THE REVIEW OF SYMBOLIC LOGIC
Volume 5, Number 1, March 2012

A PROBABILISTIC SEMANTICS FOR COUNTERFACTUALS. PART B

HANNES LEITGEB
Ludwig-Maximilians-University Munich

In memory of Horacio Arló-Costa

Abstract. This is part B of a paper in which we defend a semantics for counterfactuals which is probabilistic in the sense that the truth condition for counterfactuals refers to a probability measure. Because of its probabilistic nature, it allows a counterfactual 'if A then B' to be true even in the presence of relevant 'A and not B'-worlds, as long such exceptions are not too widely spread. The semantics is made precise and studied in different versions which are related to each other by representation theorems. Despite its probabilistic nature, we show that the semantics and the resulting system of logic may be regarded as a naturalistically vindicated variant of David Lewis' truth-conditional semantics and logic of counterfactuals. At the same time, the semantics overlaps in various ways with the non-truth-conditional suppositional theory for conditionals that derives from Ernest Adams' work. We argue that counterfactuals have two kinds of pragmatic meanings and come attached with two types of degrees of acceptability or belief, one being suppositional, the other one being truth based as determined by our probabilistic semantics; these degrees could not always coincide due to a new triviality result for counterfactuals, and they should not be identified in the light of their different interpretation and pragmatic purpose. However, for plain assertability the difference between them does not matter. Hence, if the suppositional theory of counterfactuals is formulated with sufficient care, our truth-conditional theory of counterfactuals is consistent with it. The results of our investigation are used to assess a claim considered by Hawthorne and Hájek, that is, the thesis that most ordinary counterfactuals are false.

This paper continues the project of developing a semantics for counterfactuals in terms of conditional chance which was started in "A Probabilistic Semantics for Counterfactuals. Part A." According to the semantics,

If the match were struck, it would light

is true or approximately true at a world w if and only if the conditional chance of the match lighting if struck is very high at w, where 'very high' is either analyzed as (i) 'equal to 1', or as (ii) 'close to 1' with 'close to' being a vague term (Section 2.3 of this paper will explain why this interpretation is available), or as (iii) 'greater than or equal to some (fixed) real number $1 - \alpha$' (with $0 \leq \alpha \leq \frac{1}{2}$). In cases (i) and (ii) we spoke of *truth simpliciter*, whereas we considered case (iii) to give us *approximate truth* (to degree $1-\alpha$). Either way, as explained in part A, we aim at

1. a scientifically plausible and simple probabilistic semantics for subjunctive conditionals,
2. which nevertheless turns out to stay quite close to the Lewis–Stalnaker semantics for conditionals,
3. such that the truth of $A \mathrel{\Box\!\!\rightarrow} B$ allows for exceptions, that is, where $\Box\!\!\rightarrow$ comes with an implicit ceteris paribus clause.

Received: August 10, 2010.

doi:10.1017/S1755020311000165

1, 2, and 3 are what we promised to have achieved by the end of part B. Although the semantics is a truth-conditional semantics, we also want to show that it builds on, and in some ways extends, Ernest Adams' non-truth-conditional suppositional theory for conditionals.

This is the plan for part B. One of the differences between the new semantics and Lewis' is that the so-called Centering Axioms[1] cease to count as logical truths: in Section §1 we will show why this is so, why it is not so bad, and what remnants of Centering the new semantics offers to the unconvinced. Along the way, we will also explain why and how our semantics coincides with Stalnaker's classic theory given the additional assumption of Counterfactual Determinism. In Section §2 we will suggest a second probabilistic semantics for counterfactuals. This is not because we are unhappy with the first one, but because it will throw some interesting light on the original semantics: although the second semantics will *look* quite different from the first one, it will follow from a new representation theorem for conditional probability measures that the two semantics are—in a sense explained by the theorem—*equivalent*. Since the metaphysical presuppositions of the second semantics will seem to be almost Lewisian, whereas the former semantics looks much more scientific, we will argue that in light of this representation result our probabilistic semantics may actually be understood as a naturalization of Lewis' semantics. Moreover, we will explain why validity in the sense of our truth-conditional semantics is so close to validity according to Adams' non-truth-conditional theory, we will consider a probabilistic version of Lewis' Limit Assumption, and we will see under which conditions our conditional operator may be represented in terms of Lewis'. Finally, Section §3 will consider an application of the new probabilistic semantics: As Hawthorne (2005) and Hájek (unpublished) have asked recently, are most ordinary counterfactuals false? We will see that the answer recommended by our probabilistic semantics is a negative one, at least as long as one is satisfied with the approximate truth of ordinary counterfactuals, and we will discuss what this tells us about the merits and the limitations of the semantics. A short summary and an appendix with proofs conclude the paper.

Throughout part B, we will keep referring to, and relying on, what has been established in part A. Part A also includes various references to papers in which related probabilistic semantics for counterfactuals are being developed.

§1. What becomes of the centering axioms? We turn now to the logical differences between Lewis' semantics and our probabilistic semantics, as promised. If compared to the system V, which according to section 3 of part A is sound and complete with respect to our semantics, the only logical axioms (or rather axiom schemes) that are missing from the system VC that is adequate for Lewis' semantics are:

C1 Weak Centering: $(A \mathrel{\Box\!\!\rightarrow} B) \supset (A \supset B)$

C2 (Strong) Centering: $(A \wedge B) \supset (A \mathrel{\Box\!\!\rightarrow} B)$.

C1 by itself is called the 'Weak Centering axiom'; C1 and C2 taken together are called the 'Centering axioms'.

Why do they fail in the Popper function semantics that was stated in part A? The antecedent of C1 and the consequent of C2 express a property of chance functions at worlds; however, both the consequent of C1 and the antecedent of C2 express mere truth-functional constraints on the truth-values of A and B at worlds. Since there are no logically

[1] That is, instances of: $(A \mathrel{\Box\!\!\rightarrow} B) \supset (A \supset B)$ and $(A \wedge B) \supset (A \mathrel{\Box\!\!\rightarrow} B)$.

valid—or even universally true—bridge laws which relate chances of states or events at a world w to truth-values of statements which describe the corresponding states or events at that world w, neither of the material implications in C1 or C2 holds logically. In short: $\mathfrak{P}_w(\llbracket B \rrbracket \mid \llbracket A \rrbracket)$ is simply not tied necessarily to any particular distribution of truth-values for A and B in w.

It should be noted though that the emphasis here is on *logically* valid or universal bridge laws. For in special cases, conditional chances *do* entail certain distributions of truth-values: for consider any counterfactual of the form

$$A[t] \;\Box\!\!\rightarrow_{t_r}\; B[t']$$

with t_r being the reference point, and $t < t' < t_r$—here we use the terminology of section 4.3 of part A: in particular, $A[t]$ and $B[t']$ are meant to describe events that lie in the past if seen from the viewpoint of t_r at which the relevant conditional chance of the consequent proposition given the antecedent proposition is taken. Suppose $A[t]$ is false at w: then both Centering axioms are satisfied for trivial truth-functional reasons. Now assume $A[t]$ to be true at w: then the absolute chance $\mathfrak{P}_w^{t_r}(\llbracket A[t] \rrbracket \mid W)$ is 1 since $A[t]$ is true and about the past; if $B[t']$ is false at w, the antecedent of C2 is false again by the truth function of $\wedge$, and the antecedent of C1 is false because the absolute chance $\mathfrak{P}_w^{t_r}(\llbracket B[t'] \rrbracket \mid W)$ is 0 (since $B[t']$ is about the past again) and the Ratio formula for conditional probabilities applies; finally, if also $B[t']$ is true at w, then the consequent of C1 is true truth functionally, and since the absolute chance $\mathfrak{P}_w^{t_r}(\llbracket B[t'] \rrbracket \mid W)$ is 1 in this case it follows by the Ratio formula that $\mathfrak{P}_w^{t_r}(\llbracket B[t'] \rrbracket \mid \llbracket A[t] \rrbracket) = 1$, thus also the consequent of C2 is satisfied. Hence, full Centering applies in all cases *of this type*, that is, in the case of counterfactuals whose antecedent and consequent are about the past, relative to the reference point. Such counterfactuals are often regarded as the prototypical cases of counterfactuals in the stricter sense of the word that we did not presuppose in this paper, that is, subjunctive conditionals that are against the facts in view of their antecedents describing an event in the past which did not actually happen (and the consequents of which are about the past as well). For such counterfactuals there can be no instance of either of the Centering axioms which is false. However, even for those counterfactuals it is not clear whether this is so for purely logical reasons. And in all other cases of subjunctive conditionals, there is not even a guarantee that instances of C1 and C2 are true according to the Popper function semantics, let alone that they are logically true.

