

Counterfactual Probability

Ginger Schultheis

July 2020

1 Introduction

Stalnaker's Thesis about indicative conditionals is, roughly, that the probability one ought to assign to an indicative conditional equals the probability that one ought to assign to its consequent conditional on its antecedent. The thesis seems right. If you draw a card from a standard 52-card deck, how confident are you that the card is a diamond if it's a red card? To answer this, you calculate the proportion of red cards that are diamonds—that is, you calculate the probability of drawing a diamond conditional on drawing a red card.

Skyrms' Thesis about counterfactual conditionals is, roughly, that the probability that one ought to assign to a counterfactual equals one's rational expectation of the chance, at a relevant past time, of its consequent conditional on its antecedent.¹ This thesis also seems right. If you decide not to enter a 100-ticket lottery, how confident are you that you would have won had you bought a ticket? To answer this, you calculate the prior chance—that is, the chance just before your decision not to buy a ticket—of winning conditional on entering the lottery.

The central project of this article is to develop a new uniform theory of conditionals that allows us to derive a version of Skyrms' Thesis from a version of Stalnaker's Thesis, together with a *chance-deference norm* relating rational credence to beliefs about objective chance.²

I say a *version of* Stalnaker's Thesis because it is well known that Stalnaker's Thesis itself is subject to a series *triviality results*. Assuming orthodox probability theory, it can be shown that, except in trivial cases, there is no way to interpret

¹The label 'counterfactual conditional' is misleading. Consider:

- (1) If I caught the four o'clock train today, I would make it to the meeting by five.

An utterance of (1) suggests that the speaker *leaves open* the possibility that she will catch the four o'clock train. Thus, 'counterfactual conditional' hardly seems an apt label if it is to cover conditionals like (1). Some authors use the term 'subjunctive conditional'. But this label is also misleading. It suggests that the main grammatical difference between indicative conditionals and conditionals like (1) has to do with subjunctive mood. But that is not the case. In most languages, the primary grammatical difference between indicatives and conditionals like (1) is that the latter exhibit an extra layer of past tense morphology. I will continue to use the term 'counterfactual conditional' to refer to conditionals, like (1), that contain this extra layer of past tense morphology because the term is familiar and I know of no better label.

²For discussion of Skyrms' Thesis, see Skyrms (1980), Edgington (2008), Williams (2012), Moss (2013), Schwarz (2016), Schulz (2017), and Khoo (ms).

the indicative conditional uniformly so that Stalnaker's Thesis holds universally, that is, for all rational probability functions. And I say a *version of* Skyrms' Thesis because, as I will show in §3 of this paper, that thesis also has unacceptable trivializing consequences given orthodox probability theory.

The paper opens in §2 with a discussion of Stalnaker's Thesis and, following van Fraassen (1976) and Bacon (2015), suggests an improved, *context-sensitive* version of the thesis, which I will call the *Local Thesis*. The rest of the paper breaks into two main parts. The first part (§3-§5) refutes Skyrms' Thesis and develops a context-sensitive replacement, the *Local Conditional Principal Principle*—a counterfactual analogue of David Lewis's *Principal Principle*.

The second part (§6-§8) begins by introducing a neo-Stalnakerian, uniform theory of conditionals. At a high level, my view says that all of the semantic differences between indicative conditionals and counterfactual conditionals boil down to differences in what is held fixed. When we evaluate an indicative conditional, we hold fixed all of our knowledge; when we evaluate a counterfactual, we hold fixed a contextually-determined subset of our knowledge. I show that this theory allows us to derive the Local Conditional Principal Principle from the Local Thesis and Lewis's Principal Principle. And although a full tenability result for the Local Conditional Principal Principle is beyond the scope of this paper, I will argue in the final section that there is good reason to be optimistic that the principle is indeed tenable within the Stalnakerian framework that I develop.

2 From Stalnaker's Thesis to The Local Thesis

Our eventual goal to derive a plausible, contextualist-friendly version of Skyrms' Thesis from a plausible, contextualist-friendly version of Stalnaker's Thesis and a plausible chance-deference norm. Here I introduce the contextualist-friendly version of Stalnaker's Thesis—the *Local Thesis*.

Suppose I'm a detective working on a murder case. I know that it was either the butler or the gardener. My credence that it was the gardener, on the supposition that it wasn't the butler, is high. Correspondingly, I will be confident in (2).

(2) If the butler didn't do it, it was the gardener.

Take another case. Suppose I know that the four o'clock train arrives within an hour about 75% of the time. So I am 75% confident that John will make it by five, supposing he catches the four o'clock train. Correspondingly, I will be 75% confident in (3).

(3) If John catches the four o'clock train, John will be here by five.

Examples like these are easy to multiply, and the pattern of probability assignments is robust, leading many theorists to endorse some version of Stalnaker's Thesis. Where $A > B$ stands for the indicative conditional with antecedent A and consequent B , Stalnaker's Thesis is as follows:

Stalnaker's Thesis

For any rational credence function P such that $P(A) > 0$: $P(A > B) = P(B|A)$.³

Stalnaker's Thesis, as I will understand it here, is a normative thesis. It says that, *if you're rational*, then your credence in $A > B$ is equal to your credence in B conditional on A (whenever the conditional probability is defined). For instance, in the murder case, Stalnaker's Thesis says that, if I'm rational, then my credence (2), *if the butler didn't do it, it was the gardener*, equals my conditional credence that the gardener committed the murder given that it wasn't the butler. I take no stand on whether there are irrational subjects whose credences in indicative conditionals diverge from their conditional credences.

Despite its initial plausibility, Stalnaker's Thesis is false. David Lewis (1976) showed that Stalnaker's Thesis has trivializing consequences given just two standard assumptions: (1) that rational credence functions obey the laws of probability; and (2) that the set of rational credence functions is closed under conditionalization, so that if P is rational and $P(A) > 0$, then the probability function $P(\cdot|A)$ that results from conditioning P on A is also rational. Given (1) and (2), Stalnaker's Thesis entails that whenever you think that A and B are compatible, and that A and $\neg B$ are compatible, you are certain of B conditional on the indicative conditional $A > B$. This consequence is unacceptable. To illustrate, suppose that it's compatible with my beliefs that Milo is at a picnic and in a good mood, and compatible with my beliefs that he's at a picnic and in a bad mood. Stalnaker's Thesis predicts that I should be certain that Milo is in a good mood, conditional on *if Milo is at a picnic, he's in a good mood*. In other words, if I learn the conditional *if Milo is at a picnic, he's in a good mood*, then I should be certain that he's in a good mood. But that's absurd! For all I know, he's not at a picnic; for all I know, he's not in a good mood. So if we keep (1) and (2), we have no choice but to reject Stalnaker's Thesis.

Fortunately there are limited versions of Stalnaker's Thesis that capture its intuitive motivation but are not subject to the Lewisian triviality results. The one that I will be concerned with—the *Local Thesis*—is motivated by a contextualist

³I assume the Ratio Formula: If $P(A) > 0$, then $P(B|A) = P(A \wedge B)/P(A)$.

theory of indicative conditionals. Before I state the thesis, let me say a few words to motivate contextualism independent of its connection to Stalnaker's Thesis and avoiding triviality. Contextualism about indicative conditionals is the view that what proposition is expressed by an utterance of an indicative conditional depends, in part, on a contextually-supplied body of information. Often that information is simply the speaker's knowledge. Other times it is some other body of information. For example, it may be the knowledge of some other individual or group. And sometimes the standards are more demanding than knowledge—such as being known *with certainty*. Other times they are less demanding. To allow for this variability, I refer to this contextually-supplied body of information simply as the information *associated with the context*.

Why accept contextualism? One argument comes from so-called *stand-off cases*. Consider:

Sly Pete and Mr. Stone are playing poker on a Mississippi riverboat. It is now up to Pete to call or fold. My henchman Zack sees Stone's hand, which is quite good, and signals its content to Pete. My henchman Jack sees both hands and sees that Pete's hand is rather low, so that Stone's is the winning hand. At this point, the room is cleared. A few minutes later, Zack slips me a note which says, 'If Pete called, he won,' and Jack slips me a note which says 'If Pete called, he lost.' I know that both notes come from my trusted henchmen but do not know which of them sent which note. I conclude Pete folded. (Gibbard 1981, p. 231)

According to the contextualist, Zack says something true when he writes:

(4) If Pete called, he won.

Likewise, Jack says something true when he writes:

(5) If Pete called, he lost.

Zack's conditional is true *relative to Zack's information*. Jack's conditional is true *relative to Jack's information*. Nevertheless, there is no information state—that is, no context—relative to which *both* conditionals are true.

If we're contextualists, Stalnaker's Thesis needs to be refined because it doesn't mention context. And, as we will see, these refinements are also sufficient for avoiding the triviality results. Specifically, we need to do two things. First, we need to add contextual parameters. Both the indicative conditional and the probability function need to be indexed to a context. Second, we must *coordinate* these two contextual parameters—the indicative conditional proposition on the

left side of the equation must be indexed to the same context as the probability function on the right side. The notation is as follows. I write $A >_c B$ for the proposition expressed by the indicative conditional in a given context c . I write P_c for the probability function associated with c . (To simplify, I assume that P_c is result of conditioning a uniquely rational initial credence function P_o on the information associated with context c .) The Local Thesis is as follows:

The Local Thesis

$$P_c(A >_c B) = P_c(B|A) \text{ whenever } P_c(A) > 0.$$

Suppose, for a moment, that the information associated with a given context c is the speaker's knowledge. Then the Local Thesis says that the probability that the speaker in c assigns to the proposition expressed by the indicative, relative her information—*her* indicative conditional, as I will sometimes say—is equal to the probability that she assigns to B conditional on A . But importantly, it is silent about the probability that she assigns to propositions expressed by the indicative conditional in contexts other than her own. I won't get into the details, but as Bacon (2015) and others have shown, it is for precisely this reason that the Local Thesis is not subject to Lewisian triviality results. Indeed, it is not subject to any triviality results. Building on the work of van Fraassen (1976), Bacon (2015) has shown that the Local Thesis—or, more carefully, a thesis that is very close to the Local Thesis—is tenable within a possible-worlds semantics for indicatives based on Stalnaker's *selection semantics*. I return to these tenability results in §8.^{4,5}

In the next few sections—sections §3–§5—I will set the Local Thesis to one side as I work up to my preferred formulation of Skyrms' Thesis—the *Local Conditional Principal Principle*. As we will see, that principle is similar in spirit to

⁴Note that van Fraassen himself is not explicit about how contextualism figures in his tenability results (though he does mention contextualism as a way of escaping Lewisian triviality). Nevertheless, it is natural to interpret his results within a contextualist framework. See §8 for more discussion. See also Stefan Kaufman (2005, 2005, 2009) for important work in this tradition, and Justin Khoo (ms) for a tenability result that is similar to van Fraassen's but does not rely on contextualism.

⁵There are important differences between my statement of the Local Thesis and Bacon's contextualist-friendly version of Stalnaker's Thesis. Here is Bacon's version, which he calls CP:

CP. For any rational initial probability function P_o , any (contextually-supplied) evidential accessibility relation E , and any w : $P_o(A >_E B|E(w)) = P_o(B|A \wedge E(w))$

For Bacon, the proposition expressed by an indicative conditional is relativized to an *accessibility relation*. In the formulation of the Local Thesis in the main text, the proposition expressed by an indicative conditional is relativized to an evidence proposition—a set of worlds. I agree with Bacon that there are compelling reasons to use accessibility relations rather than sets of worlds, but it greatly simplifies things to work with a version of the Local Thesis that uses sets of worlds. The central arguments of this paper do not turn on the differences between CP and the Local Thesis.

the Local Thesis. In this section, I hope to have set the basic groundwork for articulating a contextualist-friendly connection between chance and counterfactuals.

