassen, 1976, for definitions) that have the right size as semantic values. Accordingly the semantics is highly nonstandard and it is consequently not consistent with orthodox possible world accounts of conditionals (it cannot, for example, be represented by a selection function, as van Fraassen notes.) One puzzling aspect of the semantics is that conditionals of the form $A \rightarrow B_i$ can all be true, for a consistent A and a countable collection of propositions B_i , even when the B_i are jointly inconsistent. This is reminiscent of an objection to Lewis's semantics for counterfactuals, which predicts that, for each $\epsilon > 0$, if I had been taller than 2 meters, I'd have been strictly between 2 and $2 + \epsilon$ meters (given that in the actual world I'm less than 2 meters tall, which I am!) Since it is incoherent to suppose that someone's height is strictly between 2 and $2 + \epsilon$ meters for every $\epsilon > 0$, one might think that it is impossible for me to be taller than 2 meters, but this is clearly not an impossibility (see Herzberger (1979), and also Fine (2012) who makes this problem quite vivid.¹⁵)

§2. The random worlds semantics for indicatives. Our first order of business is to be a bit more precise about what counts as a model for CP. Here I will be concerned with outlining a natural logic of conditionals, and providing a selection function semantics for it. Those interested in the philosophical interpretation of this semantics can skip to the next subsection.

¹⁵ Another difference between my approach and van Fraassen's is that his logic, CE, is slightly weaker than mine. For example, the unary operator defined by the conditional $\neg A \rightarrow A$, expressing a kind of epistemic necessity, cannot be proved to be a normal modal operator in CE. More importantly, van Fraassen's construction validates modus ponens in the initial context, but when the agent updates her evidence modus ponens can end up failing at some worlds when the construction is rerun.

We want a class of connectives, \rightarrow_E , that not only supports CP but also has a reasonable conditional logic. Of particular interest is the connective \rightarrow_E obtained where the relevant evidence is tautologous – i.e. when E relates every world to every other world. I'll call this the 'ur-conditional'. When E is tautologous I shall omit the subscript altogether and I'll simply write $A \rightarrow B$.

We shall work within a modal propositional language, \mathcal{L} , consisting of the usual truth functional connectives, \neg and \supset , from which the other truth functional connectives are definable, and a special binary modal connective representing the ur-conditional, \rightarrow . I shall adopt the ordinary definitions of \wedge , \vee , \perp in terms of \supset and \neg . I shall also adopt the following shorthands:

$$\begin{aligned} A \equiv B &:= (A \supset B) \wedge (B \supset A) \\ A \leftrightarrow B &:= (A \rightarrow B) \wedge (B \rightarrow A) \\ \Box A &:= (\neg A \rightarrow \perp) \end{aligned}$$

To increase readability, as is typically done in probability theory, I shall frequently shorten $A \wedge B$ to AB .

My focus will be on the theory which I'll call L . L can be axiomatised by closing the following axioms under modus ponens (for the material conditional \supset), the rule of uniform substitution and the rules RCN and RCEA.

- RCN if $\vdash B$ then $\vdash A \rightarrow B$
RCEA if $\vdash A \equiv B$ then $\vdash (A \rightarrow C) \equiv (B \rightarrow C)$
CK $(A \rightarrow (B \supset C)) \supset ((A \rightarrow B) \supset (A \rightarrow C))$
ID $A \rightarrow A$
MP $(A \rightarrow B) \supset (A \supset B)$
CEM $(A \rightarrow B) \vee (A \rightarrow \neg B)$
C1. $(A \rightarrow B) \supset ((B \rightarrow \perp) \supset (A \rightarrow \perp))$

The first three principles correspond to a basic conditional logic, entitled CK (usually context will distinguish the logic from the principle CK.) This logic is common to pretty much all possible world approaches to conditionals and is in this sense analogous to the weakest normal modal logic K (indeed CK ensures that the unary modal operator $A \rightarrow$ is normal operator in Kripke's sense.) RCN states that conditionals whose consequents are logical truths are themselves logical truths. In conjunction with CK this ensures that what's true 'if A ' is closed under classical propositional logic (specifically, CK ensures it's closed under modus ponens for the material conditional.) Finally RCEA ensures that logically equivalent sentences can be substituted in the antecedent position. From this the intersubstitutivity of logical equivalents in any position is derivable in CK.

The next two principles should be fairly self explanatory. ID just says that if A then A . MP, on the other hand, says that indicative conditionals entail the corresponding material conditional. This is tantamount to saying that \rightarrow obeys modus ponens for the only ways for $A \rightarrow B$ to be true while $A \supset B$ to be false would be for $A \rightarrow B$ and A to be true and B false. Philosophers skeptical of modus ponens, or indeed any of the other principles, can still take interest in the tenability results. If CP is consistent with the logic L , it is certainly consistent with the weakenings of L .

Of particular note are the final two axioms, CEM and C1. The axiom CEM, short for 'conditional excluded middle', is distinctive to Stalnaker's logic of conditionals, and constitutes the primary difference between it and a similar theory due to Lewis (1973).

C1 on the other hand governs the behaviour of conditionals that are vacuously true. When we are concerned with indicatives a conditional is vacuously true when the antecedent is epistemically impossible in the relevant sense. The only cases in which $A \rightarrow \perp$ is true are cases in which the conditional is vacuously true; in these cases I'll say that A 'crashes'. Since we are focusing on the ur-conditional the only proposition ruled out by your evidence is the contradictory proposition. Thus when \rightarrow represents the ur-conditional several further principles are motivated such as C0

$$\text{C0 } (A \rightarrow \perp) \supset ((B \rightarrow C) \equiv (A \vee B \rightarrow C))$$

If A crashes, then A is inconsistent so $A \vee B$ and B ought to be equivalent and thus ought to conditionally imply the same propositions. (It is not entirely clear to me whether C0 is valid when the ur-conditional is replaced by \rightarrow_E for arbitrary evidence E so I leave that open in what follows.) It is worth noting that the system that results from replacing C1 with C0 in L has C1 as a derived theorem. C0 is therefore strictly stronger than C1.

Furthermore, if the only proposition that crashes is the inconsistent proposition then we should expect the defined operator \Box (' $\neg A$ crashes') to iterate in accordance with the modal logic S5. In particular we want:¹⁶

$$4 \ (A \rightarrow \perp) \supset (B \rightarrow (A \rightarrow \perp))$$

$$B \ A \supset (A \rightarrow \perp) \rightarrow \perp$$

Neither of these principles are motivated when \rightarrow is substituted for conditionals expressed by agents with evidence. If we were to define a \Box_E operator as $\neg A \rightarrow_E \perp$, this would express some kind of epistemic necessity which may not iterate in the way predicted by 4 and B.

Natural analogies between L and Stalnaker's logic C2 can be drawn. The most salient difference is that this logic does not have the theorem

$$\text{CSO } (\phi \leftrightarrow \psi) \supset ((\phi \rightarrow \chi) \supset (\psi \rightarrow \chi))$$

Indeed adding CSO to L collapses the logic into Stalnaker's, so in this sense we can think of L as what you get by removing CSO from C2. In my view this is a benefit of the present account: CSO has been subjected to a number of counterexamples (see Tichý, 1978 (and the variant discussed by Stalnaker in Stalnaker (1987)), Mårtensson, 1999, Tooley, 2002, and Ahmed, 2011) and is also responsible for some of the triviality results (see Stalnaker, 1976; Hájek & Hall, 1994). However this is not the venue for a full defence of this feature of the logic so I shall put it to one side for now.

A *frame* for a conditional logic is a pair $\langle W, f \rangle$ where W is a set of worlds and $f : \mathcal{P}(W) \times W \rightarrow \mathcal{P}(W) - f$ is called the 'selection function'. A model is a pair $\langle \mathcal{F}, \llbracket \cdot \rrbracket \rangle$ where \mathcal{F} is a frame and $\llbracket \cdot \rrbracket$ maps propositional letters to subsets of W . $\llbracket \cdot \rrbracket$ extends to a function from the rest of \mathcal{L} to $\mathcal{P}(W)$ as follows:

- $\llbracket \neg \phi \rrbracket = W \setminus \llbracket \phi \rrbracket$
- $\llbracket \phi \supset \psi \rrbracket = (W \setminus \llbracket \phi \rrbracket) \cup \llbracket \psi \rrbracket$
- $\llbracket \phi \rightarrow \psi \rrbracket = \{w \mid f(\llbracket \phi \rrbracket, w) \subseteq \llbracket \psi \rrbracket\}$

¹⁶ The particular formulations of these principles are due to Cian Dorr. Given our definition of \Box they are provably equivalent to the principles $\Box A \rightarrow \Box \Box A$ and $A \rightarrow \neg \Box \neg \Box A$ respectively.

A sentence, ϕ , is true in a model $\langle W, f, \llbracket \cdot \rrbracket \rangle$ iff $\llbracket \phi \rrbracket = W$, and is valid on a frame iff it's true in every model based on that frame, and valid on a class of frames iff it is valid on every member of that class.

RCEA, RCN, and CK are valid on the class of all frames. Combinations of the remaining principles are validated on the class of frames that additionally satisfy the corresponding combination of conditions from the below list:

- ID $f(A, x) \subseteq A$
- MP $x \in f(A, x)$ whenever $x \in A$.
- CEM $|f(x)| \leq 1$
- C1. If $f(A, x) \subseteq B$ and $f(B, x) = \emptyset$, $f(A, x) = \emptyset$
- C0. If $f(A, x) = \emptyset$ then $f(A \cup B, x) = f(B, x)$

In the presence of CEM the selection function always either picks out a singleton or the empty set. In this case we can modify the semantics to conform with Stalnaker's original Stalnaker (1968) semantics so that f maps us from a world, w , and a set of worlds, A , to a single possible world (namely x if $f(A, w) = \{x\}$ in the general semantics) or the unique impossible world $\#$ (if $f(A, w) = \emptyset$ in the general semantics) at which every sentence is stipulated to be true. In certain circumstances it will be useful to translate between Stalnaker's semantics and Chellas's slightly more general semantics which allows for more than one world to be selected.