Is this a problem? Although full Centering is delivered by Lewis' semantics, quite a few philosophers have opposed C2, whether in the wake of Nozick's tracking analysis of knowledge, or from considerations on indeterminism (as in Bennett, 2003, sec. 92), or for some other reason. Indeed, it is at least questionable to assume that the mere presence of facts which are described by A and B is capable of establishing a counterfactual dependence of B on A—and, for that matter, of A on B—and even so by pure logic. So maybe it is a virtue rather than a vice that according to our semantics $A \wedge B$ does not make $A \;\Box\!\!\rightarrow\; B$ true. Moreover, $A \wedge B$ not entailing $A \;\Box\!\!\rightarrow\; B$ conforms to our understanding in this paper of $A \;\Box\!\!\rightarrow\; B$ as being exception tolerant: $A \wedge B$ just by itself should not be sufficient to establish $A \;\Box\!\!\rightarrow\; B$, since $A \wedge B$ taken by itself might merely be an exception.

On the other hand, the dual kind of reasoning—$A \wedge \neg B$ should not make $A \;\Box\!\!\rightarrow\; B$ false; or: the mere presence of facts which are described by A and $\neg B$ should not be able to break the counterfactual dependence of B on A; or: C1 is not logically true—is widely regarded as counterintuitive. In fact, C1 is sometimes taken to be logically constitutive of the subjunctive 'if-then'.[2] Under which conditions may C1 fail in our semantics? In order for C1 to be false at a world w, $A \boxright B$ would have to be true and $A \supset B$ would have to be false at w. In other words: the conditional chance of B given A at w would have to be (close to) 1 while $A \wedge \neg B$ is false at w. So $A \wedge \neg B$ would have to be exceptional with respect to w's own probabilistic standards, the logical possibility of which we are reluctant to exclude within our probabilistic approach to counterfactuals.

Why do people think favorably of Weak Centering? Presumably, because in its presence it becomes possible to derive Counterfactual Modus Ponens to be logically valid, that is,

- Assume A and $A \boxright B$,
- from the latter, by Weak Centering, it follows that $A \supset B$,
- therefore, by standard Modus Ponens, B,

and similarly Counterfactual Modus Tollens:

- Assume $\neg B$ and $A \boxright B$,
- from the latter, by Weak Centering, it follows that $A \supset B$,
- therefore, by standard Modus Tollens, $\neg A$.

Counterfactual Modus Ponens is slightly odd at least pragmatically in cases in which $A \boxright B$ is a counterfactual in the stricter sense of the word that is meant to implicate (though not logically imply) the falsity of its antecedent; for it applies to situations in which both such a counterfactual and the *truth* of its antecedent are presumed.[3] In contrast, Counterfactual Modus Tollens does not suffer from the same pragmatic worries, as it gets applied only when the antecedent of the counterfactual in question is derived to be false. Note that both Indicative Modus Ponens and Indicative Modus Tollens are valid according to Adams' probabilistic semantics for conditionals,[4] which is due to the fact that in Adams' semantics factual statements are treated in the same way as conditionals with tautological antecedents.

In any case, Weak Centering, and with it Counterfactual Modus Ponens and Counterfactual Modus Tollens, are a lot to be sacrificed, and maybe even full Centering is worthy of being given a chance. In fact, perhaps one should aim for even more: if one sides with Stalnaker rather than with Lewis in their famous debate, even

[2] But there are exceptions as well: Gundersen (2004) argues that both C1 and C2 are fallacious. He also sketches the foundations of a corresponding semantics for counterfactuals which is to be based on "statistically normal worlds," that is, he suggests another probabilistic semantics for counterfactuals. However, the details of his semantics are quite different from the one in this paper and they are not worked out in any detail either. For another example of a semantics for counterfactuals in which both C1 and C2 fail to be logically true, see Menzies' (2004) Structural Equations approach which invokes a non-Lewisian Centering condition on "default" worlds that do not necessarily include the actual world. Yablo (2009) reminds us that there are also examples from natural language which put some pressure on Weak Centering: In the example he gives, both Alice and Bill go to the party. 'Well, he would be there, if she was, wouldn't he?' adds something to 'They both went', contra Strong Centering. (Maybe Bill likes Alice.) But suppose we wanted to deny just that added bit, in the belief that Bill despises Alice. 'Nuh-uh, he would not have gone if she did' might do the job, even though this conflicts with Weak Centering.

[3] McGee (1985) questions the logical validity of Modus Ponens, but the failure of Counterfactual Modus Ponens in our Popper function semantics is quite independent of McGee's counterexample, since in our case Modus Ponens is not valid even for simple, uniterated counterfactuals.

[4] However, Adams (1975) considers a possible counterexample to Modus Tollens.

Conditional Excluded Middle: $(A \,\Box\!\!\rightarrow B) \lor (A \,\Box\!\!\rightarrow \neg B)$

might be attractive, too.

We will now deal with ways of "saving" these principles and rules—or at least some of them—within our semantics. This may be achieved (i) by imposing extra semantic constraints on our Popper function models that correspond to these principles (Section 1.1), (ii) by approximating such principles by closely related logical truths without invoking any additional semantic constraints (Section 1.2), (iii) by getting something like these principles at least pragmatically from a version of the Principal Principle (Section 1.3), and (iv) by getting something like these principles on grounds of special ontic reflection properties of conditional chance (Section 1.4).

1.1. Imposing semantic constraints. Just as one can add to the set of logical truths in a possible worlds semantics for modal logic by introducing constraints on accessibility relations, it is possible to add to the set of logical truths according to our probabilistic semantics for conditional logic by imposing extra constraints on Popper functions. Since these semantic constraints are included in the defining clauses of the so-amended Popper function models, they become proper logical constraints. (We will only deal with this in the context of our truth semantics, putting our approximate truth semantics to one side for the time being.)

For instance, what we might call *Actual Determinism* corresponds to *Centering*:[5]

- Semantic constraint:
  For all $w \in W$, for all $A \in \mathcal{L}$: $\mathfrak{P}_w(\llbracket A \rrbracket \mid W)$ equals the truth value of A in w (and hence $\mathfrak{P}_w(\llbracket A \rrbracket \mid W)$ only takes values in $\{0, 1\}$).
- Characteristic axioms:
  $(A \,\Box\!\!\rightarrow B) \supset (A \supset B)$
  $(A \land B) \supset (A \,\Box\!\!\rightarrow B)$.

Note that assuming that $\mathfrak{P}_w(\llbracket A \rrbracket \mid W)$ equals the truth value of A in w does not trivialize Popper functions, for $\mathfrak{P}_w$ may still take any value whatsoever in the interval $[0, 1]$ if taken conditional on a proposition other than W.

Alternatively, if Popper function models are constrained such that state descriptions are assumed to be available in our language—that is, for every world $w \in W$ in a model one assumes that there is a formula $s_w \in \mathcal{L}$, such that $\{w\} = \llbracket s_w \rrbracket$ in that model—the following additional semantic constraints also yield Centering or Weak Centering:

- Semantic constraint:
  For all $w \in W$: $\mathfrak{P}_w(\{w\} \mid W) = 1$,
- Characteristic axioms:
  $(A \,\Box\!\!\rightarrow B) \supset (A \supset B)$
  $(A \land B) \supset (A \,\Box\!\!\rightarrow B)$

and

- Semantic constraint:
  For all $w \in W$: $\mathfrak{P}_w(\{w\} \mid W) > 0$,
- Characteristic axiom:
  $(A \,\Box\!\!\rightarrow B) \supset (A \supset B)$.

[5] The proofs for the claims in this section are simple and we summarize them only very briefly in the appendix again.

Hence, taking for granted that every world is "nationalistic" about itself, in the sense of putting one's total probabilistic weight on oneself, yields full Centering, while presupposing that every world is merely "patriotic" about itself, in the sense of putting a nonnegligible part of one's total probabilistic weight on oneself, corresponds to mere Weak Centering.

Where Actual Determinism merely constrained the absolute probability measure of each world, what might be called *Counterfactual Determinism* constrains the conditional probability measure at each world on every possible antecedent proposition; logically, it corresponds to *Conditional Excluded Middle* or Stalnaker's axiom:

- Semantic constraint:
  For all $w \in W$: $\mathfrak{P}_w(.|.)$ only takes values in $\{0, 1\}$,
- Characteristic axiom:
  $(A \mathrel{\Box\!\!\to} B) \vee (A \mathrel{\Box\!\!\to} \neg B)$.