3 Triviality for Counterfactuals: Skyrms' Thesis

The last section concerned the relationship between indicative conditionals and probability. In this section, we turn to my primary concern in this paper—the relationship between counterfactuals and chance. I begin by motivating the most natural formulation of this connection—*Skyrms' Thesis*. Then I present a new argument showing that Skyrms' Thesis has unacceptable trivializing consequences.

Suppose that I decide not to flip a fair coin at noon. And suppose I know that the coin had a 50% chance of landing heads and a 50% chance of landing tails. How confident should I be in the counterfactual (6)?

(6) If the coin had been flipped at noon, it would have landed heads.

50% seems to be the only reasonable answer.

Now imagine that I don't know the coin is fair. I divide my credence evenly between two hypotheses about the chance of heads—that the chance of heads is 30% and that the chance of heads is 60%. How confident should I be in (6) in this case? A natural answer: $(50\% \times 30\%) + (50\% \times 60\%) = 45\%$. That is, my credence in (6) should be equal to my *expectation* of the conditional chance, just before noon, of the coin landing heads conditional on being flipped.

This data motivates Skyrms' Thesis—a general principle that ties rational credences in counterfactuals to rational expectations of prior chances.⁶ To state the thesis, let t be a relevant past time; let $Ch_t(B|A) = x$ be the proposition that the chance, at t , of B conditional on A is equal to x ; and let $A \Box \rightarrow B$ stand for the counterfactual with antecedent A and consequent B . Then Skyrms' Thesis is as follows:

Skyrms' Thesis

For any rational P : $P(A \Box \rightarrow B) = \sum_x x \times P(Ch_t(B|A) = x)$

Skyrms assumption that t is always a past time is not quite right.⁷ This issue

⁶Note that Skyrms himself formulates the thesis in terms of *propensities*, rather than objective chances. Propensities provide one way of thinking about what objective chances are. But there are other interpretations, such as Lewis's own account developed in Lewis (1980) and Lewis (1994). To remain neutral about what objective chances are, I replace Skyrms' formulation in terms of propensities with one that only mentions objective chances.

⁷Take, for example:

(1) If I caught the four o'clock train today, I would make it to the meeting by five.

The probability that I assign to (1) is equal to the *present* chance of making the meeting condi-

is not especially important for my purposes, and I will often speak as though the relevant time is in the past.

Like Stalnaker's Thesis, I take Skyrms' Thesis to be a normative thesis. As I will understand the thesis here, it says that *if you're rational*, then your credence in $A \Box \rightarrow B$ is equal to your expectation of the chance, at the relevant time, of B conditional on A . In the coin case, for example, Skyrms' Thesis says that, if I'm rational, then my credence in (6), *if the coin had been flipped at noon, then it would have landed heads*, equals my expectation of the conditional chance, just before noon, of the coin landing heads conditional on being flipped.⁸

Skyrms' Thesis, I argue, has unacceptable trivializing consequences. Given orthodox probability theory, it entails that, if you give positive credence to $A \Box \rightarrow B$, and positive credence to $\neg(A \Box \rightarrow B)$, then your credence in the counterfactual $A \Box \rightarrow B$ is equal to your credence in the following proposition: the conditional chance, at t , of B given A is equal to one. We can go on to derive other absurd consequences, but this is bad enough. To illustrate, go back to the coin case. We said that I am 50% confident in (6), *if the coin had been flipped, it would have landed heads*. Skyrms' Thesis entails that I am 50% confident in the following proposition: the chance of the coin landing heads conditional on being flipped equals one. But that's absurd! We can easily imagine that I am certain that the coin is fair, which is to say that I am *not* 50% confident in the proposition that the coin has a 100% chance of landing heads, conditional on being flipped.

Here is the triviality argument. First observe that Skyrms' Thesis entails (1) and (2) below (I omit time references for readability):

- (a) For any rational P , if $P(A \Box \rightarrow B) = 1$, then $P(Ch(B|A) = 1) = 1$

tional on catching the four o'clock train. In general, when we evaluate counterfactual conditionals whose antecedents concern events that will occur at some future time, we set our credence in the counterfactual to our expectation of the present chance of the consequent given the antecedent.

⁸There are known counterexamples to Skyrms' Thesis involving *counterlegals*—counterfactuals whose antecedents concern events that violate the laws of nature. A counterlegal may have positive probability even though its antecedent has chance zero, in which case the chance of the consequent conditional on the antecedent is undefined. To deal with cases like this, one option is to use *Popper functions*, which would allow conditional chances to be well-defined even if the conditioned proposition is chance zero. Another option is to treat Skyrms' Thesis as a special case of a more general thesis stated in terms of *hypothetical* probability functions. On this view, one's credence in a counterfactual is given by the probability of the consequent conditional on the antecedent, relative to a hypothetical probability distribution that assigns positive probability to the antecedent and is suitably related to one's actual probability distribution. If this hypothetical probability distribution matches the objective chances whenever the latter are defined, Skyrms' Thesis would come out as a special case of this more general norm. (See Edgington (2008) for discussion.) For the purposes of this paper, I do not need to take a stand on how to handle counterlegals and other counterfactuals with chance-zero antecedents, so I set these aside.

(b) For any rational P , if $P(A \Boxrightarrow B) = 0$, then $P(Ch(B|A) = 1) = 0$

If you are certain that $A \Boxrightarrow B$ is true, then you are certain that the prior chance of B conditional on A is one. If you are certain that the counterfactual is false, then you are certain that the prior chance of B conditional on A is zero.

Now consider any rational probability function P . As Lewis assumed in his triviality results for Stalnaker's Thesis, I assume that the class of rational probability functions is closed under conditionalization: if P is rational, and $P(A) > 0$, then $P(\cdot|A)$ is rational.

Suppose that $P(A \Boxrightarrow B) > 0$ and $P(\neg(A \Boxrightarrow B)) > 0$. Then (a) entails (c), and (b) entails (d):

$$(c) P(Ch(B|A) = 1|A \Boxrightarrow B) = 1$$

$$(d) P(Ch(B|A) = 1|\neg(A \Boxrightarrow B)) = 0$$

And, by the Law of Total Probability, we know (e):

$$(e) P(Ch(B|A) = 1) = P(Ch(B|A) = 1|A \Boxrightarrow B) \times P(A \Boxrightarrow B) + P(Ch(B|A) = 1|\neg(A \Boxrightarrow B)) \times P(\neg(A \Boxrightarrow B))$$

(c), (d), and (e) together give us:

$$(f) P(Ch(B|A) = 1) = P(A \Boxrightarrow B)$$

Skyrms' Thesis has allowed us to derive (f) from the assumption that $P(A \Boxrightarrow B) > 0$ and $P(\neg(A \Boxrightarrow B)) > 0$. This result is unacceptable; we have no choice but to reject Skyrms' Thesis.

4 The Conditional Principal Principle

Skyrms' Thesis seemed plausible on first glance, but closer inspection revealed it to be untenable. Where do we go from here? To answer this question, I turn to the literature on *chance-deference norms*—norms governing the relationship between our credences and our beliefs about objective chance. For, viewed abstractly, Skyrms' Thesis is a kind of chance-deference norm; it tells us to defer to certain conditional chances when setting our credences in counterfactuals.⁹ My

⁹I am not the first to draw an analogy between Skyrms' Thesis and chance-deference norms. See Schulz (2017) for extended discussion. Schulz endorses a counterfactual analogue of the Principal Principle that is similar to my Conditional Principal Principle. There are important differences between my principle and the one defended by Schulz, however. One important difference

starting point is David Lewis's *Principal Principle*. After introducing this principle, I propose a counterfactual version of it—the *Conditional Principal Principle*. (A refined, context-sensitive version of this—the Local Conditional Principal Principle—will be my final proposal.)

Note that Lewis's Principal Principle is just one of many non-trivializing formulations of the norm to defer to objective chance. There are others, such as Ned Hall's *New Principle*.¹⁰ I will not defend the Principal Principle over its rivals. I only wish to describe one plausible formulation of the norm to defer to objective chance, and to construct a counterfactual analogue of that principle. I am confident that we can formulate counterfactual analogues of other chance norms—a counterfactual version of Hall's New Principle, for example—but I leave this for future research.

Let me begin with some examples to motivate Lewis's Principal Principle. Suppose I know that a fair die will be tossed in one hour. I know that it has a 50% chance of landing on an even number and a 50% chance of landing on an odd number. I have no other relevant information. To what degree should I believe that the die will land on an even number? 50% seems to be the only reasonable answer.

Now suppose I don't know that the die is fair. I divide my credence evenly between two hypotheses about the chance of even—that it's 30% and that it's 60%. To what degree should I believe that the die will land even? A natural answer: $(50\% \times 30\%) + (50\% \times 60\%) = 45\%$. My credence that the die will land even should be equal to my expectation of the chance of even.

The intuitions that I have articulated about these two cases are predicted by Lewis's Principal Principle. I am going to state Lewis's principle in a somewhat unfamiliar way. The formulation I adopt is stated in terms of conditional chances—specifically, chances conditioned on the subject's total evidence. (The more familiar formulation is not stated in terms of chances conditioned on the subject's total evidence.) There are two reasons for this: first, a statement in terms of conditional chances is more straightforwardly extendable to my final goal, which is to state a counterfactual version of the Principal Principle. Second,

is that he distinguishes two chance functions: the chance function that figures in the Principal Principle itself (the physical chances) and the one that figures in the counterfactual analogue of the Principal Principle (the counterfactual chances). The two chance functions have different properties. I do not distinguish two chance functions—the chance function that figures in the Conditional Principal Principle is the very same chance function as the one that figures in the Principal Principle. (The reason for this is that I intend to derive the Conditional Principal Principle from the Principal Principle (and the Local Thesis)). See also Fitelson (ms) for discussion of analogies between the Principal Principle and Stalnaker's Thesis.

¹⁰See Hall (1994) and Hall (2004).

the conditional chance formulation doesn't appeal to what Lewis calls *inadmissible evidence*, a theoretically fraught notion that does appear in more familiar formulations of the thesis. (In footnote 12, I prove that my formulation of the Principal Principle entails the more familiar formulation.)

To begin, I introduce the notion of an *ur-chance function*. Consider any world w . If w contains an earliest moment, then the ur-chance function of w —denoted ch_w —is a function that takes a proposition and returns its chance at the earliest moment of w . Later chance functions are defined in terms of ch_w as follows. Where $H_{t,w}$ is a complete specification of the history at w up to the moment t , then $ch_w(\cdot|H_{t,w})$ is a function that takes a proposition and returns its chance, at t , in w .¹¹ Let P_o be any reasonable initial credence function. Let ' π ' be a rigid designator that picks out a particular ur-chance function. Let ' Ch ' be a definite description for 'the initial chance, whatever it is'. Finally, let E be any total body of evidence that is compatible with the proposition $Ch = \pi$. With this notation, we state the Principal Principle as follows.

The Principal Principle

$$P_o(A|E \wedge Ch = \pi) = \pi(A|E)$$

Suppose that your total evidence is E . And suppose that you learn what the initial chance function is, which is to say that you learn $Ch = \pi$. Then, the Principal Principle says, you should adopt the opinions π would have were it given your evidence E . To use Ned Hall's metaphor, this version of the Principal Principle tells us to treat chance as an *analyst expert*. We defer to the initial objective chance function not because it has evidence than we don't have—it doesn't have any evidence at all—but because we think it's especially good at *evaluating* evidence. Upon learning what the initial chance function is, we feed it our evidence, and then defer to its *conditional opinions*—the opinions it would have, if it knew everything we know.