If we want to guarantee 4 and B as well, one can stipulate that

$$f(A, x) = \emptyset \text{ only if } A = \emptyset.$$

This of course encodes the principle that A crashes only if it's the inconsistent proposition. This condition automatically ensures C0 and (thus) C1.

So much for the ur-conditional. What of the conditionals \rightarrow_E when E represents substantial evidence? We shall use f_E to represent the selection function for this conditional where E is an accessibility relation corresponding to some evidence. Finally, given a world x , $E(x)$ will be used to represent the set $\{y \mid E x y\}$ (the function $x \mapsto E(x)$ from worlds to propositions is sometimes called a 'modal base'.) A very natural thought would be simply to define f_E in terms of E and the ur-selection function as follows:

$$f_E(A, x) = f(A \cap E(x), x)$$

This has the effect of guaranteeing that the truth value of an indicative conditional, 'if A then B ', in a context is a function of the epistemically possible A -worlds in that context (where the epistemically possible worlds are just those consistent with the contextually salient evidence E .) It is worth noting that defining $A \rightarrow_E B$ this way preserves all of the axioms of L except, possibly, for MP. If we further stipulate that E is reflexive – as indeed it probably should be given it represents knowledge or mutual knowledge – then MP holds.

As we shall see, only the first of the two models makes the above identification. This concludes our discussion of the constraints on the selection function. The following definition will be useful in what follows:

DEFINITION 2.1. *Given a frame $\langle W, f \rangle$, say that the selection function is regular if it satisfies the frame conditions for the logic L .*

A selection function is normal if it is regular and $f(A, x) = \emptyset$ iff $A = \emptyset$. Frames based on normal selection functions validate B, 4, and C0 in addition to the principles of L .

2.1. Conditional excluded middle. How should one understand the above semantics, and in particular, what does the selection function intuitively represent? The interpretation initially given to the selection function by Stalnaker in Stalnaker (1968) was that $f(A, x)$ picks out (the singleton of) the closest world to x in which A is true, where closeness is determined by some measure of similarity between worlds. This interpretation initially attracted a lot of criticism. For one thing, it requires that there be a unique closest A -world to x when the relevant notion of closeness seems to determine no such thing—any ordinary ordering of similarity would allow for ties or infinite descending chains of ever closer worlds. Another issue is that it is not clear what the relevant notion of closeness is when we are trying to evaluate indicative, as opposed to subjunctive, conditionals. Many have the intuition that indicative conditionals are in some sense epistemic and conclude that, since the notion of closeness relevant for evaluating subjunctives is irrelevant here and no epistemic notion is forthcoming, indicatives should not be analysed in terms of closeness.

Lewis's response to the first objection—that there might not be a unique closest world—in the case of subjunctives, is to relax the constraint that the selection function pick out a unique world. In terms of the constraints listed above this means relaxing the constraint that $|f(A, x)| \leq 1$. Accordingly, $f(A, x)$ must be allowed to pick out a set of closest worlds without any assumption that there must be at most one of these.¹⁷ Unfortunately this has the knock on effect of invalidating CEM.¹⁸ Lewis was primarily concerned with subjunctive conditionals, and subjunctive instances of CEM are often controversial for good reason. When we are concerned with simple past tense indicative sentences, however, CEM appears to be much harder to deny. Contrast:

1. Either the coin would land heads if it were flipped or it would land tails.
2. Either the coin landed heads if it was flipped or it landed tails.

While the former is disputable, the latter surely isn't (assuming we are not taking seriously the possibility that the coin could do anything other than land heads or land tails.¹⁹) Of course, Lewis himself does not apply his own brand of 'closest world' style semantics to indicative conditionals – my point is just that there are very good reasons not to relax the condition that $|f(A, x)| \leq 1$ in the case of simple past indicatives.

Much has been said on this, and I do not want to adjudicate between the various responses Stalnaker and others have put forth in favour of this interpretation. I will say one thing, however. One question we have been considering concerns whether there is always a unique closest A -world, and indeed whether it is even appropriate to use the notion of 'closeness' in the semantics of indicative conditionals. Another very different question asks whether CEM is valid for past tense indicatives. There is no reason to think that

¹⁷ To properly represent Lewis's semantics we'd have to go beyond the simple selection function semantics described in this section, since Lewis's semantics allows for failures of the limit assumption.

¹⁸ If $f(\llbracket A \rrbracket, x) = \{y, z\}$, $y \neq z$ and $\llbracket B \rrbracket = \{y\}$ then x belongs to neither $\llbracket A \rightarrow B \rrbracket$ nor $\llbracket A \rightarrow \neg B \rrbracket$.

¹⁹ Some indicatives don't behave like this: indicatives with 'will' in the consequent are known to behave a lot more like subjunctive conditionals. I don't want to include habitual indicatives either – sentences phrased in the simple present such as: 'if the window is left open, Granny jumps out'. Even the instance ordinary excluded middle 'Granny jumps out of the window or Granny doesn't jump out of the window' sounds nontrivial, and this is probably because we read it as saying that either Granny usually jumps out or she usually doesn't. Thus sentences in the simple present can't straightforwardly be taken to represent counterexamples to excluded or conditional excluded middle.

an answer one way or the other to the first question should determine our answer to the second, especially if indicative conditionals are not to be analysed in terms of closeness. There is consequently no reason why indicatives cannot be modeled using a selection function semantics that validates CEM, provided the selection function is not analysed in terms of closeness. (Indeed, you might take the fact that a similarity based semantics ties an intuitively correct principle to an implausible principle about similarity is a powerful argument against this kind of semantics.)

How to interpret the selection function then? Assume, with Stalnaker, that $|f(A, x)| \leq 1$. Then a more neutral way of putting things would be as follows:

There are a bunch of indices which describe the different possible way things are for all you know, one of which describes the way things actually are: x . $f(A, x)$ then simply represents the way things are if A .²⁰

Of course, for Stalnaker, the world that would have obtained if A had obtained just *is* the closest world at which A obtains (and similarly for indicatives.) But this identification is not forced on us, and one can still say everything we want to say about the semantics of conditionals by interpreting the selection function in the more neutral way.

A potentially illuminating way to think of the selection function is as picking out an antecedent world at random from the epistemically accessible worlds.²¹ In Stalnaker's theory a world is selected from the accessible antecedent worlds with an overriding preference for more similar worlds. On my understanding, however, the selection process has no preference for more similar worlds: we can think of it as having a preference for worlds that are more probable on the evidence, but this preference is not overriding but proportional to the probability.

The idea of randomly selecting a world is clearly a metaphor and not intended to provide a reductive analysis of conditionality. There are clearly many ways to select something randomly. You could picturesquely imagine God rolling a die to determine which world to select. This is not what I mean, the process of random selection is irreducibly conditional in nature. One way to randomly select a member of the set {Heads, Tails} would be to take a coin out of your pocket and flip it or spin it. Another way would be to leave the coin in your pocket and instead talk about the side that landed face up *if* it was flipped at t ; this will in some sense pick out one side at random. Here the process of random selection is partially determined by the antecedent (it's going to be a flipping rather than a spinning of the coin, for example) but the conditional morphology was essential to describing the process.

2.2. The triviality results. There are, of course, numerous triviality results affecting variants of Stalnaker's thesis. These typically come in two flavours: dynamic and static. Dynamic triviality results leave it open whether there could be a rational agent whose credences in conditionals always match their conditional credences. What they show is that if your credences are matched in this way, this will be disrupted upon changing your credences to accommodate new evidence. Static results, on the other hand, show that no rational agent could be in that kind of state in the first place.

Note that CP is designed to integrate straightforwardly with a standard Bayesian theory according to which one always updates one's beliefs by conditionalisation. The dynamic

²⁰ For subjunctive conditionals we can say that $f(A, x)$ represents the way things would have gone (at x) had A obtained.

²¹ Moritz Schulz (Forthcoming) defends something like this interpretation of counterfactuals.

triviality results, however, fail to get a hold in our setting. CP predicts that if my total evidence is E my credence in $A \rightarrow_E B$ must be my conditional credence – however, the thesis is simply silent about my credences in $A \rightarrow_E B$ once my evidence has changed from E to something stronger such as E^+ . (Although, of course, there will be another instance of CP for the conditional $A \rightarrow_{E^+} B$.)

More troubling are the static triviality results, for these purport to show that one cannot ever satisfy the conditional probability equation. One class of static triviality results rely on principles distinctive to Stalnaker's logic. For example, in Hájek & Hall (1994), it is shown that the principle CSO mentioned above causes trouble with Stalnaker's thesis if we take it to govern conditionals that contain conditionals embedded in certain ways within the antecedent.

The validity of CSO corresponds to the following constraint on selection functions:

CSO If $f(A, x) \subseteq B$ and $f(B, x) \subseteq A$ then $f(A, x) = f(B, x)$

This validates the principle CSO which would have the effect of collapsing the logic L into Stalnaker's logic C2. CSO is guaranteed on a similarity based semantics: if the closest A -world is a B -world and the closest B -world is an A -world then the closest A -world is the closest B -world. On the random world interpretation of the selection function, however, no such constraint exists and CSO is invalid. The randomly selected A world might be a B world and vice versa, but there is no guarantee that the very same world will be selected except in the special case where there is only one accessible world at which both A and B are true. Thus these static triviality results hold no sway for the present account of conditionals.