Here Popper functions $\mathfrak{P}_w$ are indeed trivialized completely in the sense of coinciding with crisp conditional probability measures that have all of their values in the set $\{0, 1\}$. For given A, $\mathfrak{P}_w(. \mid [\![A]\!])$ is simply a standard truth valuation. So it is even possible to restore Stalnaker's preferred logic of conditionals if one is willing to pay the price of Counterfactual Determinism. Or in more positive terms: Stalnaker's logic is the (counterfactually) deterministic special case of the logic for the Popper function semantics. As we shall see at the end of Section 2.4, the correspondence is even stronger: Stalnaker's whole selection function semantics may be seen as the (counterfactually) deterministic special case of the Popper function semantics.

We find that Weak Centering, Centering, and even Conditional Excluded Middle *do* become logical truths if determinism assumptions of varying strength and character are imposed semantically on Popper function models. We offer these constraints to anyone who wants to employ a probabilistic semantics such as ours but who dislikes losing some of their favorite logical axioms.

For our own purposes, regaining Weak Centering by postulating $\mathfrak{P}_w(\{w\} \mid W) > 0$ universally might be excusable. While no world w would be exceptional anymore with respect to the counterfactuals satisfied by w, we will see in Section §2 that there would still be a sense in which the truth of counterfactuals at w would allow for exceptional worlds w' *other than* w, where these worlds w' would be permitted to be arbitrarily close or similar to w (short of being identical to w). This said, even Weak Centering remains to be quite artificial in a semantics such as ours: If $\mathfrak{P}_@([\![B]\!] | [\![A]\!]) = 1$ is the case, shall this really logically exclude the existence of bad luck in the form of an actual $A \wedge \neg B$-counterexample? Only if fortune favors logicians.

1.2. Syntactic approximations of centering. Instead of adding constraints on Popper functions to the defining features of Popper function models, one may also turn to the following syntactic approximations of Lewis' Centering axioms which almost *look* like the Centering axioms themselves and which *are* logically true even in the Popper function semantics of section 2 in part A:

$$(A \mathrel{\Box\!\!\to} B) \supset (\top \mathrel{\Box\!\!\to} (A \supset B))$$

$$(\top \mathrel{\Box\!\!\to} (A \wedge B)) \supset (A \mathrel{\Box\!\!\to} B)$$

What happens here is that the factual components $A \supset B$ and $A \wedge B$ of the original Centering axioms got replaced by their counterfactual counterparts $\top \mathrel{\Box\!\!\to} (A \supset B)$ and

$\top \mathrel{\Box\!\!\rightarrow} (A \wedge B)$. From the viewpoint of Lewis' semantics, these counterfactual counterparts are in fact logically equivalent to the statements they replace, and accordingly the syntactic approximations of the Centering axioms above are logically equivalent to the original Centering axioms, respectively. Hence, while wearing Lewisian semantic glasses, one would be unable to distinguish the Centering axioms from their approximations. However, according to our probabilistic semantics, the original Centering axioms are not logically true while their approximations are. We can therefore offer to the unconvinced Lewisian, who is not willing to impose semantic constraints on Popper function models either, that whenever he or she wants to put forward instances of the Centering axioms to be logically true, they replace these statements by their counterfactual counterparts and otherwise proceed as planned. For all strictly semantic purposes, the exchange should not matter to them, and still they may employ our probabilistic semantics (for which, though, the exchange does correspond to a change of meaning).[6]

But is it plausible that $A \supset B$ might differ in truth value from $\top \mathrel{\Box\!\!\rightarrow} (A \supset B)$ for some everyday statements A and B?[7] Maybe: The truth value of 'Either Peter does not show up at the party next week or Susan shows up at the party' clearly depends just on the truth values of 'Peter does not show up at the party next week' and 'Susan shows up at the party (next week)'. However, '(Even) If $C \vee \neg C$ were to be the case, it would be true that either Peter does not show up at the party next week or Susan shows up at the party' might be regarded as adding a modal component to the material conditional which might correspond to a relevant change of meaning—perhaps one wants to say: *independently of how things are going to develop,* either Peter does not show up at the party next week or Susan shows up at the party.

But even if the mutually corresponding statement forms within the pair of Centering axioms and the pair of their counterfactual counterparts never actually differed in truth value, this would still not show that they would be *logically* equivalent. And it is not clear that a question like this can be resolved at all by mere inspection of a couple of examples from natural language, without an accompanying theoretical analysis that goes beyond the examples in question.

1.3. Saving Counterfactual Modus Ponens/Tollens pragmatically. Now we will concentrate solely on Counterfactual Modus Ponens and Counterfactual Modus Tollens. As it turns out, on the basis of a semantics such as ours one can save them at least pragmatically by invoking a conditional variant of Lewis' famous Principal Principle (cf. Lewis, 1980) which relates a rational agent's initial conditional credence function $\mathfrak{Cr}$ (say, another Popper function) with the conditional chance function $\mathfrak{P}$ at some arbitrary time t:[8]

---

6 As mentioned before, also in Adams' probabilistic semantics a factual statement A is treated semantically in the same way as its corresponding indicative conditional of the form $\top \rightarrow A$, since their degrees of acceptability are the same with respect to every epistemic probability measure.

7 We thank Leon Horsten for a very helpful exchange on this.

8 We omit '$[\![.]\!]$' in this section and for most of the next one, for the sake of simplicity. We will be equally lenient with respect to mixing metalinguistic and object-linguistic expressions of probabilistic claims as the literature on Lewis' Principal Principle is for the most part. Finally, number terms such as 'r' which occur in the probabilistic reflection principles that we are going to consider are taken to be numerals, that is, *primitive* individual constants for numbers, in order to avoid problems such as what is known in the literature as Miller's Paradox.

- By the conditional variant of the Principal Principle,

$$\mathfrak{Cr}(B \mid A \wedge \mathfrak{P}(B|A) = r \wedge C) = r$$

for any C that is admissible at t.
- Hence,

$$\mathfrak{Cr}(B \mid A \wedge \mathfrak{P}(B|A) = 1) = 1$$

by $r = 1$, and since any tautological C is certainly admissible.
- In other words, by the truth condition for counterfactuals of our probabilistic semantics:

$$\mathfrak{Cr}(B \mid A \wedge (A \mathrel{\Box\!\!\rightarrow} B)) = 1$$

(where the reference point of $A \mathrel{\Box\!\!\rightarrow} B$ is t).
- But this means that if '$\rightarrow$' is the indicative 'if-then', then by Adams' pragmatic theory for indicative conditionals,[9]

$$A \wedge (A \mathrel{\Box\!\!\rightarrow} B) \rightarrow B$$

gets assigned a degree of acceptability of 1 based on any credence function $\mathfrak{Cr}'$ that is sufficiently like the agent's initial $\mathfrak{Cr}$, in particular, any $\mathfrak{Cr}'$ that results from the initial $\mathfrak{Cr}$ by updating on typical instances of evidence.
- So if A and $A \mathrel{\Box\!\!\rightarrow} B$ are assertable according to such a $\mathfrak{Cr}'$, then B is so as well, since Adams' semantics supports Indicative Modus Ponens.
Accordingly, if $\neg B$ and $A \mathrel{\Box\!\!\rightarrow} B$ are assertable according to $\mathfrak{Cr}'$, then $\neg A$ is so as well, since Adams' semantics also supports Indicative Modus Tollens.

We conclude that there is an explanation for the intuitive attractiveness of Counterfactual Modus Ponens and Counterfactual Modus Tollens even for the likes of our probabilistic semantics for counterfactuals: though not being logically valid, these rules are still permissible *pragmatically* in the sense of preserving a high degree of acceptability in situations in which their premises are acceptable and *nothing nonadmissible is acceptable, that is, in particular, nothing which undermines the conclusion formulas B and $\neg A$ of the two argument forms in question*. This restriction on the pragmatic permissibility of these arguments is closely related to Principles of Total Evidence that one may find in Reichenbach's, Carnap's, and Hempel's work on statistical explanation and inductive reasoning; in particular, from the statistical law $\mathfrak{P}(B[x] \mid A[x]) = r$ and the factual evidence $A[a]$ one may infer $B[a]$ with inductive probability r only if $A[a]$ is the *totality of one's relevant evidence concerning the individual a*. This is simply the statistical counterpart of the defeasible pragmatic permissibility of counterfactual Modus Ponens as derived above.[10]

1.4. Saving Counterfactual Modus Ponens/Tollens by ontic reflection. Here is one final take on Counterfactual Modus Ponens and Counterfactual Modus Tollens, but this time not a pragmatic but an ontic one: Instead of having their object-linguistic material conditional versions

[9] Here, as in the following sections, we always assume conjunction to bind more strongly than any conditional operator.