Now, you might be wondering: How could this principle be useful for ordinary subjects? It tells us what to do if we learn what the entire initial chance function is like, but we're never in that situation. What we do learn are facts about the chances of *specific* propositions—the proposition that the chance that a certain die will land on an even number is 50%, for example. But, as I have stated it, the Principal Principle doesn't seem to say anything at all about ordinary cases like these. For reasons that I elaborate in a footnote, this concern fails to appreciate the strength of my formulation of the Principal Principle when com-

¹¹We can also use $ch_w(\cdot|H_{t,w})$ to define the ur-chance function at worlds that do not have earliest moments. This issue is orthogonal to the arguments in this paper so I set it aside.

bined with the laws of probability and our definition of chance.¹² My statement of the Principal Principle does, despite appearances, deliver the right predictions in ordinary cases of deference to chance. In the case of the die, for example, it says: if you're rational, and you know that the die has a 50% chance of landing even, and you have no other relevant evidence, then you are 50% confident that the die will land on an even number.

I suggest that we replace Skyrms' Thesis with a counterfactual version of the Principal Principle, which I will call *Conditional Principal Principle*. Now, my statement of the Principal Principle tells us to give the chances all of our evidence and then align our credences with the objective chances conditional on our total evidence. Clearly, this won't work in the counterfactual case. Suppose I know that I did not strike the match at noon. In that case, the chance, at noon, that the match would light conditional on my striking it *and all of my evidence* is undefined. But, we may suppose, my credence in (8) is close to one.

(8) If I had struck the match at noon, it would have lit.

So we can't give the initial chances all of our evidence, as we did with the Principal Principle. But what body of evidence should we use instead? An immediate answer that won't work: give the chances all of our evidence *minus* our evidence that the antecedent is false. This won't work because in any ordinary context in which I assert (8), I also know that the match did not light. The prior chance that the match would light conditional on being struck and *this* piece of knowledge

¹²We can show that my formulation of the Principal Principle entails Lewis's more familiar formulation. Where E is any proposition that is compatible with and wholly admissible with respect to $Ch_t(A) = x$, Lewis's formulation says:

$$(7) \quad P_o(A|E \wedge Ch_t(A) = x) = x$$

Since the t -chances are given by conditioning the ur-chance function on history up to t , we know that $Ch_t(A) = x$ is equivalent to the disjunction: $(H_t^1 \wedge Ch = \pi^1) \vee (H_t^2 \wedge Ch = \pi^2) \vee \dots \vee (H_t^n \wedge Ch = \pi^n)$, for all H_t^i and π^i such that $\pi^i(A|H_t^i) = x$. Following Meacham (2010), I will assume that E is admissible with respect to $Ch_t(A) = x$ just in case $E \wedge Ch_t(A) = x$ can be expressed as the disjunction of a subset of the $(H_t^i \wedge Ch = \pi^i)$'s associated with $Ch_t(A) = x$.

With this definition of admissibility in hand, we can show that (7) follows from my formulation of the Principal Principle:

$$\begin{aligned} P_o(A|E \wedge Ch_t(A) = x) &= P_o(A|(H_t^1 \wedge Ch = \pi^1) \vee \dots \vee (H_t^n \wedge Ch = \pi^n)) \\ &= P_o(A|(H_t^1 \wedge Ch = \pi^1)) \cdot P_o(H_t^1 \wedge Ch = \pi^1) + \dots + P_o(A|(H_t^n \wedge Ch = \pi^n)) \cdot P_o(H_t^n \wedge Ch = \pi^n) \\ &= x \cdot P_o(H_t^1 \wedge Ch = \pi^1) + \dots + x \cdot P_o(H_t^n \wedge Ch = \pi^n) \\ &= x \end{aligned}$$

The first step follows from our definitions of chance and admissibility. The second line follows from the first by the Law of Total Probability. The third line follows from the second line by my formulation of the Principal Principle. (For example, my formulation of the Principal Principle entails that $P_o(A|(H_t^1 \wedge Ch = \pi^1)) = \pi^1(A|H_t^1)$, which, by hypothesis, is equal to x .) The last line follows from the third line because the $(H_t^i \wedge Ch = \pi^i)$'s form a partition and so their probabilities sum to 1.

of mine is equal to zero. But again, we may suppose that my credence in (8) is close to one.

Here is what I suggest. Instead of giving the chances all of our evidence, we give them the subset of our evidence *that we hold fixed when we evaluate the counterfactual*. Let me take a moment to explain just what this subset is because it will be very important in what follows.

When we evaluate a counterfactual, we imagine a hypothetical scenario in which the antecedent is true and ask ourselves whether the consequent is also true in that scenario. To do this, we temporarily release some of our knowledge—our knowledge of the antecedent’s falsity, among other things. But we don’t release all of our knowledge, as philosophers have long observed. We hold much of what we know fixed. Take, for instance, Adams’ famous example:

(9) If Oswald hadn’t shot Kennedy, someone else would have.

When we evaluate (9), we tend to hold fixed our knowledge of how things went before the assassination—that Oswald acted alone, that he was not part of a conspiracy, and so forth. We clearly do not hold fixed all of our knowledge of how things went *after* the assassination—that the papers reported that Kennedy was shot, that his funeral took place, or that Johnson assumed the presidency in 1963.¹³

My *Conditional Principal Principle* says that your credence in the counterfactual $A \Box \rightarrow B$, upon learning what the ur-chance function is, should be equal to the ur-chance of B conditional on A and the evidence you’re holding fixed. To state the principle, we use the notation that we introduced to state the Principal

¹³This is not to say that we *never* hold fixed facts about history at later times. Here is a famous example due to Sidney Morgenbesser. Just before tossing a fair coin, I offer you a bet at good odds that it will land heads. You decline the bet. I toss the coin, and it lands heads. You regret your decision to decline the bet, for you know that if you *had* accepted the bet, you would have won. You are rationally confident in the counterfactual:

(10) If I had accepted the bet, I would have won.

This is so even though you know that the chance, *at the time of the antecedent*, of winning conditional on accepting the bet was only 50%. Thus, when you evaluate (10), you hold fixed your knowledge of the outcome of the toss, in addition to your knowledge of history before the toss.

Note that if the antecedent of a counterfactual concerns a long interval of time, we do not tend to hold fixed our knowledge of what took place during that interval. Here is a famous example due to John Pollock:

(11) If my coat had been stolen last year, it would have been stolen on December 31.

Although I know that it was not stolen on the first day, or the second day, or the third day, and so on, I do not hold my knowledge of any of these facts fixed when I evaluate (11). This observation of Pollock’s conflicts with David Lewis’s influential treatment of counterfactuals. Lewis says that we evaluate the counterfactual in the *most similar* antecedent worlds to actuality, and worlds where history diverges from that of the actual world at later times are *ceteris paribus* more similar to actuality than worlds where history diverges earlier.

Principle. We let P_o be any reasonable initial credence function. We let $Ch = \pi$ be the proposition that the ur-chance function is identical to π . And we let E be any total body of evidence that is compatible with the proposition $Ch = \pi$. We introduce one new piece of notation: E^- will be the set of worlds consistent with all of the information that is held fixed when evaluating the counterfactual—a strict superset of E .

The Conditional Principal Principle is as follows.

The Conditional Principal Principle

$$P_o(A \Box \rightarrow B | E \wedge Ch = \pi) = \pi(B | A \wedge E^-)$$

Suppose you learn that the initial chance function is π . Then the Conditional Principal Principle says: if you're rational, then your credence in the counterfactual $A \Box \rightarrow B$ equals the credence that π would have in B given A if π were given all of the information you are holding fixed.

Think about it this way. The information that you hold fixed is the information that you judge relevant to determining whether B would have been true if A had been true. In the example of Kennedy's assassination, for instance, you hold fixed what you know about the events leading up to Oswald pulling the trigger, as well as your general knowledge about presidential assassinations, among other things. The initial chance function π should have this information if it is to determine how likely it is that someone else shoots Kennedy supposing Oswald doesn't. You don't hold fixed that Oswald shot Kennedy, and that, as a result, nobody else did. Intuitively, this information is irrelevant to what would have happened if Oswald *hadn't* shot Kennedy. So the initial chance function π has no use for this information.

The Conditional Principal Principle says that, upon learning that the initial chance function is π , give π all of the information that you hold fixed when evaluating (9), *if Oswald hadn't shot Kennedy, someone else would have*. Then ask π : Given this information, how likely do you think it is that someone else shoots Kennedy supposing Oswald doesn't? If you're rational, the answer to this question is the credence that you assign to (9).¹⁴

¹⁴Like my formulation of the Principal Principle, my Conditional Principal Principle doesn't look like it's going to be very useful. It tells us what to do if we learn what the entire initial chance function is like, but it doesn't seem to say anything about ordinary cases. But again, this concern fails to appreciate the strength of my formulation of the Principal Principle when combined with the laws of probability and our definition of chance. We can show that my statement of the Conditional Principal Principle entails (12), where E is any proposition that is compatible with the proposition $Ch_t(B|A \wedge E^-) = x$:

$$(12) \quad P_o(A \Box \rightarrow B | E \wedge Ch_t(B|A \wedge E^-) = x) = x$$

The derivation of (12) from my formulation of the Conditional Principal Principle mirrors the

The Conditional Principal Principle avoids the triviality results that refute Skyrms' Thesis. Recall that Skyrms' Thesis implies (1) and (2) below (I omit time references for readability):

(a) For any rational P , if $P(A \Boxrightarrow B) = 1$, then $P(Ch(B|A) = 1) = 1$

(b) For any rational P , if $P(A \Boxrightarrow B) = 0$, then $P(Ch(B|A) = 1) = 0$

Taken together, (a) and (b) imply that, if you give positive credence to $A \Boxrightarrow B$ and positive credence to $\neg(A \Boxrightarrow B)$, then your credence in $A \Boxrightarrow B$ is equal to your credence in the proposition $Ch(B|A) = 1$.

The Conditional Principle escapes the triviality result because it does not entail (a) and (b). Instead, it entails (a') and (b'), where P^E is any rational initial probability function conditioned on evidence E :

(a') If $P^E(A \Boxrightarrow B) = 1$, then $P^E(Ch(B|A \wedge E^-) = 1) = 1$

(b') If $P^E(A \Boxrightarrow B) = 0$, then $P^E(Ch(B|A \wedge E^-) = 1) = 0$

But if we have (a') and (b') in place of (a) and (b), we can block the next step of the argument. For (a') and (b') do not entail (c') and (d'):

(c') $P^E(Ch(B|A \wedge E^-) = 1 | A \Boxrightarrow B) = 1$

(d') $P^E(Ch(B|A \wedge E^-) = 1 | \neg(A \Boxrightarrow B)) = 0$

And without (c') and (d') we cannot complete the argument that we used to refute Skyrms' Thesis.

5 The Local Conditional Principal Principle

We're on the right track. We have a counterfactual version of Lewis's Principal Principle that avoids the triviality result presented in §3. Still, the principle is not quite right as it stands. Counterfactuals, it is widely agreed, are context sensitive—which proposition is expressed by an utterance of a counterfactual conditional depends, in part, on the conversational context in which the utterance occurs. But the Conditional Principal Principle does not mention context. So it needs refinement. And, as we will see, the necessary refinements are also sufficient for avoiding a recent triviality result for Skyrms' Thesis due to Williams (2012).

The case for contextualism about counterfactuals is strong. Consider Quine's famous example:

proof of (7) from my formulation of the Principal Principle in footnote 12.

- (13) If Caesar had been in command in North Korea, he would have used the atom bomb.
- (14) If Caesar had been in command in North Korea, he would have used catapults.

It is easy to imagine a context in which we accept (13). It is also easy to imagine a context in which we accept (14). But we cannot imagine a context in which both (13) and (14) are acceptable. A natural explanation of these facts appeals to context-sensitivity. When we're holding fixed twentieth-century military technology, an utterance of (13) expresses a proposition that is true and an utterance of (14) expresses a proposition that is false. When holding fixed Caesar's actual competence with atomic weapons, the situation is reversed: (14) expresses a proposition that is true and (13) a proposition that is false.