Related principles of conditional logic also give rise to static triviality results. Indeed all of the principles below can be shown to cause problems analogous to the one that CSO poses (RCA, for example, is shown to be problematic in Edgington (1995)):

- CA $((\phi \rightarrow \chi) \wedge (\psi \rightarrow \chi)) \supset (\phi \vee \psi \rightarrow \chi)$
- RCA $(\phi \vee \psi \rightarrow \chi) \supset (\phi \rightarrow \chi) \vee (\psi \rightarrow \chi)$
- CM $(\phi \rightarrow \psi) \supset ((\phi \rightarrow \chi) \supset (\phi \wedge \psi \rightarrow \chi))$
- RT $(\phi \rightarrow \psi) \supset ((\phi \wedge \psi \rightarrow \chi) \supset (\phi \rightarrow \chi))$

These principles are all closely related to CSO and are validated in the similarity semantics. Indeed, given my preferred logic L , each of the above principles is provably equivalent to CSO except for RCA which is equivalent if you assume C0.²² Unsurprisingly they are all invalid in the random world semantics.

Despite the fact that these principles are strictly speaking invalid, it is worth pointing out that they enjoy a kind of pragmatic validity. For if an agent's evidence is contextually salient when the conditionals are uttered the agent will find the conclusions probable if the premises are sufficiently probable – a fact that is a straightforward consequence of CP.²³

Before we move on it is also worth noting that there are existing results that get the above logical principles and a restricted version of Stalnaker's thesis at the same time by restricting the thesis to simple conditionals in which certain iterations of conditionals in the antecedent are banned (see Van Fraassen (1976).) If one finds restrictions like this

²² The proofs of these equivalences are in Bacon (Unpublished Manuscript).

²³ That these principles are probabilistically valid was proved in Adams' (1975). An inference is valid in this sense if, roughly, however, probable you want to make the conclusion you can find a threshold such that if the agent finds the premises to be at least that probable she will find the conclusion at least as probable as the amount you wanted.

at all attractive (I do not) it is also worth noting that the results in appendix §4 show a symmetrical result: that one can have Stalnaker's thesis in full generality—i.e., with no such restrictions – with a restricted form of the above inferences instead. That is to say, one can have **CSO** and its relatives provided we restrict the sentences occurring in antecedent position of the above inferences, ϕ and ψ , to sentences that do not themselves contain conditionals.

There is another class of static triviality results that I must address. These results do not rest on principles of conditional logic, but rather show that the thesis is not satisfiable in models in which there are only finitely many worlds (Hájek, 1989) or even in models in which there are countably many worlds (Hall, 1994). In short, to satisfy **CP** one needs uncountably many worlds.

Let us begin with Hájek's result. This can be demonstrated with a fairly simple example: imagine that we just want to model the roll of a dice whose outcome we are ignorant about. Intuitively you might think that we could model this with exactly six equiprobable worlds representing each possible outcome. The problem with this is that if you want a thesis like **CP** then you need to find, for each conditional probability, a proposition with that probability. However, in a finite model there simply won't be enough propositions to go around. In the model described above, for example, every proposition has a probability of the form $\frac{n}{6}$, where n is simply the number of worlds in that proposition. However, the probability that the dice lands on a 6 given it doesn't land on a 1 is $\frac{1}{5}$ which clearly is not of the form $\frac{n}{6}$ so there is no proposition with that probability.

Note, however, that the assumption that there are only 6 epistemic possibilities in the scenario described above becomes utterly implausible once we take conditional propositions seriously. Let us consider the world in which, unbeknownst to me, the die landed on a 1. I claim that in this scenario I am not only ignorant about the outcome of the die roll, but also ignorant about some conditional facts, such as whether the dice landed on a 6 or whether it landed on a 5 if it landed on one of 5 or 6. Given **CEM** we know that even at the world where the die in fact landed on a 1, the die either landed on a 5 if it landed on a 5 or a 6 or it landed on a 6. Thus strictly speaking the 1 world should be split into two epistemic possibilities corresponding to the possibility that $D_1 \wedge ((D_5 \vee D_6) \rightarrow D_5)$ and the possibility that $D_1 \wedge ((D_5 \vee D_6) \rightarrow D_6)$ where D_n is the proposition that the die landed on n .

By considering other conditionals with antecedents that are false at the world where the die lands 1 you can argue that this world should be divided into further epistemic possibilities. Moreover, once you have recognised the existence of these further epistemic possibilities, there are new propositions you can plug in as antecedents corresponding to arbitrary sets of these epistemic possibilities. Some of these sets of epistemic possibilities do not correspond to categorical (i.e., nonconditional) propositions, so intuitively this is like considering the epistemic possibilities generated from conditionals with conditional antecedents. You can make the argument rigorous if you wish, but it should, I hope, be clear that the presence of conditional propositions ensures that in the case described the number of epistemic possibilities is infinite.

Why must the number of epistemic possibilities be uncountable? This follows from a fairly natural extension of the previous remarks. For if there are infinitely many epistemic possibilities, then there are uncountably many sets of these possibilities—that is to say, there are uncountably many propositions. Thus there are uncountably many antecedents to play around with—for each of uncountably many propositions, A , we are ignorant about what is true if A . Thus there are uncountably many things we are ignorant about.

Of course, there is another way of understanding what a 'world' is: a maximally strong categorical proposition – something which tells us the answers to ordinary questions like how dice land, but are silent about the hypothetical facts about what happened if this or that. It would be extremely puzzling if there was some argument that demonstrated that **CP** could not be satisfied in a model with finitely or countably many worlds in this sense of 'world'. Fortunately there can be no such argument – the construction in appendix §4, for example, allows you to start off with a set of worlds of any size (representing maximally strong categorical facts) and will then construct a model of **CP** in which these worlds are split into further epistemic possibilities corresponding to all the unknown conditional facts.

2.3. The tenability results. So far we have just been concerned with the interpretation of the conditional. In order to model **CP** we also need to talk about probabilities and evidence. In particular we need to enrich the frames with a class of probability functions representing the ur-priors, and a set of propositions which represent the propositions that could, in some possible world, be some agent's total evidence.

The following definition provides us with a precise framework against which we can evaluate the truth of **CP**:

DEFINITION 2.2. *A probability frame is a tuple $\langle W, \mathcal{B}, f, \Sigma, P, w \rangle$ where*

- *W is a set of worlds, where $w \in W$ represents the actual world.*
- *\mathcal{B} is a complete Boolean algebra of subsets of W , containing W , representing the evidence propositions.*
- *f maps accessibility relations to selection functions. Given E with $E(w) \in \mathcal{B}$, f_E is a regular selection function on E .*
- *Σ is a σ -algebra (a set of subsets of W containing \emptyset and closed under complements in W and countable unions.)*
- *P is a nonempty set of countably additive probability measures over Σ representing the set of rational ur-priors.*

Informally, a probability frame provides us with a set of probability measures, over a measure space $\langle W, \Sigma \rangle$, and also a collection of selection functions, f_E , relative to the same set of worlds, W , indexed by accessibility relations E . The final ingredient is a Boolean algebra of propositions, \mathcal{B} , that represent the propositions that could be one's total evidence (we can then impose the restriction $E(w)$ is a member of \mathcal{B} for each world w .) I leave it open that any proposition could be a persons total evidence; however, there are other natural constraints on evidence one might consider (see appendix A.)

We are looking for a probability frame that satisfies **CP**; such frames will be called *adequate*:

DEFINITION 2.3. *A probability frame $\langle W, \mathcal{B}, f, \Sigma, P \rangle$ is adequate if and only if*

$$Pr(\mathcal{B} \mid A \cap E(w)) = Pr(\{x \mid f_E(A, x) \subseteq \mathcal{B}\} \mid E(w)) \text{ for every } Pr \in P, \\ A, \mathcal{B} \in \Sigma \text{ and } E \text{ with } E(w) \in \mathcal{B} \cap \Sigma$$

It will often be useful to write $A \Rightarrow_E B$ instead of $\{x \mid f_E(A, x) \subseteq B\}$. We also adopt the convention of dropping the subscript when E is the vacuous evidence.

However, in all of the results I prove, I restrict attention to accessibility relations, E , that are introspective at the actual world. This just means that if Ewx , $E(w) = E(x)$.

This condition would be ensured, for example, if E was an equivalence relation. However, this seems like too strong a condition: if your knowledge, for example, can be represented by an equivalence relation then you are not only perfectly introspective, but necessarily perfectly introspective.

Our goal, then, is to construct an adequate frame. However, there are other conditions we might also want to explore. For example

NORMALITY: A frame is normal iff the ur-selection function f is a normal selection function.

This of course shows that a stronger logic than L is compatible with CP. Both the models we will consider shortly are normal. Another constraint we might want to implement is:

FULLNESS: A frame is full iff $\mathcal{B} = \mathcal{P}(W)$.

A principled reason to weaken the fullness condition would be if you thought that only categorical propositions (i.e., unconditional propositions) could be an agent's total evidence. In a full frame every proposition, categorical, or hypothetical, could in principle be an agent's total evidence. We shall return to the question of whether a conditional proposition could be an agent's total evidence in the next section. For now let me just highlight it as a possible further constraint in addition to adequacy. Finally we might also want the constraint:

HARPER'S CONDITION: $f_E(A, x) = f(A \cap E(x), x)$ where f is the ur-selection function.

The condition above is stated in Van Rooij (1999), who attributes the idea to Harper in Harper (1976).²⁴ The basic thought behind Harper's condition is this: whether an utterance of a conditional is true (at a world x) should be a function of the epistemically possible A worlds at that context (the set of epistemically accessible A -worlds is just AE .) Harper's condition is quite strong, and Stalnaker has suggested the following weaker condition (I have reformulated it from Stalnaker (2009) to match our current conventions.)

STALNAKER'S CONDITION: $f_E(A, x) \subseteq E(x)$.²⁵

Stalnaker's condition does not require that the truth of a conditional 'if A then B ' must depend only on the epistemically possible A -worlds – it is compatible that there be two distinct contexts providing evidence E and E' such that the E -accessible A worlds and the E' accessible A worlds coincide, but where $A \rightarrow_E B$ is true and $A \rightarrow_{E'} B$ is false.

We are now in a position to state the relevant theorems.