[10] Schurz (2001) develops a similar thought in his theory of so-called "normic laws" $A \Rightarrow B$ which are taken to be the phenomenological laws of evolutionary systems; in this theory, only a *default* version of Modus Ponens is valid: if *all an agent knows* are A and $A \Rightarrow B$, then the agent may reliably conclude B. Schurz explicitly notes the historical ancestors of this view.

$$(A \wedge (A \,\Box\!\!\rightarrow B)) \supset B$$

$$(\neg B \wedge (A \,\Box\!\!\rightarrow B)) \supset \neg A$$

coming out as logically true, as it is the case in Lewis' semantics, we might at least want their counterfactual counterparts

$$(A \wedge (A \,\Box\!\!\rightarrow B)) \,\Box\!\!\rightarrow B$$

$$(\neg B \wedge (A \,\Box\!\!\rightarrow B)) \,\Box\!\!\rightarrow \neg A$$

to be logically true. But so far our semantics does not subserve any interesting instances of logically true nestings of $\Box\!\!\rightarrow$. Unpacking the metalinguistic clauses that would have to hold in order for these object-linguistic statements to be true, we are led to:

$$\mathfrak{P}(B \mid A \wedge \mathfrak{P}(B|A) = 1) = 1$$

$$\mathfrak{P}(\neg A \mid \neg B \wedge \mathfrak{P}(B|A) = 1) = 1.$$

So we *would* get the nested counterfactual statements for before as logically true if we restricted ourselves to world-relative conditional chance measures $\mathfrak{P}$ for which ontic probabilistic reflection principles such as these would hold universally by a corresponding assumption on our models.

We do not know how plausible or problematic such second-order chance principles are. But it is certainly interesting to note that "counterfactualized" logical law versions of the rules of Counterfactual Modus Ponens and Counterfactual Modus Tollens are compatible with our probabilistic semantics if only conditional chance functions are subject to the right kinds of probabilistic reflection.[11]

§2. An equivalent semantics: the probabilistic limit semantics. In order to improve our understanding of the Popper function semantics and how it ought to be interpreted, it is illuminating to compare it with an alternative probabilistic semantics for counterfactuals which looks quite different but which will turn out to be equivalent in a precisely defined sense. We call it the *Limit Semantics*.

DEFINITION 2.1 $\langle W, \mathfrak{A}, ((P_i^w)_{i \in I_w})_{w \in W}, (\leq^w)_{w \in W}, [\![.]\!]\rangle$ *is a probabilistic limit model iff*

- W *is a nonempty set of possible worlds.*
- $\mathfrak{A} = \{[\![A]\!] \mid A \in \mathcal{L}\}$*, for* $\mathcal{L}$ *as explained in part A.*
- *For every* $w \in W$, $(P_i^w)_{i \in I_w}$ *is a family of absolute (finitely additive) probability measures on* $\mathfrak{A}$*;* I_w *is the index set of that family.*[12]

[11] A probabilistic reflection principle such as $\mathfrak{P}(B \mid A \wedge \mathfrak{P}(B|A) = 1) = 1$ may look harmless—and indeed one can prove it consistent in the same way as one proves the Conditional Principal Principle consistent—but it is actually more restrictive than one might think. For example, it can be shown to have the following consequence on absolute chance:

- For all $w \in W$: If $[\![A]\!] = \{w\} \in \mathfrak{A}$, $\mathfrak{P}(A) > 0$, and $r_1, r_2 > 0$ then

$$\text{if } \mathfrak{P}(\mathfrak{P}(A) = r_1) > 0 \text{ and } \mathfrak{P}(\mathfrak{P}(A) = r_2) > 0 \text{ then } r_1 = r_2.$$

[12] Here a family is just a mapping from an index set to a set of absolute probability measures. One may just as well speak of a *sequence* of absolute probability measures instead, as long as one bears in mind that an index set need not be ordered in the way in which the natural number sequence is ordered.

 https://doi.org/10.1017/S1755020311000165

- *For every* $w \in W$, $\leq^w$ *is a linear preorder on* I_w.[13]
- *The Convergence Assumption is satisfied:*
  For all $w \in W$ *and* $[\![A]\!], [\![B]\!] \in \mathfrak{A}$, *either there is no* $i \in I_w$ *such that* $P_i^w([\![A]\!]) > 0$, *or the family* $\left(\frac{P_i^w([\![B \wedge A]\!])}{P_i^w([\![A]\!])}\right)_{i \in I_w}$ *converges.*
  Remark: $\left(\frac{P_i^w([\![B \wedge A]\!])}{P_i^w([\![A]\!])}\right)_{i \in I_w}$ *is said to converge to* $x \in [0,1]$ *iff for all* $\epsilon > 0$ *there is an index* $j \in I_w$ *with* $P_j^w([\![A]\!]) > 0$, *such that for all* $i \leq^w j$ *with* $P_i^w([\![A]\!]) > 0$ *it holds that* $\left|\frac{P_i^w([\![B \wedge A]\!])}{P_i^w([\![A]\!])} - x\right| < \epsilon$. *Or in terms of an illustration:*

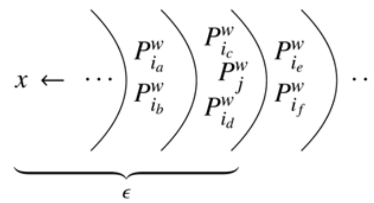

- $[\![.]\!] : \mathcal{L} \to \wp(W)$ *satisfies the following semantic rules:*
  - *Standard semantic rules for standard propositional connectives.*
  - $w \in [\![A \,\square\!\!\rightarrow B]\!]$ *iff either of the following is satisfied:*
    - *There is no* $i \in I_w$, *such that* $P_i^w([\![A]\!]) > 0$.
    - *It holds that:*

$$\lim_{i \in I_w} \left(\frac{P_i^w([\![B \wedge A]\!])}{P_i^w([\![A]\!])}\right) = 1.$$

(There is a corresponding Limit semantics for approximate truth which is analogous to the approximate truth semantics of chapter 3 of part A; since it should be clear by now how to state it, we will not do so ourselves.)

Where in the Popper function semantics each world had its own primitive conditional probability measure, in this semantics each world has its own family of absolute probability measures, that is, of probability measures in the standard sense of the word. These absolute probability measures, or rather their indices, are assumed to be ordered by a linear preorder that belongs to the world in question. Linear preorders are like proper linear orders except that they allow for ties, so that there may be two probability measures P_i^w and P_j^w with $i \leq^w j$ and $j \leq^w i$, but where this does not entail that $i = j$. As Lewis (1973b) has shown, linear preorders correspond one-to-one to sphere systems, as long as no Centering assumptions of any kind are imposed on sphere systems. In other words: one may view the semantics above as resulting from a Lewisian sphere semantics for worlds without Centering assumptions by replacing the neighboring worlds of any world w by absolute probability measures while keeping the order structure of the sphere system intact.