To layer context-sensitivity on top of the Conditional Principal Principle, we need to do two things, both of which will be familiar from when we layered context-sensitivity on top of Stalnaker's Thesis. First, we need to add contextual parameters. Both the counterfactual conditional and the information that is held fixed need to be indexed to a context. Second, we must *coordinate* these two contextual parameters, just as we saw with the Local Thesis—the counterfactual conditional proposition must be indexed to the same context as the information that is held fixed.

Note that what's held fixed doesn't depend purely on context, but also on the antecedent of the conditional. As we've seen, when we evaluate a counterfactual whose antecedent concerns a particular period of time, we hold fixed a broad range of facts about history before that time, but not after. Consider an example from Dorr (2016). Suppose John has had breakfast every day this year. You say:

- (15) If John had forgotten to have breakfast on Tuesday, that would have been the first time this year.

To evaluate (15), I hold fixed history before Tuesday—that John had breakfast on Monday, that he had breakfast on Sunday, and so forth. But plainly I do not hold fixed that he had breakfast on Tuesday. Now imagine that you had said (16) instead of (15):

- (16) If John had forgotten to have breakfast on Wednesday, that would have been the first time this year.

In that case, I *would* have held fixed that John had breakfast on Tuesday, and I would have assented to (16).¹⁵

¹⁵There are other examples of antecedent-relativity. Take, for instance, Morgenbesser's coun-

I will write $E_c^-(A)$ to refer to the information that is held fixed in context c when we are evaluating a counterfactual with antecedent A . I will write $A \Box \rightarrow_c B$ for the proposition expressed by the counterfactual in context c . I will continue to assume that P_o is the uniquely rational initial credence function. With this notation, the Local Conditional Principal Principle is as follows.

The Local Conditional Principal Principle

$$P_o(A \Box \rightarrow_c B | E_c \wedge Ch = \pi) = \pi(B | A \wedge E_c^-(A))$$

To illustrate, suppose that you are the speaker of a certain context c . The Local Conditional Principle says that, if you're rational, then upon learning that the ur-chance function is π , the credence that you assign to your counterfactual—the proposition expressed by the counterfactual, relative to your context c —is equal to the chance, relative to π , of B conditional on A *and all of the information held fixed in c , relative to antecedent A* .

Thanks to this contextual coordination, the Local Conditional Principle is not subject to a recent triviality proof due to Williams (2012). (Note that my presentation of Williams' argument differs from his own presentation; Williams' original argument targets Skyrms' Thesis, but I am interested in exploring how a version of it might be used to refute the Conditional Principal Principle. Although the details differ, the basic strategies behind the arguments are the same.) Consider a rational subject in context c who has no evidence, and thus, is not holding any evidence fixed. The Principal Principle, applied to our subject in c , entails (a) below (where $A \Box \rightarrow_c B$ is the proposition expressed by the counterfactual in c):

$$(a) P_o(A \Box \rightarrow_c B | Ch = \pi) = \pi(A \Box \rightarrow_c B)$$

Since nothing is being held fixed in c , the Local Conditional Principal Principle entails:

$$(b) P_o(A \Box \rightarrow_c B | Ch = \pi) = \pi(B | A)$$

Notice that (a) and (b) together entail (c):

$$(c) \pi(A \Box \rightarrow_c B) = \pi(B | A)$$

terfactual, repeated below:

(10) If I had accepted the bet, I would have won.

When we evaluate (10), we hold fixed the fact that the coin landed heads. But not necessarily when we evaluate (17):

(17) If I had flipped the coin with a different hand, I would have won the bet.

Now, as Williams observes, *that* equation looks a lot like Stalnaker's Thesis. The indicative conditional has been replaced with a counterfactual and the rational credence function with an objective chance function. But the Lewisian triviality results that refute Stalnaker's Thesis do not presuppose any particular interpretation of the conditional operator, nor do they depend on any particular interpretation of probability. Perhaps, then, we can use a version of Lewis's argument to refute the Conditional Principal Principle.

There are two critical lemmas in Lewis's argument, stated in terms of chance and counterfactuals below:

Lemma 1. $\pi(A \Box\rightarrow_c B|B) = 1$

Lemma 2. $\pi(A \Box\rightarrow_c B|\neg B) = 0$

If we can derive these two lemmas from the Local Conditional Principal Principle and the Principal Principle, then we can use Lewis's reasoning to derive the absurd conclusion that, if $\pi(B) > 0$ and $\pi(\neg B) > 0$, then $\pi(B|A) = \pi(B)$. That is, if the initial chance function π assigns positive probability to B , and positive probability to $\neg B$, then B is probabilistically independent of A , relative to π .¹⁶ Fortunately, Lemma 1 and Lemma 2 don't follow from the Local Conditional Principal and the Principal Principle. To see why not, consider how we might try to derive Lemma 1, following the argument in (a) to (c) above. The first step would be to obtain (a') from the Principal Principle:

$$(a') P_o(A \Box\rightarrow_c B|Ch = \pi \wedge B) = \pi(A \Box\rightarrow_c B|B)$$

The second step would be to obtain (b') from the Local Conditional Principle (in a moment we'll see that this is the step that's blocked):

$$(b') P_o(A \Box\rightarrow_c B|Ch = \pi \wedge B) = \pi(B|A \wedge B) = 1$$

The third step would be to derive Lemma 1 from (a') and (b').

The problem with this argument is that the Local Conditional Principal Principle does not entail (b'). The Local Conditional Principle requires the counterfactual proposition $A \Box\rightarrow_c B$ to be coordinated with the information that is held

¹⁶Remember that (c) says that $\pi(A \Box\rightarrow_c B) = \pi(B|A)$. So if we can show that $\pi(A \Box\rightarrow_c B) = \pi(B)$, we can conclude that $\pi(B|A) = \pi(B)$. Here is the proof of $\pi(A \Box\rightarrow_c B) = \pi(B)$ from Lemma 1 and Lemma 2.

$$\begin{aligned} \pi(A \Box\rightarrow_c B) &= \pi(A \Box\rightarrow_c B|B) \times \pi(B) + \pi(A \Box\rightarrow_c B|\neg B) \times \pi(\neg B) \\ &= 1 \times \pi(B) + 0 \times \pi(\neg B) \\ &= \pi(B) \end{aligned}$$

The step from the first line to the second uses the Law of Total Probability. The step from the second to the third uses Lemma 1 and Lemma 2.

fixed in context c , relative to antecedent A . Thus, it requires the probability of $A \Box \rightarrow_c B$, conditional on $Ch = \pi$, to be equal to $\pi(B|A \wedge B)$ *only if B is held fixed in context c* . But, by hypothesis, B is not held fixed in c —nothing is held fixed in c .

My contextualist defense of the Local Conditional Principal Principle mirrors the contextualist defense of the Local Thesis. The Local Thesis escapes Lewisian triviality by requiring the indicative conditional to be indexed to the same context as the subject's evidence. The Local Conditional Principle escapes Williams' triviality argument by requiring the counterfactual conditional to be indexed to the same context as the evidence that the subject is holding fixed. Both principles should be seen as part of a unified, contextualist approach to the probabilities of conditionals.

6 A Sketch of a Theory of Conditionals

My goal when I started this paper was to derive a plausible, contextualist-friendly version of Skyrms' Thesis from a plausible, contextualist-friendly version of Stalnaker's Thesis and a plausible chance-deference norm. We now have the first three ingredients. Our chance-deference norm is the Principal Principle. Our contextualist-friendly version of Stalnaker's Thesis is the Local Thesis. And our contextualist-friendly version of Skyrms' Thesis is the Local Conditional Principal Principle. Here I turn to the final ingredient—the theory of conditionals.

I develop a theory on which all of the semantic differences between indicative conditionals and counterfactuals boil down to differences in what is held fixed in the context in which we evaluate the conditional. Following Stalnaker and others, I say that when we evaluate indicative conditionals, we hold fixed all of our knowledge.¹⁷ And, as we've seen in previous sections, when we evaluate counterfactuals, we hold fixed a contextually-determined subset of our knowledge. Be-

¹⁷More carefully, when we evaluate indicative conditionals, we hold fixed all of the information that is associated with our context. As I mentioned earlier, this will often be the speaker's knowledge, but sometimes it will be the knowledge of some other group or individual. Moreover, sometimes we may require something more demanding than knowledge, such as being known with certainty. This sort of flexibility is needed to account for cases like Adams's famous example:

(18) If Oswald didn't shoot Kennedy, some else did.

Plausibly, I know that Oswald didn't shoot Kennedy. Nevertheless, I am not holding this knowledge fixed when I evaluate the indicative conditional. Perhaps that's because we are only holding fixed what I know with certainty, and in any context in which I utter (18), I don't count as knowing that Oswald shot Kennedy with certainty. Or perhaps I do know this proposition with certainty but I am not *presupposing* that it is true (in the sense outlined in Stalnaker (2002)) for the purposes of the conversation. See Holguín (forthcoming) for extended discussion of cases like this.

cause indicative conditionals and counterfactual conditionals, on my view, differ only in what is held fixed, there is a systematic connection between the truth conditions for indicatives and the truth conditions for counterfactuals. Roughly, a counterfactual is true, relative to our present context, just in case the corresponding indicative conditional is true relative to the information we are holding fixed. In this section, I show how to implement this idea within a Stalnakerian *selection semantics* framework for conditionals.¹⁸

Stalnaker's theory is a uniform theory of conditionals. He states the truth conditions for conditionals in terms of a contextually-supplied *selection function* f . This is a function that takes a world w , and an antecedent A , and yields a world where A is true—the *selected A-world*, relative to w . Then Stalnaker says that, a conditional, whether indicative or subjunctive, is true at a world w just in case the selected antecedent-world, at w , is a consequent world.

To adopt a uniform theory of conditionals is not, of course, to say that indicatives and counterfactuals have the same meaning. They do not. Indicative conditionals are about epistemic possibilities; counterfactuals usually concern possibilities that are incompatible with our knowledge. Adams's famous minimal pair highlights the contrast:

- (9) If Oswald didn't shoot Kennedy, someone else did.
- (19) If Oswald hadn't shot Kennedy, someone else would have.

While (19) strikes us as a dubious claim about an alternative course of history, (9) looks straightforwardly true. The difference seems to be that while (9) is about how the world must have been, given what we now know, if Oswald wasn't the shooter, (19) is about how the world would have been had history taken a different course.

How do we account for these differences within a uniform theory? Stalnaker proposes that the selection function we use to evaluate indicative conditionals is subject to a special constraint: roughly, the selected antecedent-world must be an epistemically possible world. Here is a precise statement of the constraint on indicative selection functions.

Stalnaker's Constraint

If $A \cap E_c \neq \emptyset$, then if $w \in E_c$, $f_c(w, A) \in E_c$.

This says: If A is compatible with the information associated with context c , then for any world w in E_c , the selected A -world, at w , is also in E_c . Stalnaker's con-

¹⁸I choose Stalnaker's framework because, as van Fraassen (1976) and others have argued, Stalnaker's distinctive logic for conditionals—specifically, the principle of *Conditional Excluded Middle*—is needed if we want to vindicate the Local Thesis.

straint captures the sense in which, when assessing indicatives, we hold fixed all of our knowledge. To evaluate an indicative, we check whether the consequent is true at an epistemically possible antecedent world—that is to say, at a world where *everything we know is true*. Counterfactuals, Stalnaker says, are not subject to this constraint; their selection functions may reach outside the set of epistemically possible worlds.

I am going to account for the differences between indicatives and counterfactuals in a different way. It is clear that we cannot uphold Stalnaker’s Constraint, in its current form, for counterfactuals. One response is to dispense with the constraint altogether, as Stalnaker seems to suggest. But another response is to replace it with something else. On an abstract level, it is not hard to see what the replacement should be. For indicatives, Stalnaker’s Constraint requires that the selected antecedent-world be one where everything we’re holding fixed when we evaluate *the indicative* is true—a world where *everything* we know is true. For counterfactuals, the selected antecedent-world should be one where everything we’re holding fixed when we evaluate the *counterfactual* is true—a world where *some* of what we know is true, the part we’re holding fixed.