THEOREM 2.4. *There is an adequate frame that satisfies NORMALITY and HARPER'S CONDITION (and therefore also STALNAKER'S CONDITION).*

²⁴ Actually van Rooij states his theory in terms of a contextually salient *proposition* rather than an accessibility relation. He states that $f_E(A, x) = f_{\emptyset}(AE, x)$ when $AE \neq \emptyset$. When $AE = \emptyset$, however, van Rooij stipulates that $f_E(A, x) = f_{\emptyset}(A, x)$. On this interpretation $A \rightarrow_E \perp$ can be false even if A is inconsistent with E : this has the effect of making \rightarrow_E satisfy MP even at the worlds inconsistent with E .

²⁵ Stalnaker also qualifies this with the condition that A and $E(x)$ are consistent. In the models considered here the stronger thing stated above also holds so I have left out the qualification.

THEOREM 2.5. *There is an adequate frame that satisfies NORMALITY, FULLNESS and STALNAKER'S CONDITION.*

See the appendices for the proofs.²⁶

§3. Conclusion. In summary, then, we have proposed a thesis, **CP**, connecting the probabilities of conditionals to conditional probabilities that predicts the instances of Stalnaker's thesis that seems intuitively right without predicting the instances that are intuitively incorrect; this thesis extends to account for intuitions about iterated conditionals. Moreover, the proposal is specifically designed to address the dynamics of belief in a way that is consistent with a standard Bayesian theory of updating via conditionalisation.

It is also worth noting that while the resulting theory admits a possible worlds semantics based on selection functions, the theory is not compatible with one prominent interpretation of the selection function based on the idea that the selected world be an antecedent world which is, in some sense, minimally different from the actual world. On the similarity interpretation the resulting logic would be slightly stronger than my own, for it would include the principle **CSO**. Let me end the paper by making a few remarks about this.

The reading of the selection function I proposed was that $f_E(A, x)$ selects an epistemically accessible A -world at random in a way that may or may not select the closest accessible A world to x . The process by which the world is 'selected at random', however, cannot be given an explication in nonconditional terms: it cannot be understood more informatively than just the (E -accessible) world that describes how things are if A , which is random only in the sense that we don't know which world this is (when what we know is given by E).²⁷

One might object to this proposal on the grounds that, unlike the similarity account, we do not get a reductive, or even an informative account of conditionals out of the analysis. Of course this is no objection to someone who never set out to give a reductive analysis, but more importantly, it is not clear that the similarity analysis enjoys this apparent advantage either. Stalnaker (Stalnaker, 1987, pp. 126–132) explicitly states that the pretheoretic notion of similarity plays no role in fixing the truth conditions of conditional statements. Much of the motivation for the abstract analysis in terms of orderings with certain constraints is to provide rationale for formal properties on the selection function that validate desirable modes of inference such as **CSO**. The ordering relevant for evaluating conditionals is therefore not an antecedently understood notion of similarity, but one specifically guided by a pre-existing understanding of conditionality.²⁸

Insofar as the similarity analysis is motivated by the desirability of principles such as **CSO**, Stalnaker's attitude toward this principle is surprisingly noncommittal. He writes, for

²⁶ One might wonder if it's possible for an adequate frame to be both full and satisfy Harper's condition together. The answer to this is no: see Korzukhin's 'Triviality Results' <https://courses.cit.cornell.edu/tk283/Triviality.pdf> (unpublished).

²⁷ Schulz (Forthcoming) argues for something like this way of understanding the selection function in the case of counterfactuals, although I read him as taking the notion of 'random selection' to give us a more informative grip on conditionals than I do.

²⁸ Another worry you might have is that to many people it seems clear that some notion of similarity is important for evaluating subjunctive conditionals, and that our analysis of indicatives belies the parallels between the two cases. My view is that if similarity does make its way into the analysis of subjunctive conditionals, it is through the distinctive behaviour of modals like 'will' and 'would' that appear in these constructions. A natural view would be that the counterfactual selection function selects an A world at random from among the most similar A worlds. This also has the side effect of validating **CEM** without imposing implausible constraints on similarity.

example, that ‘the arguments for condition (3) [i.e. **CSO**] are far from decisive’ (ch7 ft5) and, after considering an apparent counterexample to it, suggests that ‘the question is one of how to distribute the burden of explanation between pragmatics and semantics’ and that ‘to some extent the issue may be one of simplicity and efficiency of formulation rather than substance’. Stalnaker also points toward many good inferences involving **CSO** and related principles. (I find these examples less than convincing, however, since alternative accounts of the goodness of these inferences are available. **CP** already predicts that anyone who is certain in $A \leftrightarrow B$ and in $A \rightarrow C$ in at a given context must be certain in $B \rightarrow C$. Moreover, while the axiom **CSO**, a compound of conditionals, might not always be fully probable in a context, there is a threshold below which its probability cannot fall, according to **CP**, so that one can never be in a context where one can outright assert its negation.)

An example we discussed earlier brings out the difference between Stalnaker’s gloss and mine quite nicely (this point is due to Edgington (2014).) In the example where Bob knows that a car has travelled between 0 and 100 miles but does not know how far, it seems natural to say that he also doesn’t know whether the car travelled, say, 63 miles if it passed the 50 mile mark, or 89. Indeed it’s natural to think that he takes it to be just as probable that it went 63 miles if it passed the 50 mile mark as that it went any other number of miles between 50 and 100. This fits my interpretation fairly naturally, where the world that describes how far the car went is picked randomly from the worlds where it travelled between 50 and 100 miles, with no preference for one world over any other (and in particular, no preference for more similar worlds.) A similarity analysis would rather suggest that at any world at which the car actually went less than 50 miles, the conditional ‘the car went exactly 50 miles if it went at least 50 miles’ would be true. Since it went less than 50 miles in half the accessible worlds, this makes it much more probable that it went 50 exactly miles if it went at least 50 miles, than that it went any other number of miles.

Although neither account can give an analysis of the selection function reductively in terms of similarity or random selection, the present proposal has an advantage over the similarity analysis. For I can say something substantial about the conceptual role of the ur-selection function that a similarity theorist cannot. While, of course, we can both say something about the logical role of conditionals, this will not rule out a material analysis since that satisfies all of the logical principles we have mentioned in this paper. However, the ur-selection function has a distinctive role in thought which neither the material analysis nor the similarity analysis can accommodate.²⁹ Supposing that my present epistemic state at world x is represented by a prior Pr and an epistemic accessibility relation E . Then the ur-selection function is subject to satisfying the following rational constraint: $Pr(\{y \mid f_E(A, y) \in B\} \mid E(x)) = Pr(B \mid A \cap E(x))$, where f_E is defined from the ur-selection function (this could be filled out using Harper’s condition, or some other way.) Writing $Cr(\cdot)$ for $Pr(\cdot \mid E(x))$ (my present rational credences, relative to evidence E) and $f_E^{-1}(A, B)$ for $\{y \mid f_E(A, y) \in B\}$ (my ‘personal’ conditional, which plays a special role in my thinking) this just simplifies to:

$$Cr(f_E^{-1}(A, B)) = Cr(B \mid A).$$

Of course, this is just a restatement of **CP**. Note, however, that one can perfectly well ask whether there is a selection function that plays this distinctive role in thought without taking any stance about how it relates to the semantics of indicative sentences in any

²⁹ In the case of the similarity analysis, it is exactly the validity of the principle **CSO** that prevents it from satisfying this role.

language.³⁰ The idea is that we first get a grip on the selection function this way, and then we ask whether this is clearer than Stalnaker's abstract orderings, and whether it is more suited to play a role in a theory of conditional language.

If this is right then it opens up the possibility of providing a theory of indicatives that is based on probabilistic considerations rather than similarity. These are no more than programmatic remarks, of course, yet I hope that the above results will open up an avenue of research (outlined in Stalnaker (1970)) that has been long abandoned.

§4. Appendix A: Tenability result with Harper's condition. In this section we will construct a model for CP. Several things that are worth noting about this model

1. All of the selection functions are determined by the ur-selection function in accordance with Harper's condition: $f_E(A, x) = f(A \cap E(x), x)$.
2. The ur-selection functions is normal so that $f(A, x) = \emptyset$ only if $A = \emptyset$, thus it satisfies **B**, **4** and **C0**.
3. The model is not full: the set of evidence propositions, \mathcal{B} is a strict subset of the set of all proposition $\mathcal{P}(W)$. Here \mathcal{B} intuitively represents the categorical (nonconditional) propositions.
4. P is a fairly rich set of probability functions. In fact, for every probability function Pr over the nonconditional propositions, \mathcal{B} , there is a unique probability function in P whose restriction to \mathcal{B} is Pr .

The third point is particularly worthy of note. According to the informal gloss, \mathcal{B} represents the set of propositions that could, in some possible world, be an agent's total evidence. Since some propositions do not belong to \mathcal{B} according to this model it follows that there are some propositions that could not be an agent's total evidence. Fortunately there is an intuitive interpretation of this feature of the model.

In the model we start off with an initial set of objects, which we can think of as possible worlds, and the propositions in \mathcal{B} can be identified with arbitrary sets of these worlds. We can think of a possible world as determining all the ordinary facts concerning where objects are located, and so on and so forth, but *not* the conditional facts. For example a possible world might determine that a particular coin, C , isn't flipped on a particular occasion, but it won't determine whether the coin will land heads or tails if it is flipped at that occasion. Thus there will be two epistemic possibilities, corresponding to the same worldly facts (i.e. the same possible world), and according to one the coin lands heads if it is flipped, and according to the other it lands tails if flipped. In general, then, \mathcal{B} represents nonhypothetical/non-conditional propositions and can be represented by sets of possible worlds whilst the full set of propositions, including hypothetical propositions, and can be represented by sets of epistemically possible worlds.