The central idea behind the Limit semantics is now the following: $A \,\square\!\!\rightarrow B$ is true in w if and only if (the trivial case) either the absolutely probability of A is never positive

[13] That is: $\leq^w$ is reflexive, transitive, and: (Linearity) For all $u, v \in W$, $u \leq^w v$ or $v \leq^w u$.

for any absolute probability measure that is revolving around w, or (the nontrivial case) the conditional probability of B given A tends towards 1 as being given by those absolute probability measures revolving around w which assign positive probability to A and for which therefore the Ratio formula of conditional probability is applicable.[14]

The trivial case is nothing but the probabilistic version of Lewis' trivial case of truth for $A \;\Box\!\!\rightarrow\; B$ in w if A is impossible as seen by w.[15] In order for the nontrivial case to go through, we assume first of all that the relevant conditional probabilities converge if taken along the given linear preorder at w: this is the task of the Convergence Assumption above. Secondly, it has to be made clear what 'convergence' ought to mean here; after all, although each $\frac{P_i^w(\llbracket B \wedge A \rrbracket)}{P_i^w(\llbracket A \rrbracket)}$ is a real number, the underlying order $\leq^w$ of indices i is an arbitrary linear preorder which may differ from the order of the natural numbers substantially; so it is not just 'convergence of a real number sequence' that is asked for. Fortunately, the theory of Moore–Smith sequences (or nets) in Functional Analysis offers precisely what we need: the corresponding definition of convergence that we have stated above leaves convergence with the same intuitive properties as convergence for real number sequences while applying to linear preorders of conditional probabilities as planned. The only adaptation necessary in our context is to restrict attention only to those members of $\left(\frac{P_i^w(\llbracket B \wedge A \rrbracket)}{P_i^w(\llbracket A \rrbracket)}\right)_{i \in I_w}$ for which $P_i^w(\llbracket A \rrbracket)$ is greater than 0.[16]

We find that the Limit semantics is nothing but the probabilistic version of the Lewisian standard semantics without Centering assumptions: each world determines an ordering of absolute probability measures rather than of worlds; where Lewis defines $A \;\Box\!\!\rightarrow\; B$ to be true in w if the A-worlds around w eventually "converge on" B's truth (with a Stalnaker-type Limit Assumption not being presupposed, as intended by Lewis), we now have an actual real number limit in its place; that is it.

Hence, for the moment, think of this semantics as Lewis' semantics probabilified. Especially, prima facie, supply the orderings for absolute probability measures with a Lewisian interpretation in terms of closeness or weighted similarity—though more will have to be said about this below.

2.1. From the Popper function semantics to the limit semantics by representation. The Limit semantics relies on semantic structures which appear to be quite different from the conditional probability measures that have been assigned to worlds by the original Popper function semantics. However, in a sense explained by the following general

---

[14] A variant of this truth condition in the context of nonmonotonic consequence relations may be found in the "Sequence Semantics" of Leitgeb (2004, p. 171).

[15] We should note that—like Lewis' semantics—both our Popper function semantics and the Limit semantics trivialize so-called counterpossible conditionals with impossible antecedents, if the underlying set W of possible worlds is taken to be the set of metaphysically possible worlds or a subset thereof. *Any* such conditional is then evaluated true for trivial reasons, which is certainly questionable. If this is taken to be a problem, then one way out would be to choose W in a less restrictive manner. Accordingly, given a Popper function $\mathfrak{P}_w$, there might then be propositions X which are impossible and yet X would be "normal" with respect to $\mathfrak{P}_w$, that is, $\mathfrak{P}(W \backslash X \mid X) \neq 1$. See Brogard & Salerno (2007) for a corresponding quasi-Lewisian semantics with impossible worlds in which counterpossibles are not necessarily trivially true.

[16] Take all other members of the family $\left(\frac{P_i^w(\llbracket B \wedge A \rrbracket)}{P_i^w(\llbracket A \rrbracket)}\right)_{i \in I_w}$ to be undefined; so such families of real numbers are really partial functions.

representation theorem, the semantics in fact turns out to be *equivalent* to the Popper function semantics:[17]

THEOREM 2.2

- *Every family $(P_i)_{i \in I}$ of finitely additive probability measures (on one and the same countable algebra $\mathfrak{A}$) which satisfies the Convergence Assumption with respect to a given linear preorder $\leq$, represents a Popper function $\mathfrak{P}$ (on $\mathfrak{A}$), where the representation is given by:*

  Repr *If there is an $i \in I$, such that $P_i(X) > 0$, then*

  $$\mathfrak{P}(Y|X) = \lim_{i \in I} \left(\frac{P_i(Y \cap X)}{P_i(X)} \right)$$

  Otherwise, $\mathfrak{P}(Y|X) = 1$.

- *Every Popper function (on a countable algebra $\mathfrak{A}$) can be represented by means of **Repr** by a family of finitely additive probability measures (on $\mathfrak{A}$) which satisfies the Convergence Assumption with respect to some linear preorder $\leq$.*

By the representation theorem, all Popper functions—and only Popper functions—may be represented as families of absolute probability measures, such that (in the nontrivial case) the values of Popper functions are represented by limits of conditional probabilities as being determined from the absolute probability measures by means of the Ratio formula.

[17] The proof is stated in the appendix again. Our representation theorem is an improvement of previous representation results for Popper functions. Rényi's (1955) and Csaszar's (1955) representation theorems do not represent Popper functions in terms of probability measures but in terms of measures which may take any value up to ∞. van Fraassen's (1976) elegant representation theorems rely on other "tricks": either he adds axiomatic constraints to the standard axioms for Popper functions and then represents only the so-constrained conditional probability functions, or the representing absolute probability functions that he constructs are not all defined on one and the same algebra. Spohn's (1986) representation theorem presupposes σ-additivity as a further axiom, which in our context would be unfortunate since it is unclear whether chances obey σ-additivity: if chances have formal properties which are much *like* those of limiting frequencies, then σ-additivity will fail for them; cf. Schurz & Leitgeb (2008) (see Arló-Costa & Parikh, 2005, for further worries about postulating σ-additivity for Popper functions). In order for our representation theorem to go through, no assumptions over and above the original Popper function axioms are necessary at all, nor are there any constraints on what the representing families of absolute probability measures have to be like except for the Convergence Assumption. This is possible because we represent Popper functions in terms of *limits* of ratios, unlike any of the results just cited. This said, it should be noted that our proof of Theorem 2.2 requires that $\mathfrak{A}$ is countable, which is definitely a restriction. However, a corresponding countability assumption is also contained in Popper's original axiomatization of conditional probability functions, and in a context such as ours, in which the relevant algebra $\mathfrak{A}$ is generated from (as is standard) a countably infinite language $\mathcal{L}$, the assumption is unproblematic. Theorem 2.2 also explains nicely under which circumstances Popper function values reduce to standard ratios: if $\lim_i P_i(X) > 0$, then $\lim_i \left(\frac{P_i(Y \cap X)}{P_i(X)}\right) = \frac{\lim_i P_i(Y \cap X)}{\lim_i P_i(X)}$ by elementary calculus, and hence conditional probabilities are ratios of absolute probabilities in this case. In the other case, '$\frac{\lim_i P_i(Y \cap X)}{\lim_i P_i(X)}$' is undefined while '$\lim_i \left(\frac{P_i(Y \cap X)}{P_i(X)}\right)$' is not, by the Convergence Assumption. Note that from the proof of Theorem 2.2 it follows that one could actually strengthen the second claim ("completeness of representation") of the theorem by replacing 'linear preorder' by 'linear (partial) order'. The same replacement could be applied to the first claim ("soundness of representation") of the theorem, too, but that would weaken the claim accordingly.

But this representation is precisely what is needed in order to transform the Popper function semantics from section 3 of part A into the new Limit semantics and vice versa: simply represent the Popper function of world w by a corresponding family of absolute probability measures, and adapt the truth condition for counterfactuals by exchanging Popper function values by their corresponding limits of ratios of absolute probabilities. More precisely, taking into account the time relativity discussed in section 4.3 in part A, for each reference point t_r, represent the Popper function for t_r at w by a corresponding family of absolute probability measures which are all assigned the same reference point t_r. In this way, the truth condition of the Popper function semantics is translated into the truth condition of the new Limit semantics (and the same would apply to the two corresponding semantics for approximate truth). A fortiori, by combining theorem 4 of part A and Theorem 2.2 above, the system V of conditional logic is sound and complete with respect to the Limit semantics, too.[18]

It is tempting to give this representation result a reductive interpretation, as in: the theorem demonstrates that Popper functions are really *nothing but* limits of ratios of absolute probability measures. But this would be an overinterpretation as long as one did not have good additional independent reasons for thinking so. What the theorem actually does, if taken by itself, is to establish a correspondence—a bridge principle—between primitive conditional probability measures and linearly preordered families of absolute probability measures (which satisfy the Convergence Assumption), without entailing that either of those would be "nothing but" the other. If anything, one ought to give the result an epistemic reading: It is possible to *understand* Popper functions in terms of ordered families of absolute probability measures, just as it is possible to *understand* ordered families of absolute probability measures in terms of Popper functions. This said, we *do* think there are in fact good independent reasons to believe that Popper functions are at least epistemically prior to ordered families of absolute probability measures; we will return to this point at the end of Section §2. So we *do* suggest that the intended ordered family of absolute probability measures at a world is "determined" from the intended Popper function at that world by means of Theorem 2.2, since our best way of getting epistemic access to families like that is by first getting epistemic access to Popper functions. At the same time we believe that some properties of Popper functions may be more easily grasped in terms of common sense concepts if facilitated by the conceptual apparatus of families of absolute probability functions and their linear preorders, rather than on the basis of the primitive concept of conditional probability alone.