To implement this idea, I propose that a conditional, whether indicative or subjunctive, is evaluated relative to a *conditional information function*. This is a function s that takes an information state E and delivers a selection function that is *Stalnakerian relative to E* —a selection function that satisfies the constraints that Stalnaker imposes on indicative selection functions, relative to information state E . The constraint that matters for my purposes is a generalized version of Stalnaker’s Constraint, *Generalized Stalnaker’s Constraint* (I leave the others to a footnote):¹⁹

¹⁹The other four constraints are:

Success. $s(E)(w, A) \in A$ if $A \neq \emptyset$.

Minimality. $s(E)(w, A) = w$ if $w \in A$

Absurd. Where γ is an absurd world that makes all sentences true, $s(E)(w, A) = \gamma$ if and only if $A = \emptyset$

CSO. If $s(E)(w, A) \in B$ and $s(E)(w, B) \in A$, then $s(E)(w, A) = s(E)(w, B)$.

Success is needed to secure the validity of Identity, the principle that *if A , then A* is always true; See Mandelkern (2020) for extended discussion. Minimality secures the validity of Modes Ponens. Absurd secures a form of Conditional Non-Contradiction. CSO is needed to validate a host of intuitively compelling inference patterns. Note, however, that there is a tension between CSO and the Local Thesis. We can use an argument due to Stalnaker (1976) to show that the Local Thesis trivializes if we assume CSO. But importantly, Stalnaker’s proof relies on instances of the Local Thesis that involve conditionals with conditional antecedents. There are two possible responses to this argument. One is to dispense with CSO; this strategy is advocated by Bacon (2015). A different response is to reject the fully general version of the Local Thesis and replace it with a version that is restricted to conditionals with non-conditional antecedents. Importantly, if we

Generalized Stalnaker's Constraint

If $E \cap A \neq \emptyset$, then for all $w \in E$, $s(E)(w, A) \in E$.

The only difference between indicative and subjunctive conditionals, on my view, is that they supply different *arguments* to the conditional information function. For an indicative conditional, the argument to the conditional information function is E_c , the set of worlds compatible with everything we know. This gives us the following semantic entry, which is roughly equivalent to Stalnaker's own theory of indicative conditionals:

Indicative Selection Semantics

$$\llbracket A > B \rrbracket^{c,w,s} = 1 \text{ iff } s(E_c)(w, A) \in B$$

This says: An indicative conditional is true at a world w , relative to a context c and conditional information function s , just in case s takes E_c —the information associated with c —to a selection function that takes w and the antecedent A to a world where the consequent B is true. The selection function is Stalnakerian relative to E_c so it satisfies the Generalized Stalnaker's Constraint relative to E_c . This means that we evaluate an indicative conditional by checking whether the consequent holds at an antecedent world that is compatible with everything we know.

For counterfactuals, the informational argument to the conditional information function is the set of worlds consistent with what we are holding fixed, relative to the antecedent of the counterfactual.²⁰ (Remember that what we hold fixed for counterfactuals varies by antecedent.) The semantic entry is as follows:

Counterfactual Selection Semantics

$$\llbracket A \Box \rightarrow B \rrbracket^{c,w,s} = 1 \text{ iff } s(E_c^-(A))(w, A) \in B$$

This says: A counterfactual is true at a world w , relative to a context c and conditional information function s , just in case s takes $E_c^-(A)$ —the information that we hold fixed, relative to antecedent A —to a selection function that takes w and the antecedent A to a world where the consequent B is true. The selection function is Stalnakerian relative to $E_c^-(A)$ so it satisfies the Generalized Stalnaker's Constraint relative to $E_c^-(A)$. This means that we evaluate a counterfactual con-

accept only the restricted version of the Local Thesis, then we can only derive a restricted version of the Local Conditional Principle. If, instead, we reject CSO and accept the Local Thesis in full generality, then we can derive the Local Conditional Principle in full generality. I am inclined towards the second strategy, but I do not have the space to offer a full defense of that choice here. Note that the remarks that I make in §8 about tenability apply only to the restricted version of the Local Thesis.

²⁰I am indebted to David Boylan for extensive discussion about the formal relationship between E_c and E_c^- in a Stalnakerian selection semantics. An important question is how to derive these meanings compositionally. See the conclusion for a brief discussion.

ditional by checking whether the consequent holds at an antecedent world that is compatible with everything we're holding fixed.

On my theory, both indicatives and counterfactuals are governed by Generalized Stalnaker's Constraint, and there are no other differences between the selection functions that we use to interpret the two kinds of conditional. As a result, there is a close connection between the truth conditions for indicatives and the truth conditions for counterfactuals. To make this connection precise, let me introduce some notation. Consider a context c . $E_c^-(A)$ is, to repeat, the set of worlds compatible with everything we're holding fixed in c , relative to antecedent A . Let c^- be a hypothetical context in which our information is characterized by $E_c^-(A)$. In other words, $E_c^-(A) = E_{c^-}$. My theory predicts:²¹

$$\llbracket A \Box \rightarrow B \rrbracket^{c,s} = \llbracket A > B \rrbracket^{c^-,s}$$

This says: The proposition expressed by the counterfactual $A \Box \rightarrow B$ relative to $\langle c, s \rangle$ is identical to the proposition expressed by the indicative $A > B$ relative to $\langle c^-, s \rangle$. In the next section, we will see that this fact plays a crucial role in deriving the Local Conditional Principal Principle from the Principal Principle and the Local Thesis.

We have my neo-Stalnakerian uniform theory conditionals in place. I will close this section by giving three brief arguments for my uniform theory of conditionals, on which both indicatives and counterfactuals are subject to Generalized Stalnaker's Constraint, and there are no other differences between indicative and counterfactual selection functions.²²

First is an abductive argument based on the main claims of this paper. We have good reason to believe that some version of Skyrms' Thesis is true. I argue that we can derive this principle from the Local Thesis, the Principal Principle, and the unified semantics that I propose, on which both indicatives and subjunctives are subject to Generalized Stalnaker's Constraint. This gives us some reason to believe that the premises of that derivation are true. Since one of the premises is my uniform theory of conditionals, we have some reason to believe that this uniform theory is right.

A second, closely related argument concerns the fact that the probability one assigns to a counterfactual is often equal to the probability that one assigned to

²¹*Proof.* Suppose $A \Box \rightarrow B$ is true relative $\langle c, w, s \rangle$. By the Counterfactual Selection Semantics, it follows that $s(E_c^-(A))(w, A) \in B$. Then, by the Indicative Selection Semantics and the definition of c^- , it follows that $A > B$ is true relative to $\langle c^-, w, s \rangle$.

Now suppose that $A > B$ is true relative to $\langle c^-, w, s \rangle$. By the Indicative Selection Semantics and the definition of c^- , it follows that $s(E_c^-(A))(w, A) \in B$. Then, by the Counterfactual Selection Semantics, it follows that $A \Box \rightarrow B$ is true relative $\langle c, w, s \rangle$.

²²Thanks to Harvey Lederman and Matt Mandelkern for discussion about these arguments.

the corresponding indicative conditional at an earlier time. Suppose I know the Lakers are playing the Clippers in the NBA semi-finals, and that whoever wins that series will go on to the NBA finals and play the Celtics. Before the series starts I am confident in the indicative conditional (20).

(20) If the Lakers beat the Clippers, they will win the NBA championship.

The series between the Lakers and the Clippers concludes and the Clippers have won. I now endorse the counterfactual:

(21) If the Lakers had beat the Clippers, they would have won the NBA championship.

The probability I now assign to the counterfactual (21) at the conclusion of the series matches the probability I assigned to the indicative (20) at the start of the series. My theory easily accounts for this observation. If we assume that the information I hold fixed when evaluating (21) is identical to my total evidence at the earlier time when evaluating (20), then, on my theory, the proposition I am evaluating now *just is* the proposition I was evaluating then. And if these are just the same propositions, then of course I assign them equal probability.

A final argument concerns *presupposition*. Contemporary research about presupposition starts from the idea that the presuppositions of a clause must be satisfied relative to their *local contexts*. The local context of an embedded clause is, very roughly, the information information that is *already available*—the information that we can draw on to evaluate the clause—in the course of processing the sentence. Schlenker (2009) develops an algorithm for calculating local contexts, which says, very roughly, that the local context for an embedded clause is the strongest proposition that you can add to that clause without changing the truth-value of the whole sentence at any world in the global context. Mandelkern and Ramoli (2017) have shown that, given Stalnaker’s semantics for indicatives, this algorithm rightly predicts that the local context for the antecedent of an indicative conditional is (using my notation) E_c , the set of worlds consistent with the information associated with the context. Stalnaker’s Constraint plays a critical role in their argument. For, together with Stalnaker’s other constraints on selection functions, Stalnaker’s Constraint entails that, for any $w \in E_c$, and any A compatible with E_c , $f(w, A) = f(w, A \cap E_c)$. And once this constraint is in place, it is not hard to show that adding E_c to the antecedent won’t change the truth-value of the conditional at worlds compatible with our information.

Now, the local context for the antecedent of a counterfactual clearly isn’t E_c . Counterfactual antecedents are often inconsistent with what we know. So their local contexts can’t contain all of our information. But they do seem to contain

some of our information. Take an example from Heim (1992). You and I both know that Mary went to the party. I'm wondering whether she attended with her partner John. You're pretty sure John didn't attend, and you say:

(22) If John had attended too, I would have seen him.

An utterance of (22) is predicted to be felicitous only if the presupposition of its antecedent—that a salient individual attended the party—is satisfied relative to its local context. Hence, our theory of local contexts had better predict that the local context of the antecedent of (22) entails that a salient person—in this case, Mary—attended the party. In light of examples like (22), a natural hypothesis is that the local context of the antecedent of a counterfactual is the set of worlds consistent with what we're holding fixed.²³ Generalized Stalnaker's Constraint will play a central role in deriving this prediction, just as we saw with indicatives. For again, together with the other constraints, Stalnaker's Constraint entails that, for any $w \in E_c^-(A)$, and any A compatible with $E_c^-(A)$, $f(w, A) = f(w, A \cap E_c^-(A))$. Once this constraint is in place, adding $E_c^-(A)$ to the antecedent won't change the truth-value of the conditional at worlds compatible with our information.

7 Deriving the Local Conditional Principal Principle

Now that I have outlined my theory of conditionals, I am ready to show how we can use that theory to derive the Local Conditional Principal Principle from the Local Thesis and the Principal Principle.²⁴

So that we have everything in front of us, here is the Local Conditional Principal Principle:

Local Conditional Principal Principle

$$P_o(A \Box \rightarrow_c B | E_c \wedge Ch = \pi) = \pi(B | A \wedge E_c^-(A))$$

Remember that P_o is the uniquely rational initial credence function; E_c is the information associated with context c ; and $E_c^-(A)$ is the information that is held fixed, relative to antecedent A .

²³Heim (1992) makes a similar suggestion. She says: '...the antecedent of a counterfactual is not really added to an *empty* context, but to one which is in some sense a *revision* of the common ground c . It results from c by suspending some of the assumptions in c ; i.e., it is a superset of c .'

²⁴See Moss (2013) for a derivation of a version of Skyrms' Thesis for *future-directed* subjunctive conditionals from the Principal Principle. Moss does not show how to extend her argument to the case of past-directed subjunctive conditionals. But, as she notes, this does not mean that her argument has no implications for past subjunctives. If the proposition expressed by an earlier utterance of a future-directed subjunctive is the very same proposition as the proposition expressed by a current utterance of a past-directed subjunctive, then constraints on credences in future-directed subjunctives will entail constraints on credences in past-directed subjunctives. In many ways, then, the project of this essay is quite friendly to Moss's framework.