Why then, couldn't an arbitrary hypothetical proposition, say the proposition that the coin C will land heads if it's flipped, be an agent's total evidence? A common observation for views accepting CEM is that conditionals like these give rise to a curious epistemic phenomenon: in this case it doesn't seem to be possible to find out whether the coin will land heads if flipped when the coin is never flipped. For example, if you accept conditional excluded middle then either C will land heads if it is flipped, or it will land tails, but in worlds where the coin is not flipped it is impossible to obtain further evidence to settle

³⁰ One does not, for example, need to be a contextualist to raise this question.

the question of which way it would land if flipped. Philosophers subscribing to the law of conditional excluded middle have conjectured that hypothetical propositions like this are a special source of indeterminacy (e.g., Stalnaker, 1981) Whether or not this is so we can certainly agree that we must be ignorant in the scenario described, much as we would be in the face of vagueness or indeterminacy. The basic intuition is that one can have any credence you like regarding the completely determinate nonhypothetical facts, but once you have fixed your credences in those propositions your credences over the rest of the space of propositions is fixed. For example, if you know that C is fair and will not be flipped, then you are forced to have a credence of a half in the proposition that C will land heads if flipped. The situation here is similar to the analogous situation with vague propositions. Once you know someone has a certain borderline number of hairs, N , you are forced to be uncertain, to some degree, in the proposition that that person is bald.

Of course, we learn conditionals all the time; it is important to keep in mind that this fact is completely consistent with the thesis that our total evidence is never conditional. According to the logic L , when one learns that A and B then one learns the conditional stating that if A then B , and when you learn that A and $\neg B$ you can rule out the conditional if A then B . But in these cases your total evidence (AB and $A \wedge \neg B$ respectively) is strictly stronger than the conditional facts you've learnt. In other cases we know conditionals even when we are ignorant about the antecedent and consequent. Even when you do not know whether the fuse will blow or the light will go off, it is quite reasonable to assert that if the fuse blows the light will go off. But in these cases it is natural to think that your assertion is only appropriate when you know a stronger strict conditional (say, that in all nomically possible worlds in which the fuse blows the light goes off.) When you do not know the strict conditional, such as in the case of the coin flip, it is not appropriate to assert the indicative conditional, even if it is in fact true.

4.1. The construction. The following construction uses the ideas developed by van Fraassen's in his 'Bernoulli–Stalnaker' models from Van Fraassen (1976). However, van Fraassen's models do not satisfy the principle CP for two reasons. Firstly, there is only one conditional connective that satisfies a variant of the conditional to conditional probability link, whereas CP states something much more general (that some form of the link holds for each conditional connective you can express in some context or other.) Secondly in van Fraassen's model the probability conditional to conditional probability link holds only for special conditionals and does not extend to iterated conditionals of various sorts. The following construction is, in a loose sense, the result of iterating van Fraassen's construction ω_1 many times.

The construction begins with an initial set of possible worlds, W , which intuitively can be thought of as representing maximally specific things that can be said about the world without mentioning conditional facts (i.e., facts about what will happen if this or that happens.) The set W_∞ then extends this set, dividing members of W into epistemic possibilities according to the kind of hypothetical distinctions you can make. Epistemic possibilities can be thought of as ordered pairs of ordinary worlds and sequences of worlds, with the latter encoding all the conditional facts that hold at that epistemic possibility. An ordered pair of a world and a sequence is isomorphic another sequence with an extra initial element—thus epistemic possibilities will just be represented as sequences of possible worlds.

Let us put this into practice. Assume that the initial set of states, W , that do not involve conditional facts is given and is countable. The set of worlds in our model will be the set

$W_\infty = W^{\omega_1} = \{\pi \mid \pi : \omega_1 \rightarrow W\}$. We shall set $\mathcal{B} = \{A \times W_\infty \mid A \subseteq W\}$. \mathcal{B} is isomorphic to $\mathcal{P}(W)$ and is thus a complete Boolean algebra. It is easy to see that \mathcal{B} embeds into the larger algebra of all propositions, which we shall denote $\mathcal{B}_\infty = \mathcal{P}(W_\infty)$.

Given our initial space W , define the following sequence of sets for $\alpha < \omega_1$

- $W_\alpha = W^{\omega^\alpha}$

That is, W_α represents the set of all ω^α sequences of members of W . Since $\omega^\alpha < \omega_1$ whenever $\alpha < \omega_1$ it follows that an element of W_α will be isomorphic to an initial segment of a member of W_∞ . Note also the following consequences of this definition:

- $W_0 = W$
- $W_{\alpha+1} \cong W_\alpha^\omega = W_\alpha \times W_\alpha \times W_\alpha \times \dots$

In what follows we shall adopt a practice of identifying products which are isomorphic to subsets of W_∞ , allowing us, for example, to identify $A \times W_\infty$ with a subset of W_∞ whenever A is contained in some W_α .

The sets W_α for $\alpha < \omega_1$ help us describe the measurable sets.

DEFINITION 4.1. *Suppose X is a set of subsets of W_∞ . Then $cl(X)$ is the closure of X under the operations of countable unions and intersections, and complements relative to W_∞ .*

The measurable sets, which we shall denote Σ_∞ , can be thought of as being approximated by an infinite sequence of σ -algebras, $\Sigma_\alpha \subset \mathcal{B}_\infty$ for $\alpha < \omega_1$.

- $\Sigma_0 = \{A \times W_\infty \mid A \subseteq W\}$
- $\Sigma_{\alpha+1} = cl\{A_0 \times \dots \times A_n \times W_\infty \mid A_i \times W_\infty \in \Sigma_\alpha \text{ for } 0 \leq i \leq n\}$
- $\Sigma_\gamma = cl(\bigcup_{\alpha < \gamma} \Sigma_\alpha)$

Note that $\Sigma_{\alpha+1}$ is generated by sets of the form $A_0 \times \dots \times A_n \times W_\infty$ where $A_i \subseteq W_\alpha$. Each of these generating sets consists of an ω_1 sequence such that an initial finite number of elements belong to W_α and the rest belong to W . This is, of course, just equivalent to an ω_1 sequence of elements of W whenever $\alpha < \omega_1$: it is just equivalent to n successive ω^α -sequences of elements of W followed by an ω_1 -sequence of elements of W , which is itself an ω_1 -sequence of elements of W . Bearing this equivalence in mind we can see from the construction that an arbitrary member of Σ_α will be of the form $A \times W_\infty$ where $A \subseteq W_\alpha$. It is straightforward to show

PROPOSITION 4.2. $\Sigma_\alpha \subseteq \Sigma_\beta$ if $\alpha \leq \beta$

Now we turn to our definition of Σ_∞ , the set of measurable sets.

DEFINITION 4.3. *A set $A \in \mathcal{B}_\infty$ is measurable iff $A \in \Sigma_\alpha$ for some α . We denote the set of measurable sets $\Sigma_\infty := \bigcup_{\alpha < \omega_1} \Sigma_\alpha$.*

It should now become apparent why we chose the ordinal ω_1 in our definitions: it is due to this choice that our measurable sets are closed under countable unions so that Σ_∞ is a σ -algebra.

DEFINITION 4.4. *If A is measurable then the rank of A is the smallest α such that $A \in \Sigma_\alpha$. We shall write this: $rank(A) = \alpha$. If A is not measurable then $rank(A) = \infty$.*

It is now time to define the ur-selection function for a $A \in \mathcal{B}_\infty$ of rank α (possibly identical to ∞). If A is nonempty let τ_A be any member of A (it doesn't matter which.)

$$f(A, \pi) = \begin{cases} \pi[\omega^\alpha .i] & \text{where } i \text{ is the smallest number such that } \pi[\omega^\alpha .i] \in A \\ \tau_A & \text{if there is no such number and } A \neq \emptyset \\ \# & A = \emptyset \end{cases}$$

Here $\pi[\alpha] : \omega_1 \rightarrow W$ is given by the function $\pi[\alpha](\beta) = \pi(\alpha + \beta)$ (i.e. $\pi[\alpha]$ is just the ω_1 sequence you get by lopping off the first α members of π .) Note that this is formally reminiscent of Stalnaker’s semantics: $f(A, \pi)$ represents the closest world to π which belongs to A , where closeness depends on how small the ‘ i ’ is—the crucial difference is that the notion of closeness at play here depends on the rank of the antecedent, A . It is easy to verify that f is normal. In particular, the second condition – that $f(A, \pi) = \tau$ if A is nonempty and there is no A -world in the sequence of $\pi[\omega^\alpha .i]$ ’s—is to ensure that $f(A, \pi)$ does not output the impossible world unless $A = \emptyset$ (if we were to replace τ_A with $\#$ in the definition we get a merely regular selection function.) In order to obtain f_E we simply identify $f_E(A, x)$ with $f(A \cap E(x), x)$ in accordance with Harper’s condition.

PROPOSITION 4.5. \emptyset is measurable and if A, B and A_0, A_1, A_2, \dots are measurable then so is $W_\infty \setminus A, A \Rightarrow B$ and $\bigcup_n A_n$.

We now define the set, P , of ur-priors. For simplicity we have assumed that W is countable so that every subset of W can be treated as a measurable set (although it would be simple enough to drop this assumption and work with an initial σ -algebra over W instead.) We shall show that every regular countably additive probability function Pr on the powerset algebra on W extends to the measurable sets over \mathcal{B}_∞ . We then identify P with the set of all such probability functions generated this way.

Suppose that Pr is a regular countably additive probability function on \mathcal{B} . For $\alpha \leq \omega_1$ we define Pr_α over Σ_α as follows.

- $Pr_0 = Pr$
- $Pr_{\alpha+1}(A_0 \times \dots \times A_n \times W_\infty) = Pr_\alpha(A_0 \times W_\infty) \dots Pr_\alpha(A_n \times W_\infty)$; $Pr_{\alpha+1}$ extends to the rest of $\Sigma_{\alpha+1}$ via Carathéodory’s extension theorem.
- $Pr_\gamma(A) = Pr_\alpha(A)$ when $A \in \Sigma_\alpha$ for $\alpha < \gamma$. This extends to the rest of Σ_γ by Carathéodory’s extension theorem.