We have glossed over one additional problem here: Theorem 2.2 does not determine for a given Popper function a *unique* family of absolute probability measures that represents it; it only establishes the existence of such a family. Indeed, it is easy to see that uniqueness must fail, so linearly ordered families of absolute probability measures have some surplus structure if compared to Popper functions. However, all families of absolute probability measures which represent one and the same Popper function in the sense of **Repr** must at least be indistinguishable with respect to conditional probabilities (and hence absolute

[18] It would be possible to modify the Limit semantics in the way that weaker systems of conditional logic would be sound and complete with respect to the so modified semantics. In order to do so, linearity would have to be dropped, the Convergence Assumption would have to be adapted, and a choice function would have to be introduced by which for every proposition $[\![A]\!]$ and every world w some maximal descending path through $\leq^w$ would be chosen: $A \,\square\!\!\rightarrow B$ would then have to be evaluated true in w if the conditional probability of $[\![B]\!]$ given $[\![A]\!]$ tended towards 1 along that path.

probabilities) *in the limit*. So while the exact ordering of a representing family of absolute probability measures is not determined uniquely by Theorem 2.2, the representation conditions at least impose serious constraints on that ordering.

2.2. Keeping our promises. Why is Theorem 2.2 interesting philosophically? We hope to have given good reasons (mainly in part A) to believe that promise 1 on p. 85 is satisfied by the Popper function semantics, that is:

1. The Popper function semantics conveys a natural and simple probabilistic semantics for subjunctive conditionals.

What Theorem 2.2 adds to this is that although the Popper function semantics relies on a primitive notion of conditional probability, it may still be understood as being based on a notion of comparative similarity—formally: linear preorders or sphere systems—as employed in the Limit semantics. So it should be possible to interpret the Popper function semantics in a quasi-Lewisian manner based on the representing ordered families of absolute probability measures, with the only modification that now absolute probability measures take over the role of worlds, and especially the actual absolute probability measure takes over the role of the actual world as being the center of the "sphere system" of that world:

2.a $\mathfrak{P}_w(.|W)$ may thus be interpreted as the *actual* absolute probability measure from the viewpoint of w.[19]

2.b $\mathfrak{P}_w(Y|X) = 1$ iff the *more similar* an absolute probability measure is to the actual absolute probability measure, the *closer* the conditional probability of Y given X that it determines is to 1.

So promise 2 on p. 85 is satisfied, too: the Popper function semantics is actually very close to Lewis' semantics, despite initial appearance to the contrary. Obviously, there are still lots of differences, not just because of the probabilistic nature of our two semantical systems but also because it is not clear at all that the ordering of absolute probability measures at worlds that is determined from intended conditional probability measures by Theorem 2.2 conforms to anything like Lewis' (1979) well-known heuristics for comparative similarity of worlds. But at least some of the formal properties of the two kinds of similarity or closeness are alike; and as we will argue below, it may be a virtue of our theory if it actually deviates in some respects from Lewis' intuitive suggestions for how to understand similarity or closeness. At the same time, not *only* formal properties might carry over from Lewis' intended closeness relations for worlds to the closeness of absolute probability measures as given by intended primitive conditional probability measures through Theorem 2.2. For instance, holding the temporal reference point fixed, absolute probability measures "close by" the actual absolute probability measure for that reference point might be assumed to

[19] Indeed, by the Convergence Assumption, one can show for every Limit model and every world $w \in W$: if there is no absolute probability measure P_i^w the index i of which is the unique "globally" least index in I_w according to $\leq^w$, then the model can be extended by such a measure, which is then given by $\lim_{i \in I} P_i(.)$, without affecting any limiting ratios. When we were speaking before, informally, of absolute probability measures P_i^w as revolving around the world w, it would have been more appropriate maybe to describe them as revolving around the uniquely determined actual absolute probability measure of w. Note that if one of the semantic constraints were being imposed which were shown to amount to Centering in Section 1.1, then the actual absolute probability measure of w would in fact *coincide* with the truth value assignment that is determined by w.

assign values which are close to crisp chances (0 or 1) to more and later events of the form *at time ... it happens that ...* than absolute probability measures which are "remote from" the actual absolute probability measure in the relevant ordering; but this remains to be seen, and if it is so at all, it ought to be argued for on the basis of properties of primitive conditional chance. On the other hand, in some respects, relative closeness of absolute probability measures as determined from an intended conditional probability measure by the lights of Theorem 2.2 may be expected to differ from relative closeness according to Lewis and maybe also from standard temporal intuitions on closeness. Here is an example: If one reconsiders our probabilistic take on the "hat story" in section 4.7 of part A from the viewpoint of the Limit semantics (in combination with Theorem 2.2), then since

$$\mathfrak{P}_{@}\left([\![\text{My coat is in Fence's shop (at some } t' > t + \Delta t)]\!] \mid [\![A^1_t \vee A^2_{t+\Delta t}]\!]\right)$$

was far from 1 there, it cannot be the case that absolute probability functions which assign positive probability to $A^2_{t+\Delta t}$ approximate the actual absolute probability measure *more closely* than the absolute probability functions which assign positive probability to A^1_t, for then it would be the case that

$$\mathfrak{P}_{@}\left([\![\text{My coat is in Fence's shop (at some } t' > t + \Delta t)]\!] \mid [\![A^1_t \vee A^2_{t+\Delta t}]\!]\right) =$$
$$\mathfrak{P}_{@}\left([\![\text{My coat is in Fence's shop (at some } t' > t + \Delta t)]\!] \mid [\![A^2_{t+\Delta t}]\!]\right) \approx 1$$

by the story again. And this is so *although* A^1_t necessitates an earlier deviation from the actual course of history than $A^2_{t+\Delta t}$, which one might have thought might show up in terms of closeness. The general moral to be drawn from this is: Let oneself be guided by the conditional chances in question—which in this case are determined by the hat story—rather than by some preestablished intuitions on similarity or closeness. Use the former to determine the latter by the representation theorem, rather than the other way round.

The correspondence between the Popper function semantics and the Limit semantics that is established by Theorem 2.2 also enables us to see, finally, in what sense the Popper function semantics is more tolerant with respect to exceptions than Lewis' semantics, over and above the failure of Centering (assuming that Lewis' standard of closeness and the standard of closeness that is determined by Theorem 2.2 are sufficiently comparable):

3. The truth of $A \,\Box\!\!\rightarrow B$ in w, according to the Popper function semantics, allows for $P^w_i([\![\neg B]\!] \mid [\![A]\!]) > 0$ for all i with $P^w_i([\![A]\!]) > 0$, if looked at from the viewpoint of the Limit semantics.

This means that if $A \,\Box\!\!\rightarrow B$ is true in a world w, then independently of how close one gets to the actual probability measure of w, it may be that the conditional chance for an A-world to be an exceptional $\neg B$-world is never actually 0. So it is not just that exceptional $A \wedge \neg B$-worlds need not be wiped out completely, it is not even the case that the chance of $\neg B$ given A is bound to be 0 from some index $i \in I_w$. We hope this makes it clear enough that also promise 3 on p. 1 has been kept.

We have thus applied the conceptual machinery of ordered families of absolute probability measures in order to analyze properties of Popper functions in different terms; but finding a way to describe one and the same fact by two different sets of concepts can be progress, too. In Section 2.3 we will now finally defend the availability of the interpretation of 'a Popper function value being very high' understood as a vague term, which we had introduced in section 3 of part A, and we will point to an interesting logical correspondence between, on the one hand, validity in the sense of either the Popper function semantics or the Limit semantics, and on the other, validity in the sense of Adams' semantics for conditionals. Sections 2.4 and 2.5 will strengthen the correspondence between the Popper function semantics/the Limit semantics and Lewis' semantics, while Section 2.6 will make it more precise how and why these probabilistic semantics differ from Lewis'.