My derivation of the Local Conditional Principal Principle will rely on three principles. We have already seen two of these principles—the Local Thesis and the Principal Principle, repeated below.

Local Thesis

$$P_o(A >_c B|E_c) = P_o(B|A \wedge E_c)$$

The Principal Principle

$$P_o(A|E \wedge Ch = \pi) = \pi(A|E)$$

The third is a principle we have not yet seen that concerns the relationship between E_c , our actual information, and $E_c^-(A)$, the information we hold fixed when we evaluate a counterfactual with antecedent A . That principle is:

Independence

$$P_o(A \Box\rightarrow_c B|E_c^-(A)) = P_o(A \Box\rightarrow_c B|E_c)$$

Independence says that the probability of $A \Box\rightarrow_c B$ conditional on $E_c^-(A)$ (the evidence we hold fixed in context c) is equal to the probability of $A \Box\rightarrow_c B$ conditional on E_c (our evidence in context c).

To see why this assumption is warranted, remember what $E_c^-(A)$ is supposed to represent. The information that you hold fixed when you evaluate a counterfactual with antecedent A is the information that you judge relevant to determining what would have happened if A had been true. The information that you do not hold fixed is information you do not judge relevant to determining what would have happened if A had been true. Recall the case of Kennedy's assassination. You hold fixed a broad range of facts about what happened before Oswald pulled the trigger—that Oswald acted alone, that he was not part of a conspiracy, and so forth. You do not hold fixed what happened after the assassination—that Oswald shot Kennedy, that nobody else shot Kennedy, or that Johnson assumed the presidency in 1963. Consider a rational subject who knows everything you're holding fixed, and nothing more—that is, a rational subject whose total evidence consists of everything you know about history before Oswald pulled the trigger, and nothing you know about history after. Independence says that the probability that this subject assigns to the counterfactual (9), *if Oswald hadn't shot Kennedy, nobody else would have*, equals the probability that you assign to the counterfactual. In other words, learning what you're not holding fixed—that is, learning that Oswald shot Kennedy, that nobody else shot Kennedy, that Johnson assumed the presidency in 1963, and so forth—should not change her view about the counterfactual (9). For if it did, then E_c (your total evidence) must know something that $E_c^-(A)$ (the evidence you're holding fixed) doesn't know and that

you judge relevant to determining what would have happened if Oswald hadn't shot Kennedy. But in that case you should have been holding it fixed!²⁵

To make things concrete, I will show how the derivation of the Local Conditional Principal Principle works for a specific example. Suppose that on Monday at noon you are deciding whether to buy a lottery ticket. For simplicity, I will assume that you know $Ch = \pi$. Suppose that you decide not to purchase the ticket. Later you are evaluating the counterfactual:

(24) If you had bought the ticket, you would have lost.

Let c be the context in which you are evaluating (24). E_c is your current information and, where Buy is the proposition that you buy the ticket, $E_c^-(Buy)$ is the information you hold fixed when evaluating (24). Let c^- be your context at noon, just before deciding not to buy the ticket. I will assume that the information associated with c^- just is the information that you hold fixed in c —specifically, everything you knew before deciding not to buy the ticket and nothing you have learned since.

Our derivation begins with an instance of the Principal Principle (where Win is the proposition that you win the lottery).

$$(a) P_o(Win|Buy \wedge E_c^-(Buy)) = \pi(Win|Buy \wedge E_c^-(Buy))$$

(a) says that your credence, at noon, that you win the lottery, conditional on buying the ticket, is equal to the initial chance of winning conditional on buying a ticket and your total evidence at noon. (This follows from the Principal Principle because we have stipulated that you know that the initial chance function is π .)

Next, observe that (b) follows from (a) and the Local Thesis:

$$(b) P_o(Buy >_{c^-} Win|E_c^-(Buy)) = \pi(Win|Buy \wedge E_c^-(Buy))$$

²⁵Independence places constraints on the relationship between E_c (your evidence) and $E_c^-(A)$ (what you hold fixed, relative to antecedent A). It might be helpful to look at a specific case in which Independence is satisfied. Often when you're evaluating a counterfactual $A \Box \rightarrow B$, the relationship between E_c and $E_c^-(A)$ is the following: $E_c = E_c^-(A) \cap (\neg A \wedge \neg B)$. That is, your current knowledge is the result of intersecting what you hold fixed with the negation of the antecedent and the negation of the consequent. Take the coin case. I decide not to flip a fair coin. I am evaluating the counterfactual (6), *if the coin had been flipped, it would have landed heads*. I hold fixed all of my knowledge *except* for my knowledge of the fact that I did not flip the coin and that, as a result, it did not land heads. If this case, Independence says (where $E_c^-(Heads)$ is the set of worlds consistent with what I hold fixed):

$$(23) P_o(Heads \Box \rightarrow_c Flip|E_c^-(Heads)) = P_o(Heads \Box \rightarrow_c Flip|E_c^-(Heads) \cap (\neg Flip \wedge \neg Heads))$$

In the next section, I will show that that this instance of Independence holds in van Fraassen's *Stalnaker-Bernoulli* models—the models that van Fraassen (1976) and Bacon (2015) use to establish the tenability of the Local Thesis.

Your credence, at noon, in $Buy >_{c^-} Win$ —the proposition expressed by the indicative conditional relative to your information at noon—is equal to the initial chance of winning conditional on buying a ticket and your total evidence at noon.

On page 22, I showed that the counterfactual $Buy \Box \rightarrow_c Win$ is the very same proposition as the indicative $Buy >_{c^-} Win$. This means that the probability of the former is always equal to the probability of the latter. Thus, (c) follows from (b):

$$(c) P_o(Buy \Box \rightarrow_c Win | E_c^-(Buy)) = \pi(Win | Buy \wedge E_c^-(Buy))$$

Your credence, at noon, in $Buy \Box \rightarrow_c Win$ —the proposition expressed by the counterfactual relative to your *present* context, after deciding not to buy the lottery ticket—is equal to the initial chance of winning conditional on buying and your total evidence at noon.

Next, we appeal to Independence, which says that the probability of $Buy \Box \rightarrow_c Win$, conditional on $E_c^-(Buy)$ —what you hold fixed—is equal to your credence $Buy \Box \rightarrow_c Win$, conditional on you E_c —your total evidence.²⁶ Applying Independence to (c) gives us:

$$(d) P_o(Buy \Box \rightarrow_c Win | E_c) = \pi(Win | Buy \wedge E_c^-(Buy))$$

We said that $E_c^-(Buy)$ —your evidence at noon—entails $Ch = \pi$. Since E_c entails $E_c^-(Buy)$, it follows that E_c also entails $Ch = \pi$. Thus, (d) entails (e):

$$(e) P_o(Buy \Box \rightarrow_c Win | E_c \wedge Ch = \pi) = \pi(B | A \wedge E_c^-(Buy))$$

And (e) is an instance of the Local Conditional Principal Principle.

We have shown that under the assumption that you know $Ch = \pi$, the Local Conditional Principal Principle follows from the Principal Principle, the Local Thesis, and Independence.²⁷

²⁶One way to secure Independence in this example is to assume that: $E_c = E_c^-(Buy) \cap \neg Buy$. (See §8 for explanation.) This assumption seems plausible given the setup of the case.

²⁷Here's how it goes when $E_c^-(Buy)$ doesn't 'know' the chance of Win conditional on Buy . Let c_{ch} be the context that results from updating the information in c, E_c , with the proposition $Ch = \pi$. I will assume that what's held fixed in this new context is the intersection of what's held fixed in c and $Ch = \pi$. So we want to show that:

$$P_o(Buy \Box \rightarrow_{c_{ch}} Win | E_c \wedge Ch = \pi) = \pi(Win | Buy \wedge (E_c^-(Buy) \wedge Ch = \pi))$$

We begin with an instance of the Principal Principle:

$$(a) P_o(Win | Buy \wedge (E_c^-(Buy) \wedge Ch = \pi)) = \pi(Win | Buy \wedge E_c^-(Buy))$$

Next, (a) entails (b):

$$(b) P_o(Win | Buy \wedge (E_c^-(Buy) \wedge Ch = \pi)) = \pi(Win | Buy \wedge (E_c^-(Buy) \wedge Ch = \pi))$$

The details of the derivation are somewhat involved, so let me take a moment to walk through it in a more informal way. Suppose you are evaluating a counterfactual with antecedent A and consequent B . Consider a subject whose total evidence is the evidence that you hold fixed (and, we will assume, who knows all of the chance facts). By the Principal Principle, her credence in B given A equals the conditional chance of B given A , and by the Local Thesis, her credence in B given A equals her credence in her indicative conditional—that is, the proposition expressed by the indicative conditional, relative to her information. Thus it follows that her credence in her indicative is equal to the conditional chance of B given A .

Now, according to the uniform semantics for conditionals that I have proposed, the proposition expressed by the indicative, relative to her information, is equivalent to the proposition expressed by the corresponding counterfactual $A \Box \rightarrow B$, relative to your context. This means that her credence in her indicative is equal to her credence in your counterfactual. But remember that she has all of the information that you have and that you judge relevant to evaluating the counterfactual. Thus, it stands to reason that *your* credence in your counterfactual should equal *her* credence in your counterfactual.

If your credence in your counterfactual is equal to her credence in your counterfactual, which, in turn, is equal to her credence in her indicative, then your credence in your counterfactual is equal to her credence in her indicative. And we have already seen that her credence in her indicative is equal to the conditional chance of B given A . So, putting everything together, it follows that your credence in your counterfactual is equal to the conditional chance of B given A , just as the Local Conditional Principal Principle requires.

(The reason that (a) entails (b) is that it is a consequence of the Principal Principle that the initial chance function knows that it is the initial chance function: $\pi(Ch = \pi) = 1$.)

Let c^- be any context such that the information associated with c^- is the information that is held fixed in c_{ch} . Then (b) and the Local Thesis entail:

$$(c) \quad P_o(Buy >_{c^-} Win | E_c^-(Buy) \wedge Ch = \pi) = \pi(Win | Buy \wedge (E_c^-(Buy) \wedge Ch = \pi))$$

By the theory of conditionals outlined in §6, (c) entails (d):

$$(d) \quad P_o(Buy \Box \rightarrow_{c_{ch}} Win | E_c^-(Buy) \wedge Ch = \pi) = \pi(Win | Buy \wedge (E_c^-(Buy) \wedge Ch = \pi))$$

And finally, (d) and Independence entail:

$$(e) \quad P_o(Buy \Box \rightarrow_{c_{ch}} Win | E_c \wedge Ch = \pi) = \pi(Win | Buy \wedge (E_c^-(Buy) \wedge Ch = \pi))$$

8 Looking Forward: Tenability

Bas van Fraassen (1976) and Andrew Bacon (2015) have shown that the Local Thesis is *tenable* within a Stalnakerian semantic framework. There are non-trivial models in which the Local Thesis holds for the Stalnaker conditional. On an abstract level, it is not hard to see how to extend these results to establish the tenability of the Local Conditional Principle. That principle follows from the Principal Principle, the Local Thesis, and Independence. So if there are non-trivial models in which all three of these principles hold, then there are non-trivial models in which the Local Conditional Principal Principle holds—that is, the Local Conditional Principal Principle is tenable. (Note: I am only going to talk about *simple conditionals* in this section—conditionals with non-conditional antecedents and consequents.)