Write Pr_∞ for Pr_{ω_1} . Observe, from the construction of Pr_∞ , that for any $\alpha < \omega_1$ and $A_0, \dots, A_k \subset W_\alpha$ $Pr_\infty(A_0 \times \dots \times A_k \times W_\infty) = Pr_\infty(A_0 \times W_\infty)Pr_\infty(A_1 \times W_\infty) \dots Pr_\infty(A_k \times W_\infty)$

We are now in a position to prove our main theorem.

THEOREM 4.6. *The frame $\langle W_\infty, \mathcal{B}, f, \Sigma_\infty, P \rangle$ is adequate.*

In particular, if Pr is a countably additive regular probability function over W then $Pr_\infty \in P$ and $Pr_\infty(A \Rightarrow_E B \mid E(\tau)) = Pr_\infty(B \mid A \cap E(\tau))$ whenever $Pr(A) > 0, E(\pi) \in \mathcal{B}$ for all π and A, B and $E(\pi)$ are measurable.

Proof. We begin by showing the result for the ur-selection function. Suppose that $rank(A) = \alpha$ so that $A = A' \times W_\infty$ for some $A' \subseteq W_\alpha$.

Assume that $Pr_\infty(A) > 0$. Thus $A \neq \emptyset$ so according to our definition $\pi \in A \Rightarrow B$ if and only if the smallest A world in the sequence $(\pi[\omega^\alpha .i])_i$ is a B world or there are no A -worlds in this sequence and τ_A is a B world. In other words, if and only if $\pi[\omega^\alpha .0] = \pi \in A \cap B$ or $\pi \notin A$ but $\pi[\omega^\alpha .1] \in A \cap B$ or $\pi[\omega^\alpha .0] \notin A, \pi[\omega^\alpha .1] \notin A$ and $\pi[\omega^\alpha .2] \in A \cap B$ or \dots or $\pi[\omega^\alpha .i] \notin A$ for any i and $\tau_A \in B$.

Let R be the set of π with $f(A, \pi) = \tau_A$. Thus $A \Rightarrow B = (A \cap B) \cup (\overline{A'} \times (A \cap B)) \cup (\overline{A'} \times \overline{A'} \times (A \cap B)) \cup \dots \cup R = \bigcup_n (\overline{A'}^n \times (A \cap B)) \cup R$. Here I am using \overline{X} to denote the complement of X .

Note that $R \subseteq (\overline{A'})^\omega \times W_\infty$ which has probability 0 whenever $Pr_\infty(A) > 0$. Since we are calculating a union of disjoint sets we have

$$\begin{aligned} Pr_\infty(A \Rightarrow B) &= \sum_{n < \omega} (Pr_\infty(\overline{A'})^n \cdot Pr_\infty(A \cap B)) = \frac{Pr_\infty(A \cap B)}{1 - Pr_\infty(\overline{A'})} = \frac{Pr_\infty(A \cap B)}{Pr_\infty(A)} \\ &= Pr_\infty(B \mid A) \quad \square \end{aligned}$$

The above demonstrates the result for the ur-conditional: $Pr_\infty(A \Rightarrow B) = Pr_\infty(B \mid A)$. It remains to show that $Pr_\infty(A \Rightarrow_E B \mid E(\tau)) = Pr_\infty(B \mid A \cap E(\tau))$ for accessibility relations E . In what follows I restrict attention to accessibility relations, E , and worlds τ such that $E(\tau) = E(\pi)$ whenever $E\tau\pi$. When this holds say that E is locally an equivalence relation at τ . Intuitively these correspond to worlds where the evidence concerning what the evidence is complete at the world w (for example, if the salient evidence is just my current knowledge, then this means that I know what I do and don't know.) The restriction to these cases is purely an idealization – I do not think it is a general fact that the evidence available in a context always behaves like this. Whether these idealizations can be relaxed is a question I shall leave to future work.

We shall begin by showing how to write $(A \Rightarrow_E B) \cap E(\tau)$ as a disjunction of disjoint sets as in the previous proof. Suppose that E is locally an equivalence relation at τ and that $E(\tau) \in \mathcal{B}$. For short let us write X for $E(\tau)$ and suppose that the rank of A is α . Recall that since X is a member of \mathcal{B} , $X = X' \times W_\infty$ for some $X' \subseteq W_\alpha$.

LEMMA 4.7. *Given the above definitions, $(A \Rightarrow_E B) \cap X = \bigcup_n (\overline{A'X'}^n \times ABX) \cap X$*

Proof. Suppose $\pi \in (\overline{A'X'}^n \times ABX) \cap X$ for some n . That means that $\pi \in X$, $\pi[\omega^\alpha.n] \in ABX$ and $\pi[\omega^\alpha.m] \notin A$ and $\notin X$ for $m < n$.

Since $\pi[\omega^\alpha.n] \in X = E(\tau)$ this means that n is the smallest number such that $\pi[\omega^\alpha.n]$ is both a member of A and $E(\tau)$: i.e., $f(A \cap E(\tau), \pi) = \pi[\omega^\alpha.n]$. Since $E\tau\pi$, $E(\tau) = E(\pi)$ so $f(A \cap E(\pi), \pi) = \pi[\omega^\alpha.n]$. Moreover, since $\pi[\omega^\alpha.n]$ is B world, $\pi \in A \Rightarrow_E B$. Since, by assumption, $\pi \in X$ this shows one inclusion.

Now suppose that $\pi \in (A \Rightarrow_E B) \cap X$. Suppose that $f(A \cap E(\pi), \pi) = \pi[\omega^\alpha.n]$. Since $\pi \in X = E(\tau)$ we know that $E\tau\pi$, and thus that $E(\pi) = X = E(\tau)$. $f(A \cap E(\pi), \pi) = f(A \cap X, \pi) \in ABX$ which means that $\pi \in \overline{A'X'}^n \times ABX$ (just as in the last theorem.) Thus $\pi \in \bigcup_n (\overline{A'X'}^n \times ABX) \cap X$ completing the proof. \square

We shall now demonstrate that CP holds in this model.

THEOREM 4.8. *If $X \in \Sigma_0$ and $Pr \in P$ then $Pr(A \Rightarrow_E B \mid X) = Pr(B \mid AX)$.*

Proof. Suppose that $X = X' \times W_\infty$ where $X' \subseteq W$. Then in general, for any α , and $A_0 \dots A_k \subseteq W_\alpha$, $(A_0 \times \dots \times A_k \times W_\infty) \cap X = (A_0 \cap (X' \times W_\alpha)) \times A_1 \times \dots \times A_k \times W_\infty$.

By 4.7 $(A \Rightarrow_E B) \cap X = \bigcup_n (\overline{A'(X' \times W_\alpha)}^n \times ABX) \cap X$. By the above observation this amounts to $ABX \cup \bigcup_{n>0} (\overline{A'(X' \times W_\alpha)} \cap X' \times W_\alpha) \times (\overline{A'(X' \times W_\alpha)}^{n-1} \times ABX)$, which simplifies to $ABX \cup \bigcup_{n>0} (\overline{A' \cap X' \times W_\alpha}) \times (\overline{A'(X' \times W_\alpha)}^{n-1} \times ABX)$

Now $Pr((A \Rightarrow_E B) \cap X) = Pr(ABX) + \sum_n Pr(\bigcup_{n>0} (\overline{A' \cap X' \times W_\alpha}) \times (\overline{A'(X' \times W_\alpha)}^{n-1} \times ABX))$. Simplifying we get $Pr(ABX) + \sum_{n>0} Pr(\overline{AX}) Pr(\overline{AX})^{n-1} Pr(ABX)$

$= Pr(ABX) + Pr(\bar{A}X) \sum_{n>0} Pr(\bar{A}X)^{n-1} Pr(ABX)$. And finally, as in theorem 4.6 this amounts to $Pr(ABX) + Pr(\bar{A}X) \frac{Pr(ABX)}{1-Pr(\bar{A}X)} = Pr(ABX) + Pr(\bar{A}X)Pr(B|AX)$.

To get $Pr(A \Rightarrow_E B | X)$ we simply divide this number by $Pr(X)$, which simplifies to $Pr(AB | X) + Pr(\bar{A} | X)Pr(B|AX)$. This reduces to $Pr(AB | X) + Pr(B | AX) - Pr(A | X)Pr(B | AX)$ using $Pr(\bar{A} | X) = 1 - Pr(A | X)$. Note that by the definition of conditional probability $Q(AB) = Q(A)Q(B | A)$, so $Pr(AB | X) = Pr(A | X)Pr(B | AX)$. Thus the last expression cancels out to $Pr(B | AX)$ as required. \square

§5. Appendix B: Tenability result with fullness. Here we construct instead a probability frame that is normal, full, and satisfies Stalnaker's condition. However, unlike the previous construction, this construction does not satisfy Harper's condition. Here it will be useful to use Stalnaker's original selection function semantics in which f maps us into W (so $f(A, x)$ picks out a world instead of a singleton of a world. When A crashes, $f(A, x)$ picks out a distinguished object, $\#$, the impossible world, instead of the empty set.) In this model we use only probability functions defined over the real numbers – to distinguish these we shall use Greek letters ' μ ', ' ρ ' and so on, to denote measures on the reals with ' λ ' being reserved for the standard Lebesgue measure.

Given a probability space $\langle W, \Sigma, \mu \rangle$ we define a subspace of W to be those spaces of the form $\langle X, \Sigma \cap \mathcal{P}(X), \mu(\cdot | X) \rangle$ with $X \in \Sigma$. I shall write μ_X for $\mu(\cdot | X)$ and Σ_X for $\Sigma \cap \mathcal{P}(X)$.