2.3. Validity in our semantics, validity in Adams' semantics, and high probability. We are back again to Adams' semantics for conditionals. Here is another sense in which our semantics builds on, and extends, Adams' semantics, as claimed at the beginning of this paper. Even in a probabilistic—and hence quantitative—theory of conditionals such as Adams', ultimately a step towards a discrete—qualitative—treatment of conditionals must be taken. For even if an agent has assigned some quantitative degree of acceptability to an indicative conditional $A \rightarrow B$, the agent must eventually decide when, and if, he or she is permitted to assert that conditional in conversation. Of course, ideally, its exact degree of acceptability might still show up in the firmness with which $A \rightarrow B$ gets asserted (or, less ideally, with what loudness), but there still has to be a point at which that subject would be willing to assert $A \rightarrow B$; the continuum of degrees of acceptability does not correspond to some "existential speech act continuum" between, as it were, the nonexistence of an assertion and its more or less firmly instantiated existence. Otherwise we would be asserting all statements whatsoever all the time, only with differing degrees of firmness! One way to handle this transition from degrees of acceptability to assertability *simpliciter* is by considering "high conditional probabilities," where 'high' is regarded as semantically vague, at least as a necessary condition on assertability in normal circumstances.[20] Normally, only if the conditional probability that belongs to an indicative conditional is high, that is, close to 1, an agent is licensed to assert that conditional. While Adams does not have so much to say about high probability as licensing an act of plain assertion, he gives a detailed and mathematically precise account of high probability in his semantic (or maybe pragmatic) theory of probabilistically valid arguments—the "logic of high probability," as he calls it in Adams (1986), where he investigates them in terms of so-called metalinguistic high probability predicates: Inferences with conditionals are logically valid in this sense if they necessarily *preserve* high conditional probabilities. The corresponding entailment relation may be introduced and explained in different ways (including an uncertainty sum condition and an epsilon–delta continuity criterion), but in our context the following one is the most salient.[21]

Consider arguments of this simple conditional form:[22]

$$\begin{array}{c} A_1 \rightarrow B_1 \\ \vdots \\ A_n \rightarrow B_n \\ \hline C \rightarrow D \end{array}$$

By Adams' theory, such arguments are probabilistically valid if and only if whenever all the conditional probabilities that correspond to the premises are very high, the same holds

[20] In Section 2 of part A we already dealt with assertability *simpliciter* briefly when we showed that the two pragmatic meanings of counterfactuals that we had uncovered resulted in very similar constraints on the recipient of a counterfactual assertion.

[21] For yet another way of introducing probabilistic validity Adams style by quantification over Popper functions with an intended epistemic interpretation, see McGee (1994).

[22] In Adams' system, conditionals cannot be nested.

for the conditional probability that corresponds to the conclusion; and this gets analyzed formally in the way that for all infinite sequences $P_1, P_2, P_3, \ldots$ of absolute probability measures on the corresponding language it holds:[23]

If

$$P_i(B_1|A_1) \quad \text{tends towards 1 for } i \to \infty,$$

$$\vdots$$

$$P_i(B_n|A_n) \quad \text{tends towards 1 for } i \to \infty,$$

then

$$P_i(D|C) \quad \text{tends towards 1 for } i \to \infty$$

where if $P_i(\varphi) = 0$ then $P_i(\psi|\varphi)$ is regarded to be equal to 1.

So a premise's or conclusion's conditional probability being "high" really means that the conditional probability tends towards 1 along an infinite sequence of probability measures, without necessarily being equal to 1 from some index i. But this should sound familiar by now from our Limit semantics in which the truth of a counterfactual at w consists in its conditional probability tending towards 1 if following the order of indexed absolute probability measures for w, without necessarily being equal to 1 from some index i.

This invites the following two conclusions: First of all, although our Limit semantics is truth-conditional, and hence the validity of an argument is given by strict truth preservation, validity in our sense coincides with probabilistic validity in Adams' sense in the case of purely conditional arguments, since the assumed truth of conditional premises and the entailed truth of conditional conclusions amount to their conditional probabilities going to 1 in the limit, just as in Adams' semantics. The only differences are that where in the Limit semantics the ordering of absolute probability measures is given relative to some world, this is not so in Adams' semantics; and where we interpret these probability measures ontically he interprets them epistemically. Indeed, our semantics *properly* extends Adams' by allowing for arbitrary nestings of conditionals, whether in the premises or the conclusions. Consequently, it should not be surprising anymore that Adams' logical system for indicative conditionals is just the flat fragment of the logical system V for subjunctive conditionals that is sound and complete with respect to the Limit semantics. The same logical correspondence holds between our original Popper function semantics and Adams', in light of our representation result Theorem 2.2.[24]

Secondly, if Adams is justified in analyzing the meaning of the vague term 'high conditional probability' in terms of limiting conditional probabilities, then we should be justified in doing so, too. So if in the Popper function semantics we say that

$$\mathfrak{P}_w([\![B]\!] \mid [\![A]\!]) = 1$$

---

23 For simplicity, we avoid '$[\![.]\!]$' again.

24 There remains to be one logical difference, however: Since Adams treats factual statements A in the same way as conditionals $\top \to A$—which is not so in either of our semantics—factual statements interact with indicative conditionals according to Adams' theory differently from the way in which statements without $\Box\!\!\rightarrow$ interact with subjunctive conditionals in our semantics.

within our semantic rule for counterfactuals, this may be understood as

$$\lim_{i \in I_w} \left(\frac{P_i^w([\![B \wedge A]\!])}{P_i^w([\![A]\!])} \right) = 1$$

in light of Theorem 2.2, which—following Adams (1986)—may be taken to say

the conditional probability of B given A is very high.

with 'very high' being vague.

Or, perhaps, if the ordering of I_w is expressed in terms of similarity, it means:

if a measure P is similar to the actual probability measure, then

$$P([\![B]\!] \mid [\![A]\!]) > 1 - \epsilon \text{ for small } \epsilon$$

with both 'similar' and 'small' being vague. No real number $\epsilon > 0$ would be small enough to yield a threshold value by which one could interpret 'very high' as '$> 1-\epsilon$' without also taking into account how similar the probability measure is to the actual one; the threshold remains only vaguely determined.

From the viewpoint of the Popper function semantics: A conditional probability equal to 1 definitely amounts to truth; a conditional probability $< 1-\epsilon$ for any real number ϵ greater than zero definitely amounts to falsity (though approximate truth is not ruled out yet); but the theory purports that there is a blurred area of transition in between which may be taken to consist in a blurred interval of infinitesimal quantities.[25] This blurred area shows up in the Limit semantics in the way that if $\lim_{i \in I_w} \left(\frac{P_i^w\left([\![B \wedge A]\!]\right)}{P_i^w\left([\![A]\!]\right)} \right) = 1$, then this really covers two possible cases: (i) Either $\frac{P_i^w\left([\![B \wedge A]\!]\right)}{P_i^w\left([\![A]\!]\right)} \Big) = 1$ from some index i. (ii) Or there is no index i from that $\frac{P_i^w\left([\![B \wedge A]\!]\right)}{P_i^w\left([\![A]\!]\right)} \Big)$ would be 1, so that the fractions merely become 1 in the limit. By representation, both cases would amount to the very same condition '$\mathfrak{P}_w([\![B]\!] | [\![A]\!]) = 1$'

[25] In fact, one can show that there is yet another equivalent probabilistic semantics for subjunctive conditionals in which each possible world comes equipped with one and only probability measure—as in the Popper function semantics—but where any such probability measure is an absolute one—as in the Limit semantics. However, such probability measures then have to be allowed to assign *nonstandard* reals to propositions. Given that, one can formulate the corresponding truth condition for counterfactuals as follows: $w \in [\![A \mathrel{\Box\!\!\rightarrow} B]\!]$ iff either of the following is satisfied:

- $P^w([\![A]\!]) = 0$.
- It holds that:
$$\frac{P^w\left([\![B \wedge A]\!]\right)}{P^w\left([\![A]\!]\right)} > 1 - \epsilon$$
for some *infinitesimal* ϵ.

From a representation theorem for Popper functions in terms of nonstandard probability measures (cf. McGee, 1994; Halpern, 2001) it follows that this semantics is equivalent to the Popper function semantics in a similar sense as the limit semantics is equivalent to it (see also Lehmann & Magidor, 1992, on a nonstandard probability semantics for nonmonotonic consequence relations.) In such a semantics, for each true conditional there are always sufficiently small *exact* thresholds above which its corresponding conditional probability amounts to truth—but those are nonstandard thresholds. 'Close to 1' then means: being equal to 1 or "almost" equal to 1, that is, *being 1 up to some infinitesimal number*.

on the Popper functions $\mathfrak{P}_w$ that is represented by $(P_i^w)_{i \in I_w}$. But for the same reason, from the viewpoint of the Limit semantics, one may read '$\mathfrak{P}_w([\![B]\!]|[\![A]\!]) = 1$' in either of two ways which correspond to (i) and (ii) from before: (i) $\mathfrak{P}_w([\![B]\!]|[\![A]\!])$ is precisely 1. (ii) $\mathfrak{P}_w([\![B]\!]|[\![A]\!])$ is merely close to 1 (without being precisely 1). And if both cases are to be subsumed under one heading, we may say: $\mathfrak{P}_w([\![B]\!]|[\![A]\!])$ is close to 1. This finally justifies our interpretation of '$\mathfrak{P}([\![B]\!]|[\![A]\!]) = 1$' in terms of 'the probability of B given A is very high, or close to 1', where 'very high' and 'close' are vague terms, which we suggested in section 3 of part A. This is yet another sense in which we may understand our semantics for conditionals to build on, and extend, Adams' semantics for conditionals.