Here is a simplified overview of van Fraassen’s *Stalnaker-Bernoulli* models.²⁸ We evaluate conditional sentences relative to *sequences* of worlds. The first world in the sequence represents all of the *non-conditional* propositions that are true at the sequence—that is, all of the facts that can be specified without mentioning conditionals. And the rest of the sequence represents the conditional facts. To construct a Stalnaker-Bernoulli model, we begin with a set of worlds $\mathcal{I}$, which I will take to be the set of worlds compatible with all of the non-conditional information in a given context. We define $O_{\mathcal{I}}$ as the set of all sequences of worlds in $\mathcal{I}$. For example, if $\mathcal{I} = \{w_1, w_2, w_3\}$, then:

$$O_{\mathcal{I}} = \{\langle w_1, w_2, w_3 \rangle, \langle w_1, w_3, w_2 \rangle, \langle w_2, w_1, w_3 \rangle, \langle w_2, w_3, w_1 \rangle, \langle w_3, w_1, w_2 \rangle, \langle w_3, w_2, w_1 \rangle\}$$

A non-conditional sentence A is true at a sequence just in case the first world in that sequence is an A -world. A conditional $A > B$ is true at a sequence just in case the first A -world in the sequence is also a B -world. Suppose, for instance, that A is true at w_1 and w_2 , but false at w_3 , and that B is true at w_1 , but false at w_2 and w_3 . Then the conditional $A > B$ is true at $\langle w_1, w_2, w_3 \rangle$, $\langle w_1, w_3, w_2 \rangle$, and $\langle w_3, w_1, w_2 \rangle$, and false at the other three sequences.

(Note that we can represent the information that a sequence of worlds carries using more familiar Stalnakerian machinery: specifically, a pair $\langle w, f \rangle$ con-

²⁸I am heavily indebted to lecture notes from Justin Khoo and Paolo Santorio for this presentation. See Khoo and Santorio, ‘Lecture Notes: Probabilities of Conditionals in Modal Semantics.’ The models in the main text make various simplifying assumptions. They assume that all worlds have the same probability. They also assume that the initial set of worlds $\mathcal{I}$ is finite. Finally, they assume that a conditional is true at a sequence just in case the first A -world in the sequence is a B -world. If we want to handle conditionals with conditional antecedents, this definition has to be amended.

sisting of a world (which specifies the non-conditional facts) and a selection function (which specifies the conditional facts). Take, for instance, the sequence $S_1 = \langle w_1, w_2, w_3 \rangle$. This sequence corresponds to the pair $\langle w_1, f \rangle$ where f is a selection function satisfying two constraints: (i) for any $w \in \mathcal{I}$: if $w \in A$, then $f(w, A) = w$; and (ii) for any $w \in \mathcal{I}$: if $w \notin A$, then $f(w, A) =$ the first A -world in S_1 . More generally, each sequence in $O_{\mathcal{I}}$ corresponds to a pair consisting of a world in $\mathcal{I}$ and a selection function that is Stalnakerian relative to $\mathcal{I}$ —that is, a selection function that satisfies Stalnaker’s constraints relative to $\mathcal{I}$.)

To model the probabilities of conditionals, van Fraassen provides a recipe for taking us from a probability function P defined over $\mathcal{I}$ to a probability function P' defined over $O_{\mathcal{I}}$. He shows that the resulting probability function (1) *extends* P in the sense that $P'(A) = P(A)$ for all non-conditional A , and (2) for simple conditionals, the probability of the conditional is the corresponding conditional probability (whenever the conditional probability is defined). Although van Fraassen himself is not explicit about the role of context-sensitivity, there is a natural way of interpreting his results within a contextualist framework. If you feed the construction E_c and P_c —the information associated with c , and the probability function associated with c , respectively—the construction will output an interpretation of the conditional $A >_c B$ and an extended probability function P'_c such that Stalnaker’s equation holds for $A >_c B$ relative to P'_c . This establishes the tenability of the Local Thesis for simple conditionals.

Because P' extends P , if we start with a probability function P that obeys the Principal Principle with respect to non-conditional sentences, then P' will also obey the Principal Principle with respect to non-conditional sentences. So there is no obstacle to upholding both the Principal Principle (with respect to non-conditional sentences) and the Local Thesis.

My derivation of the Conditional Principal Principle also relied on the principle of Independence, repeated below.

Independence

$$P_o(A \Box \rightarrow_c B | E_c^-(A)) = P_o(A \Box \rightarrow_c B | E_c)$$

Independence can be reformulated as a principle about indicative conditionals. Let c^- be a hypothetical context in which our information is characterized by $E_c^-(A)$. On my theory, the proposition expressed by the counterfactual, relative to context c , is identical to the proposition expressed by the indicative conditional, relative to context c^- . So Independence becomes:

Indicative Independence

$$P_o(A >_{c^-} B | E_c^-(A)) = P_o(A >_c B | E_c)$$

Stated as a principle about indicatives, it is easy to show that certain instances of this principle are consistent with the Local Thesis. Indeed, we can show that a special case of the principle is entailed by the Local Thesis, given a Stalnakerian semantics for the conditional. Let $E_c = E_c^-(A) \cap \neg A$, and assume that $P_o(\neg A | E_c^-(A))$. Then Independence says that the proposition expressed by the indicative conditional, relative to c^- , is probabilistically independent of the negation of its antecedent, relative to the information in c^- . And this fact is a well-known consequence of the Local Thesis, assuming Stalnaker's selection semantics for conditionals.²⁹

Of course, showing that Independence holds for this particular choice of E_c and $E_c^-(A)$ does not show much. That's because it's not normally the case that $E_c = E_c^-(A) \cap \neg A$. Suppose I know that neither you nor your partner went to the party last night, and that you often attend parties together. When I evaluate:

(25) If you had gone to the party, your partner would have gone to the party.

I don't hold fixed that you didn't go, but I also don't hold fixed that your partner didn't go. So we don't get to my present knowledge by intersecting what I hold fixed with the proposition that you did not go to the party—the negation of the counterfactual's antecedent. We get to my present knowledge by intersecting what I hold fixed with the proposition that *neither you nor your partner* attended the party—the conjunction of the negation of the antecedent and the negation of the consequent.

Often our current knowledge results from intersecting what we're holding fixed, relative to some proposition A , with some proposition Q that is *stronger* than $\neg A$. It would be good to show that, for any such Q , the conditional $A >_{c^-} B$ is probabilistically independent of Q relative to $P_o(\cdot | E_c^-(A))$. Formally, where Q is any proposition entailing $\neg A$ such that $P_o(Q | E_c^-(A)) > 0$:

$$P_o(A >_{c^-} B | E_c^-(A)) = P_o(A >_{c^-} B | E_c^-(A) \cap Q)$$

²⁹Let $P_{c^-} = P_o(\cdot | E_c^-(A))$. Note that $P_{c^-}(A >_{c^-} B | A) = P_{c^-}(A >_{c^-} B)$ just in case $P_{c^-}(A >_{c^-} B | \neg A) = P_{c^-}(A >_{c^-} B)$. So it suffices to show that $P_{c^-}(A >_{c^-} B | A) = P_{c^-}(A >_{c^-} B)$.

$$\begin{aligned} P_{c^-}(A >_{c^-} B | A) &= \frac{P_{c^-}(A >_{c^-} B \wedge A)}{P_{c^-}(A)} \\ &= \frac{P_{c^-}(A \wedge B)}{P_{c^-}(A)} \\ &= P_{c^-}(B | A) \\ &= P_{c^-}(A >_{c^-} B) \end{aligned}$$

The step from the second line to the third line relies on Stalnaker's logic for the conditional—specifically, the principle of *Strong Centering*, which says that $(A > B \wedge A)$ is true just in case $(A \wedge B)$ is true. Strong Centering follows from Minimality given Stalnaker's assumption that there is always a unique selected antecedent-world.

Stated in terms of counterfactuals, this becomes:

$$P_o(A \Boxrightarrow_c B | E_c^-(A)) = P_o(A \Boxrightarrow_c B | E_c^-(A) \cap Q)$$

This says: The probability of the proposition expressed by the counterfactual in c , conditional on what's held fixed in c , is equal to the probability of the proposition expressed by the counterfactual in c , conditional on the intersection of what's held fixed in c and Q , where Q is any proposition that entails $\neg A$ (and is assigned positive probability by $P_o(\cdot | E_c^-(A))$). In the appendix, I show that this fact holds in our simplified Stalnaker-Bernoulli models. This establishes the tenability of many plausible instances of Independence with respect to these models.

Let's look at a simple example. Suppose that $E_c^-(A) = \{w_1, w_2, w_3, w_4\}$. Suppose that A is true at w_1 and w_2 , but false at w_3 , and w_4 . And suppose that B is true at w_1 and w_3 , but false at w_2 and w_4 . If we assume that each world has equal probability, then the probability $A >_{c-} B$, relative to $E_c^-(A)$, is equal to the proportion of sequences of $O_{E_c^-(A)}$ whose first A -world is B -world. It is easy to verify that $O_{E_c^-(A)}$ contains 24 sequences and that 12 of these sequences are such that their first A -world is a B -world. So the probability of the conditional $A >_{c-} B$, relative to $E_c^-(A)$, is $1/2$.

Now consider a $\neg A$ -entailing factual proposition—say, $\neg A \wedge \neg B$. This proposition is true at all and only the sequences in $O_{E_c^-(A)}$ whose first world is w_4 :

$$\begin{aligned} &\langle w_4, w_1, w_2, w_3 \rangle, \langle w_4, w_1, w_3, w_2 \rangle, \langle w_4, w_2, w_1, w_3 \rangle \\ &\langle w_4, w_2, w_3, w_1 \rangle, \langle w_4, w_3, w_1, w_2 \rangle, \langle w_4, w_3, w_2, w_1 \rangle \end{aligned}$$

There are six sequences in total beginning with w_4 , three of which are such that their first A -world is a B -world. So the conditional is true at half of sequences beginning with w_4 , which is to say that the probability of $A >_{c-} B$, conditional on $\neg A \wedge \neg B$, is again $1/2$. The conditional is probabilistically independent of $\neg A \wedge \neg B$, relative to $E_c^-(A)$. Formally we have shown that:

$$P_o(A >_{c-} B | E_c^-(A)) = P_o(A >_{c-} B | E_c^-(A) \cap (\neg A \wedge \neg B))$$

Stated in terms of counterfactuals, this says:

$$P_o(A \Boxrightarrow_c B | E_c^-(A)) = P_o(A \Boxrightarrow_c B | E_c^-(A) \cap (\neg A \wedge \neg B))$$

The same will be true of any factual Q entailing $\neg A$. Zoom in on the set of sequences in $O_{E_c^-(A)}$ that make Q true. The proportion of sequences in this new set whose first A -world is B -world will be equal to $1/2$.

We have seen that there are Stalnaker-Bernoulli models—of the simplified

variety that I have presented in this section—in which each of the Principal Principle, the Local Thesis, and Independence holds. In each of these, the Local Conditional Principal Principle holds, too. I leave a full tenability proof—one that dispenses with the simplifying assumptions I have made here—to future research.

9 Conclusion

The project of this article has been to sketch a neo-Stalnakerian, uniform theory of conditionals that allows us to derive a plausible, contextualist-friendly version of Skyrms' Thesis (the Local Conditional Principal Principle) from a plausible, contextualist-friendly version of Stalnaker's Thesis (the Local Thesis) and a plausible chance-deference norm (the Principal Principle). I close by outlining two questions for future research.

One question is about chance. I used the Principal Principle to derive a version of Skyrms' Thesis. But the Principal Principle is just one candidate chance-deference norm—one way of formalizing the claim that one's credences ought to be guided by objective chances. Other candidates are Hall's *New Principle* and Dorst's *Trust Principle*.³⁰ One area of future research involves determining whether we can derive counterfactual versions of these principles from the Local Thesis and their non-conditional counterparts.