We need to employ a notion from measure theory – that of a measure-preserving map:

DEFINITION 5.1. *Let X and Y be subspaces of W . A map, $t : X \rightarrow Y$, is measure preserving on the spaces $\langle X, \mu_X \rangle, \langle Y, \mu_Y \rangle$ iff (i) $t^{-1}(A)$ is measurable in X when A is in Y and (ii) $\mu_X(t^{-1}(A)) = \mu_Y(A)$ for each A in Y 's sigma-algebra.*

As usual, the preimage of a set A under the function f , written $f^{-1}(A)$, is defined as $\{x | f(x) \in A\}$.

DEFINITION 5.2. *A selection function, f , is **stretchy** on a probability space $\langle W, \Sigma, \mu \rangle$ iff for every measurable $A \in \Sigma$, the restriction of $f(\bar{A}, \cdot)$ to A , $f(A, \cdot) : A \rightarrow A$, is measure preserving on the spaces $\langle A, \mu_{\bar{A}} \rangle, \langle A, \mu_A \rangle$. Here \bar{A} just means $W \setminus A$.*

PROPOSITION 5.3. *Suppose that there exists a tuple $\langle W, \Sigma, \mu, t_A \rangle$ satisfying the following conditions:*

1. Σ is a σ -algebra over W ,
2. μ a probability measure over Σ and for each nonempty $A \subseteq W$,
3. $t_A : \bar{A} \rightarrow A$, for each $A \subseteq W$,
4. t_A is measure preserving on $\langle \bar{A}, \mu_{\bar{A}} \rangle, \langle A, \mu_A \rangle$ whenever $\mu(A) \in (0, 1)$

Then the selection function f defined as $f(A, \cdot) = id_A \cup t_A$ is stretchy, where id_A is the identity function on A . More precisely, f , as defined below, is stretchy:

$$f(A, x) = x \text{ if } x \in A \text{ and } f(A, x) = t_A(x) \text{ if } x \in \bar{A} \text{ provided } A \text{ is nonempty}$$

$$f(\emptyset, x) = \#.$$

Any set A in Σ which has measure in $(0, 1)$ is stretched out onto its complement by f in a way that preserves the measure of its measurable subsets. (In the models we consider any pair of sets, X and Y , with measures in $(0, 1]$, can be stretched on to the other.)

Note also that f is normal and thus will validate CEM, MP, ID, 4, B, and C0. By construction $f(A, x) \in A$ and $f(A, x) = x$ whenever $x \in A$. But notice further that

A crashes ($f(A, x) = \#$) only if $A = \emptyset$, so the principles C0, B and 4 for crashing are validated in this kind of model as well.³¹ So the logic of stretchy selection functions of this type is at least $L+C0+4+B$. Whether the logic of stretchy selection functions generated this way is exactly this logic bears further investigation.

It should be clear that Stalnaker's thesis holds for any stretchy selection function. If $\mu(A) = 0$ then Stalnaker's thesis vacuously holds. If $\mu(A) = 1$ then (i) $\mu(B | A) = \mu(B)$ and (ii) $t_A^{-1}(B) \subseteq \bar{A}$ has measure 0 so $\mu(f^{-1}(A, B)) = \mu(id_A^{-1}(B)) + \mu(t_A^{-1}(B)) = \mu(AB) + 0 = \mu(B)$. Suppose that $\mu(A) \in (0, 1)$. Note that $f^{-1}(A, B) = id_A^{-1}(B) \cup t_A^{-1}(B) = AB \cup t_A^{-1}(B)$. Note that $\mu(t_A^{-1}(B) | \bar{A}) = \mu(B | A)$, since t_A is measure preserving, so $\mu(t_A^{-1}(B)) = \mu(B | A)\mu(\bar{A})$. Thus $f^{-1}(A, B)$ has a measure of $\mu(AB) + \mu(B | A)\mu(\bar{A}) = \mu(B | A)$.

5.1. Existence of a model. Here we construct a full model, $\langle W, \mathcal{B}, f, \Sigma, P \rangle$, for CP. In this model worlds will be identified with real numbers, with the constraint that the selection functions f_E are stretchy on each world $E(x)$.

- $W := [0, 1], w \in [0, 1]$.
- $\mathcal{B} := \mathcal{P}([0, 1])$
- $P := \{\lambda\}$ where λ is the Lebesgue measure on $[0, 1]$.
- Σ is the Lebesgue measurable subsets of $[0, 1]$.
- f_E is a normal selection function on E which is additionally stretchy on $\langle X, \Sigma_X, \lambda_X \rangle$ for every measurable $X = E(x)$ with positive measure.

One thing to note about this model is that I have only specified one ur-prior, λ . This does not appear to be an essential restriction—the following proof works with any measure isomorphic to the Lebesgue measure, so we could equally well expand P to $\{\lambda' | \langle [0, 1], \Sigma, \lambda' \rangle \cong \langle [0, 1], \Sigma, \lambda \rangle\}$ —i.e., the set of measures on $[0,1]$ with the Lebesgue measurable sets, that are isomorphic to the Lebesgue measure.³² For example, while the Lebesgue measure is generated by stipulating that the length of an interval (a, b) is $b - a$, the measure one gets by stipulating that the ‘length’ of the interval (a, b) be given by $b^2 - a^2$ is isomorphic to λ even though it is a very different measure.

I have thus explicitly defined every aspect of the model except for the selection functions, f_E , for accessibility relations E . I shall once again assume that E is locally an equivalence relation around the actual world w . The only thing to prove, then, is that we can find a stretchy selection function defined on the restricted conditional probability space over $E(w)$, where w is the actual world. Any extension of this function to the whole space $[0, 1]$ which only maps worlds to accessible worlds, and therefore in particular maps $E(w)$ to itself, will be a function of the desired type.³³ Moreover, by construction, it will be a stretchy selection function relative to the probability space gotten by conditioning on $E(w)$.

³¹ In my view neither C0 B nor 4 are valid; however, for the purposes of showing that a reasonable logic is consistent with Stalnaker's thesis this does not matter as every sublogic is also shown to be consistent.

³² Here the relevant notion of isomorphism is the existence of an invertible measure preserving function between the two spaces.

³³ It would also be quite easy to ensure that the function is stretchy relative to $E(x)$ for any x where E is locally an equivalence relation by constructing a stretchy selection function on $E(x)$ for each such x and extending them jointly to the whole space.

By proposition 5.3 it suffices to show that for any measurable E (of positive finite measure) and measurable $A \subseteq E$ where $A, E \setminus A$ have positive finite measure, we can find a measure preserving function from $E \setminus A$ to A , t_A . Indeed, we shall go one further and show that for *any* two measurable sets of reals, X and Y , of positive and finite measure there is a measure preserving function, t , from X to Y .

For existence of a model it thus suffices to prove the following³⁴

THEOREM 5.4. *Given any two measurable sets of reals, X and Y , of positive and finite measure there is a measure preserving function, t , from X to Y .*

This is all we need to construct the relevant stretchy selection functions. If \bar{A} and A have positive measure we can use this theorem to choose a measure preserving map t_A from \bar{A} to A .

In what follows I will need to talk about the Lebesgue measure, λ , and the renormalised Lebesgue measures on X and Y , $\lambda_X(\cdot) = \lambda(\cdot)/\lambda(X)$ and $\lambda_Y(\cdot) = \lambda(\cdot)/\lambda(Y)$. However, since this notation becomes hard to follow I shall rename the latter two measures as μ_X and μ_Y for ease of reading.

The basic idea for the proof of this theorem is to construct a pair of measure preserving maps, $f : X \rightarrow [0, 1]$ and $h : [0, 1] \rightarrow Y$, which can be composed to form a measure preserving map from X to Y . Things are more transparent if we define h in terms of a another measure preserving map, $g : Y \rightarrow [0, 1]$. Here is how we define them:

- $f : X \rightarrow [0, 1]$
- $g : Y \rightarrow [0, 1]$
- $h : [0, 1] \rightarrow Y$
- $f(x) = \mu_X((-\infty, x] \cap X)$
- $g(y) = \mu_Y((-\infty, y] \cap Y)$
- $h(a) = \begin{cases} y & \text{if there is exactly one } y \text{ such that } g(y) = a \\ a & \text{otherwise} \end{cases}$

here a can be any old member of Y , it does not matter which. We will also make use of the following property of the Lebesgue measure.

Nifty fact: the Lebesgue measure, λ , is regular. This means that:

1. $\lambda(S) = \inf\{\lambda(O) \mid S \subseteq O, O \text{ is open}\}$
2. $\lambda(S) = \sup\{\lambda(C) \mid C \subseteq S, C \text{ is closed}\}$

LEMMA 5.5. *f and g are measure preserving on open (and therefore closed) sets.*

Proof. Since g is defined exactly analogously to f it suffices to show that f is measure preserving on open sets.

Firstly note that by construction $\mu_X(f^{-1}((a, b))) = b - a$.

Let O be an open set. Since O is open, it may be written as a countable union of disjoint intervals, $\bigcup_i (a_i, b_i)$. So $\mu_X(f^{-1}(O)) = \mu_X(f^{-1}(\bigcup_i (a_i, b_i))) = \mu_X(\bigcup_i f^{-1}((a_i, b_i))) = \sum_i \mu_X(f^{-1}((a_i, b_i))) = \sum_i (b_i - a_i) = \lambda(O)$ as required.

Now let C be a closed set, so $C = [0, 1] \setminus O$ for some open set O . So $\lambda(C) = 1 - \lambda(O) = 1 - \mu_X(f^{-1}(O)) = 1 - \mu_X(f^{-1}([0, 1] \setminus C)) = 1 - (1 - \mu_X(f^{-1}(C))) = \mu_X(f^{-1}(C))$. So f is measure preserving on closed sets too. \square

³⁴ I am indebted to Gareth Davies here for some helpful suggestions regarding this proof.