Another question is about semantics. I've given a semantics for counterfactuals on which the meaning of a counterfactual is closely related to the meaning of an indicative conditional. A full defense of this theory would require showing how to derive this meaning compositionally. I am optimistic about the prospects of this project if one adopts a certain approach to the role of tense in counterfactuals. Let me conclude by saying something about the approach I favor. Plausibly, a 'would'-conditional is composed of a 'will'-conditional under a past tense operator. (For defense of this claim, see, for example, Ippolito (2013).) There are two main hypotheses about what this past tense operator does, one of which I take to be particularly promising: the *past-as-modal* view, on which the past tense is interpreted as a modal.³¹ (The alternative approach is the *past-as-past* view, on which the past tense has its usual temporal meaning in counterfactuals. See Khoo (2015) for a defense of this approach.) Inspired by Schulz (2014), one understanding of the past-as-modal view says that the past tense shifts the information state relative to which we interpret the embedded indicative condi-

³⁰See Kevin Dorst (2020). Note that Dorst's principle is formulated as a principle about deference to one's own (future) evidence, but Ben Levinstein (ms) advocates adopting Trust for deference to chance.

³¹See, for example, Iatridou (2000) and Schulz (2014) for defenses of this approach.

tional. Specifically, the past tense operator shifts the information state from the set of worlds consistent with our actual information to the set of worlds consistent with what we hold fixed. The result is that a counterfactual is true, relative to our present context, just in case the corresponding indicative conditional is true relative to the information we're holding fixed when we evaluate the counterfactual. This, of course, is exactly what my uniform theory of conditionals predicts: the only difference between indicative conditionals and counterfactuals is the information that is held fixed when we evaluate the conditional.³²

³²Thanks to Zach Barnett, Fabrizio Cariani, Harvey Lederman, and Daniel Rothschild for helpful conversations. Thanks to Kevin Dorst, Branden Fitelson, Simon Goldstein, Sarah Moss, Bernhard Salow, Paolo Santorio, and Robbie Williams for feedback on earlier drafts. Special thanks to David Boylan, Melissa Fusco, Arc Kocurek, Matt Mandelkern, and Milo Phillips-Brown for extensive feedback throughout the project.

10 Appendix

Begin with some notation and terminology.

- Let X be any finite set of worlds.
- Let S be the set of all sequences of worlds in X .
- Let $A >_S B$ be the set of all sequences in S whose first A -world is a B -world.
- Let S_x be the set of sequences in S whose first element is x .
- For any $x \in X$, let X_{-x} be $X - \{x\}$. Let S_{-x} be the set of all sequences of worlds in X_{-x} .

We begin by showing:

Claim 1. For any x such that $x \in X$ and $x \notin A$: $\frac{|A >_S B|}{|S|} = \frac{|A >_{S_x} B|}{|S_x|}$

We will show Claim 1 by proving two sub-claims that together entail Claim 1. Those claims are (1) and (2) below, for any x such that $x \in X$ and $x \notin A$:

1. $\frac{|A >_S B|}{|S|} = \frac{|A >_{S_{-x}} B|}{|S_{-x}|}$
2. $\frac{|A >_{S_{-x}} B|}{|S_{-x}|} = \frac{|A >_{S_x} B|}{|S_x|}$

Proof of (1).

Let $|X_{-x}| = n$. Then $|S_{-x}| = n!$ and $|S| = (n + 1)!$. We know that for any sequence $o \in S_{-x}$, there are exactly $n + 1$ sequences in S that preserve the order of the elements in o . Roughly, that is because, for any $o \in S_{-x}$, there are $n + 1$ places where we can insert x : at the beginning of the sequence, after the first element, after the second element, and so forth. For example, consider $o = \langle w_1, w_2, w_3, \dots, w_n \rangle$. There are $n + 1$ sequences in S that preserve the order of the elements of o :

$$\langle x, w_1, w_2, w_3, \dots \rangle$$

$$\langle w_1, x, w_2, w_3, \dots \rangle$$

$$\langle w_1, w_2, x, w_3, \dots \rangle$$

$$\langle w_1, w_2, w_3, x, \dots \rangle$$

And so forth. For each of these sequences, the first A -world in the sequence will be a B -world just in case the first A -world in o —the original sequence—is a B -world. So, for any $o \in A >_{S-x} B$, there are exactly $n+1$ sequences in $A >_S B$. This means that we can reason as follows:

$$\begin{aligned}
\frac{|A >_S B|}{|S|} &= \\
\frac{|A >_{S-x} B|(n+1)}{(n+1)!} &= \\
\frac{|A >_{S-x} B|(n+1)}{(n+1)(n!)} &= \\
\frac{|A >_{S-x} B|}{n!} &= \\
\frac{|A >_{S-x} B|}{|S-x|} &
\end{aligned}$$

Proof of (2).

The sequences in S_x are the same as the sequences in S_{-x} , except that x is tacked on to the beginning of each. Let $f : \langle w_1, \dots, w_n \rangle \mapsto \langle x, w_1, \dots, w_n \rangle$. Then f is a bijection from S_{-x} to S_x as well as from $A >_{S-x} B$ to $A >_{S_x} B$. So Claim 2 immediately follows: $\frac{|A >_{S-x} B|}{|S-x|} = \frac{|A >_{S_x} B|}{|S_x|}$

We have shown Claim 1. Next we want to show Claim 2 (where $|S_Q|$ is the set of sequences in S whose first world is a Q -world and $|A >_{S_Q} B|$ be the set of sequences in S_Q whose first A -world is a B -world):

Claim 2. Where Q is any proposition that entails $\neg A$: $\frac{|A >_S B|}{|S|} = \frac{|A >_{S_Q} B|}{|S_Q|}$

Proof of Claim 2.

Let $Q = \{x_1, \dots, x_n\}$. We know:

- $|S_Q| = |S_{x_1}| + \dots + |S_{x_n}|$
- $|A >_{S_Q} B| = |A >_{S_{x_1}} B| + \dots + |A >_{S_{x_n}} B|$

Then we can reason as follows:

$$\begin{aligned}
& \frac{|A >_{S_Q} B|}{|S_Q|} &= \\
& \frac{|A >_{S_{x_1}} B| + \dots + |A >_{S_{x_n}} B|}{|S_{x_1}| + \dots + |S_{x_n}|} &= \\
& \frac{|A >_{S_{x_1}} B|(n)}{|S_{x_1}|(n)} &= \\
& \frac{|A >_{S_{x_1}} B|}{|S_{x_1}|} &= \\
& \frac{|A >_S B|}{|S|} &=
\end{aligned}$$

The third line follows from the second because (a) $|S_{x_1}| = |S_{x_2}| = \dots = |S_{x_n}|$ and (b) $|A >_{S_{x_1}} B| = |A >_{S_{x_2}} B| = \dots = |A >_{S_{x_n}} B|$. And Claim 1 secures the inference from the fourth line to the fifth line.

11 References

1. Bacon, Andrew (2015). 'Stalnaker's thesis in context.' *Review of Symbolic Logic* 8, pp. 131 - 163.
2. Bradley, Richard (2000). 'A Preservation Condition for Conditionals.' *Analysis* 60, pp. 219-222.
3. Dorst, Kevin (2020). 'Evidence: A Guide for the Uncertain.' *Philosophy and Phenomenological Research* 100, pp. 586-632.
4. Edgington, Dorothy (2008). 'Counterfactuals.' *Proceedings of the Aristotelian Society* 108, pp. 1-21.
5. von Fintel, Kai (1998). 'The presupposition of subjunctive conditionals.' In Orin Percus & Uli Sauerland (eds.), *The interpretive tract*, Vol. 25, pp. 25-44. Cambridge, Massachusetts: MIT Working Papers in Linguistics.
6. van Fraassen, Bas C. (1976). 'Probabilities of conditionals.' In *Foundations of probability theory, statistical inference, and statistical theories of science*, Springer, pp. 261-308.
7. Gillies, Anthony (2009). 'On the truth conditions for If (but not quite only If).' *Philosophical Review* 118, 325 - 349.
8. Hall, Ned (1994). 'Correcting the Guide to Objective Chance.' *Mind* 103, pp. 504-18.
9. Hall, Ned (2004). 'Two Mistakes about Credence and Chance.' *Australasian Journal of Philosophy* 82, pp. 93-111
10. Harper, William L., Robert Stalnaker & Glenn Pearce (1981). *Ifs: Conditionals, Belief, Decision, Chance, and Time*. D. Reidel Publishing Company, Dordrecht.
11. Holguin, Ben (2020). 'Knowledge in the Face of Conspiracy Conditionals.' *Linguistics and Philosophy*.
12. Iatridou, Sabine (2000). 'The Grammatical Ingredients of Counterfactuality.' *Linguistic Inquiry* 31, pp. 231 - 270.
13. Ippolito, Michela (2013). 'Subjunctive Conditionals: A Linguistic Analysis.' *Linguistic Inquiry Monograph* (Series 65), Cambridge: MIT Press.

14. Kaufmann, Stefan (2004). 'Conditioning against the grain: Abduction and indicative conditionals.' *Journal of Philosophical Logic* 33, pp. 583-606.
15. Kaufmann, Stefan (2005). 'Conditionals predictions: A probabilistic account.' *Linguistics and Philosophy* 28, pp. 181-231.
16. Kaufmann, Stefan (2009). 'Conditionals right and left: Probabilities for the whole family.' *Journal of Philosophical Logic* 38, pp. 1-53.
17. Khoo, Justin (2015). 'On Indicative and Subjunctive Conditionals.' *Philosopher's Imprint* 15, pp. 1-40.
18. Khoo, Justin (2020). 'The meaning of *If*.' Ms.
19. Levinstein, Ben (2020). 'Accuracy, Deference, and Chance.' Ms.
20. Lewis, David (1976). 'Probabilities of conditionals and conditional probabilities.' *Philosophical Review* 85, pp. 297-315.
21. Lewis, David (1980). 'A Subjectivist's Guide to Objective Chance.' In Jeffrey (ed.) *Studies in Inductive Logic and Probability*, Vol. 2, pp. 263-93. Berkeley: University of California Press.
22. Lewis, David (1994). 'Humean Supervenience Debugged.' *Mind* 103, pp. 473-490.
23. Mandelkern, Matthew (2019). 'If p, then p!' Ms.
24. McGee, Vann (1989). 'Conditional probabilities and compounds of conditionals.' *Philosophical Review* 98, 485-541.
25. Moss, Sarah (2013). 'Subjunctive Credences and Semantic Humility.' *Philosophy and Phenomenological Research* 87, pp. 251 - 278
26. Schulz, Katrin (2014). 'Fake tense in conditional sentences: A modal approach.' *Natural Language Semantics* 22, pp. 117 - 144.
27. Schulz, Moritz (2017). *Counterfactuals and Probability*. Oxford: Oxford University Press.
28. Schwarz, Wolfgang (2016). 'Subjunctive Conditional Probability.' Forthcoming in *Journal of Philosophical Logic*.
29. Skyrms, Bryan (1980). 'The Prior Propensity Account of Subjunctive Conditionals.' In Harper et al. (1981).

30. Stalnaker, Robert (1968). 'A Theory of Conditionals.' In Nicholas Rescher (ed.) *Studies in Logical Theory* (American Philosophical Quarterly Monographs 2), Oxford: Blackwell. pp. 98-112.
31. Stalnaker, Robert (1975). 'Indicative conditionals.' *Philosophia* 5, pp. 269-286.
32. Stalnaker, Robert (2002). 'Common ground.' *Linguistics and Philosophy* 25, pp. 701-721.
33. Stalnaker, Robert and Richard Jeffrey (1994). 'Conditionals as random variables.' In Ellery Eells, Brian Skyrms & Ernest W. Adams (eds.) *Probability and Conditionals: Belief Revision and Rational Decision*, Cambridge: Cambridge University Press.
34. Williams, J. Robert G. (2012). 'Counterfactual Triviality: A Lewis-Impossibility Argument for Counterfactuals.' *Philosophy and Phenomenal Research* 85, pp. 648 - 670