THEOREM 5.6. *f and g are measure preserving.*

Proof. Let $S \subseteq [0, 1]$ be a measurable set. Then by regularity (form 1) and the fact that f is measure preserving on opens sets we have:

$$\lambda(S) = \inf\{\lambda(O) \mid S \subseteq O, O \text{ is open}\} = \inf\{\mu_X(f^{-1}(O)) \mid S \subseteq O, O \text{ is open}\} \geq \mu_X(f^{-1}(S))$$

Then by regularity (form 2) and the fact that f is measure preserving on closed sets we have:

$$\lambda(S) = \sup\{\lambda(C) \mid C \subseteq S, C \text{ is closed}\} = \sup\{\mu_X(f^{-1}(C)) \mid C \subseteq S, C \text{ is closed}\} \leq \mu_X(f^{-1}(S))$$

So $\lambda(S) = \mu_X(f^{-1}(S))$ as required. The argument that g is measure-preserving is exactly analogous. \square

Now to finish the argument we have

THEOREM 5.7. *h is measure preserving.*

Proof. Suppose that $Z \subseteq Y$.

Our strategy will be to show that $\mu_Y(Z) = \mu_Y(g^{-1}(h^{-1}(Z)))$. This suffices since $\mu_Y(g^{-1}(h^{-1}(Z))) = \lambda(h^{-1}(Z))$ by the fact that g is measure preserving. Here goes.

$g^{-1}(h^{-1}(Z)) = \{y \mid g(y) \in h^{-1}(Z)\} = \{y \mid \exists!z : g(z) = g(y) \text{ and } z \in Z\} = Z \setminus \{y \mid g(y) = g(z) \text{ for some } z \neq y\} = Z \setminus g^{-1}(\{\alpha \mid |g^{-1}(\{\alpha\})| > 1\})$.

Now note that the set $S := \{\alpha \mid |g^{-1}(\{\alpha\})| > 1\}$ is countable. We can map S injectively into \mathbb{Q} as follows: if $\alpha \in S$, then since $|g^{-1}(\{\alpha\})| > 1$ there is a rational number, q , strictly inside the convex hull of $g^{-1}(\{\alpha\})$. So we can map α to q . This mapping is injective because g is increasing: if $\alpha < \beta$ then the convex hull of $g^{-1}(\{\alpha\})$ and of $g^{-1}(\{\beta\})$ overlap at most at a boundary point (since, if $\alpha < \beta$, $g(x) = \alpha$ and $g(y) = \beta$ then $x \leq y$) and we have chosen q not to be a boundary point.

Now, of course, $\{\alpha\}$ has Lebesgue measure 0, so $\mu_Y(g^{-1}(\{\alpha\})) = 0$ since g is measure preserving. So $g^{-1}(\{\alpha \mid |g^{-1}(\{\alpha\})| > 1\})$ is a countable union of null sets, and is thus a null set. So putting this all together we have $\mu_Y(g^{-1}(h^{-1}(Z))) = \mu_Y(Z \setminus g^{-1}(\{\alpha \mid |g^{-1}(\{\alpha\})| > 1\})) = \mu_Y(Z) - 0 = \mu_Y(Z)$.

So $\mu_Y(Z) = \mu_Y(g^{-1}(h^{-1}(Z)))$. \square

This completes the proof. To obtain a measure preserving map, t , from X to Y we simply let $t = h \circ f$.

§6. Acknowledgements. Many thanks to Gareth Davies, Cian Dorr, Kenny Easwaran, John Hawthorne, the Coalition of Los Angeles Philosophers and audiences at the University of California, Irvine, the Eighth Barcelona Workshop in the Theory of Reference on Conditionals and the Formal Epistemology Workshop 2014. Thanks also to two anonymous reviewers and the editor Robert van Rooij for making several helpful suggestions regarding presentation and content.

BIBLIOGRAPHY

- Adams, E. W. (1975). *The Logic of Conditionals: An Application of Probability to Deductive Logic*, Vol. 86. Berlin: Springer.
- Ahmed, A. (2011). Walters on conjunction conditionalization. *Proceedings of the Aristotelian Society (Hardback)*, **111**, 115–122.

- Bacon, A. (Unpublished Manuscript). Conditional logics accommodating Stalnaker's thesis.
- Bennett, J. F. (2003). *A Philosophical Guide to Conditionals*. New York: Oxford University Press.
- Bradley, R. (2011). Multidimensional possible-world semantics for conditionals. *Philosophical Review*, **121**(4) pp. 539–571.
- Edgington, D. (1986). Do conditionals have truth conditions? *Crítica: Revista Hispanoamericana de Filosofía*, **18**(52), 3–39.
- Edgington, D. (1995). On conditionals. *Mind*, **104**(414), 235–329.
- Edgington, D. (2014). Conditionals. In Zalta, E. N., editor. *The Stanford Encyclopedia of Philosophy* (2014 edition). Berlin: Springer .
- Fine, K. (2012). A difficulty for the possible worlds analysis of counterfactuals. *Synthese*, **189**(1), 29–57.
- Gibbard, A. (1981). Two recent theories of conditionals. In Harper, W. L., Stalnaker, R., & Pearce, G., editors. *IFS: Conditionals, Belief, Decision, Chance and Time*, Dordrecht: D. Reidel Publ. Co., pp. 211–248.
- Hájek, A. (1989). Probabilities of conditionals — revisited. *Journal of Philosophical Logic*, **18**(4), 423–428.
- Hájek, A., & Hall, N. (1994). The hypothesis of the conditional constraint of conditional probability. In Eells, E. & Skyrms, B., editors. *Probability and conditionals: belief revision and rational decision*, Cambridge: Cambridge University Press, p. 75.
- Hall, N. (1994). Back in the cccp. In Eells, E. & Skyrms, B., editors. *Probability and conditionals: belief revision and rational decision*, Cambridge: Cambridge University Press, p. 75.
- Harper, W. L. (1976). Ramsey test conditionals and iterated belief change (a response to stalnaker). *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science: Volume II Foundations and Philosophy of Statistical Inference Proceedings of an International Research Colloquium held at the University of Western Ontario, London, Canada, May 10–13, 1973*. Berlin: Springer, pp. 117–135.
- Herzberger, H. G. (1979). Counterfactuals and consistency. *Journal of Philosophy*, **76**(2), 83–88.
- Jeffrey, R., & Edgington, D. (1991). Matter-of-fact conditionals. In Eells, E. & Skyrms, B., editors. *Proceedings of the Aristotelian Society, Supplementary Volumes*, New Jersey: John Wiley & Sons, Inc., pp. 161–209.
- Jeffrey, R., & Stalnaker, R. (1994). Conditionals as random variables. *Probability and conditionals: belief revision and rational decision*, Cambridge: Cambridge University Press, p. 31.
- Kaufmann, S. (2009). Conditionals right and left: Probabilities for the whole family. *Journal of Philosophical Logic*, **38**(1), 1–53.
- Kratzer, A. (1986). Conditionals. *Chicago Linguistics Society*, **22**, 1–15.
- Kratzer, A. (2012). *Modals and Conditionals: New and Revised Perspectives*, Vol. 36. Oxford: Oxford University Press.
- Lewis, D. (1973). *Counterfactuals*. Oxford: Basil Blackwell.
- Mårtensson, J. (1999). *Subjunctive conditionals and time: A defense of the classical approach* Ph.D thesis, University of Gothenburg, Gothenburg.
- McGee, V. (1989). Conditional probabilities and compounds of conditionals. *The Philosophical Review*, **98**(4), 485–541.
- Nolan, D. (2003). Defending a possible-worlds account of indicative conditionals. *Philosophical Studies*, **116**(3), 215–269.

- Russell, J., Hawthorne, J., & Buchak, L. (Forthcoming). Groupthink. *Philosophical studies*.
- Santos, P. (2008). Context-sensitivity and (indicative) conditionals. *Disputatio*, **2**(24), 1–21.
- Schulz, M. (Forthcoming). Counterfactuals and arbitrariness. *Mind*.
- Stalnaker, R. (1968). A theory of conditionals. *Studies in logical theory*, **2**, 98–112.
- Stalnaker, R. (1975). Indicative conditionals. *Philosophia*, **5**(3), 269–286.
- Stalnaker, R. (1976). Letter to van Fraassen. In Harper, W. L., & Hooker, C. A., editors. *Foundations of probability theory, statistical inference, and statistical theories of science*, Vol. 1, Berlin: Springer, pp. 302–306.
- Stalnaker, R. (1981). A defense of conditional excluded middle. In Harper, W. L., Stalnaker, R., & Pearce, G., editors. *IFS: Conditionals, belief, decision, chance, and time*, Dordrecht: D. Reidel Publ. Co., pp. 87–104.
- Stalnaker, R. (2009). Conditional propositions and conditional assertions. In Egan, A., & Weatherson, B., editors. *Epistemic Modality*, Oxford: Oxford University Press, 227–248.
- Stalnaker, R. (1970). Probability and conditionals. *Philosophy of Science*, **37**, 64–80.
- Stalnaker, R. (1987). *Inquiry*. Cambridge: MIT Press.
- Tichý, P. (1978). A new theory of subjunctive conditionals. *Synthese*, **37**(3), 433–457.
- Tooley, M. (2002). Backward causation and the stalnaker-lewis approach to counterfactuals. *Analysis*, **62**(275), 191–197.
- Van Fraassen, B.C. (1976). Probabilities of conditionals. *Foundations of probability theory, statistical inference, and statistical theories of science*, Vol. 1, pp. 261–308.
- Van Rooij, R. (1999). Gibbard's problem: The context dependence of conditional statements. *Proceedings Dutch-German workshop on nonmonotonic reasoning*.
- Weatherson, B. (2009). Conditionals and indexical relativism. *Synthese*, **166**(2), 333–357.

DEPARTMENT OF PHILOSOPHY
 UNIVERSITY OF SOUTHERN CALIFORNIA
 3709 TROUSDALE PARKWAY
 LOS ANGELES, CA 90089
 E-mail: abacon@usc.edu