Ideally, we would be able to support Adams' account of the vague term 'high' (or 'close to 1') as applying to probabilities by some form of Limit semantics for vague expressions in general. There are some approaches to vagueness which might go some way towards achieving this: First of all, the idea of precisification that is underlying the supervaluation semantics for vague terms comes reasonably close to the idea of a conditional probabilities getting ever closer to 1; but of course the semantic clause for $\Box\!\!\rightarrow$ in the Limit semantics is not actually a supervaluationist one. Secondly, Edgington's (1996) theory of vagueness in terms of degrees of truth or "verities" which have a probabilistic structure is modelled after Adams' semantics for conditionals; but then again what we need above is not so much probabilities by which vague terms may be analyzed semantically but rather probability function signs which themselves determine their values vaguely. In any case, we leave the defense of the Limit semantics as a proper semantics for the vague term 'high probability' on the basis of some Limit semantics for vague terms *in general* as an open challenge; for now we will carry on as if this had been achieved. At least it is very plausible that the informal expression 'probability close to 1' that gets used in everyday English is in fact vague; and maybe the formalization of this expression in terms of limiting or infinitesimal probabilities (or Popper function values as understood through Theorem 2.2) may at least be said to be superior to the more common formalization that philosophers usually would think of at that point, that is, of having an exact probability above some exact and predetermined threshold value. Instead, we used such exact threshold values when we defined the approximate truth of counterfactuals (relative to these thresholds) in section 3 of part A.

2.4. The limit assumption. The Limit semantics may be extended or modified in various ways; by representation, some of these moves—though not all—can be made in the Popper function semantics, too, but even then, due to its affinity to Lewis' semantics, the Limit semantics might make these extensions or modifications more easily interpretable in everyday terms.

Just as in Lewis' semantics for counterfactuals, also in the Limit semantics a *Limit Assumption* may be introduced, which may be shown to lead to another equivalent semantics. For every $w \in W$, let $<^w$ be the strict linear preorder that is determined by $\leq^w$ by means of: $u <^w v$ iff $u \leq^w v$ and $v \not\leq^w u$. Then this is what the probabilistic Limit Assumption on Limit models looks like:

- For every $w \in W$, and for every $[\![A]\!] \in \mathfrak{A}$ for which there is an $i \in I_w$ such that $P_i^w([\![A]\!]) > 0$, there is at least one index $i_{min} \in I_w$ which is minimal with respect to $<^w$ among all indices $i \in I_w$ for which it is the case that $P_i^w([\![A]\!]) > 0$, that is, there is no index $i <^w i_{min}$ that has this property.

With the Convergence Assumption still in place, it follows that if i_{min} and i^*_{min} are both minimal among indices $i \in I_w$ for which $P_i^w([\![A]\!]) > 0$, then $P^w_{i_{min}}$ and $P^w_{i^*_{min}}$ must assign the same probabilities conditional on $[\![A]\!]$. Hence, given both the Convergence Assumption and the Limit Assumption, the truth condition for subjunctive conditionals simplifies to:

- $w \in [\![A \mathrel{\Box\!\!\rightarrow} B]\!]$ iff either of the following two conditions is satisfied:
  - There is no $i \in I_w$, such that $P_i^w([\![A]\!]) > 0$.
  - Let i_{min} be any index that is minimal with respect to $<^w$ among all indices $i \in I_w$ for which it is the case that $P_i^w([\![A]\!]) > 0$: then it holds that

$$\lim_{i \in I_w} \left(\frac{P_i^w([\![B \wedge A]\!])}{P_i^w([\![A]\!])} \right) = \frac{P^w_{i_{min}}([\![B \wedge A]\!])}{P^w_{i_{min}}([\![A]\!])} = 1.$$

In light of the following representation theorem, the proof of which builds on the proof of Theorem 2.2, the by now very familiar system V of conditional logic is sound and complete also with respect to the Limit semantics *with* Limit Assumption:

THEOREM 2.3

- *Every family $(P_i)_{i \in I}$ of finitely additive probability measures (on one and the same countable algebra $\mathfrak{A}$) which satisfies the Convergence Assumption with respect to a given linear preorder $\leq$, and which also satisfies the Limit Assumption, represents a Popper function $\mathfrak{P}$ (on $\mathfrak{A}$) where the representation is given by:*

Repr$_2$ *If there is an $i \in I$, such that $P_i(X) > 0$, let i_{min} be any index that is minimal with respect to $<$ among all indices i for which it is the case that $P_i(X) > 0$; then*

$$\mathfrak{P}(Y|X) = \lim_{i \in I} \left(\frac{P_i(Y \cap X)}{P_i(X)} \right) = \frac{P_{i_{min}}(Y \cap X)}{P_{i_{min}}(X)}.$$

Otherwise, $\mathfrak{P}(Y|X) = 1$.

- *Every Popper function (on a countable algebra $\mathfrak{A}$) can be represented by means of* **Repr**$_2$ *by a family of finitely additive probability measures (on $\mathfrak{A}$) which satisfies the Convergence Assumption with respect to some linear preorder $\leq$, and which also satisfies the Limit Assumption.*[26]

Should the Limit Assumption be adopted? This is hard to decide. Assume that all we know about an intended ordered family of absolute probability measures is that it represents a given intended conditional probability measure: then families *with* Limit Assumption do not seem worse off than families for which the Limit Assumption fails; if anything, it might be the other way round, since families with Limit Assumption are easier to handle and thus favorable on pragmatic grounds. Only if there are ways to epistemically access intended ordered families of absolute probability measures which are independent of our access to conditional probability measures—which we doubt—or maybe in light of some general

[26] As in the case of Theorem 2.2, the proof of Theorem 2.3 shows that one could strengthen the second claim ("completeness of representation") of the theorem by replacing 'linear preorder' by 'linear (partial) order' again. The same replacement could also be applied to the first claim ("soundness of representation") of the theorem, but once again that would weaken the claim accordingly.

logical, methodological, or probabilistic arguments, we might have reasons for preferring a Limit semantics with Limit Assumptions to the original Limit semantics (as introduced at the beginning of Section §2). Similar considerations apply *vice versa*.

In the case *contra* the Limit Assumption, the original Lewisian worries about infinitely descending chains of distances between worlds—in our case: absolute probability measures—for which no natural "limits of maximal similarity" seem to be in sight might play a role. Given the Limit Assumption, it would no longer be the case that conditional probabilities could tend towards 1 along an ordered family of absolute probability measures without ever actually being 1 at a final measure on that the antecedent receives positive absolute probability.

In the case *pro* the Limit Assumption, logical considerations such as the following one might be an issue: While in all of our probabilistic semantics finite Agglomeration (the rule of conjunction) is valid, that is,

$$A \mathrel{\Box\!\!\rightarrow} B, A \mathrel{\Box\!\!\rightarrow} C \vdash A \mathrel{\Box\!\!\rightarrow} B \wedge C$$

(as it is in Adams' logic of conditionals), this is not so for its countably infinite version

$$A \mathrel{\Box\!\!\rightarrow} B_1, A \mathrel{\Box\!\!\rightarrow} B_2, \ldots, A \mathrel{\Box\!\!\rightarrow} B_n, \ldots \vdash A \mathrel{\Box\!\!\rightarrow} (B_1 \wedge B_2 \wedge \ldots \wedge B_n \wedge \ldots)$$

even when the countably infinite conjunction in the conclusion were expressible in some way in $\mathcal{L}$. However, the Limit semantics *with* Limit Assumption and *with* the additional assumption of Countably Infinite Additivity for all absolute probability measures involved would even make *countably infinite* Agglomeration valid. (This would not be so for the Limit semantics without Limit Assumption but with Countably Infinite Additivity.) But of course versions of Agglomeration with yet greater cardinality would still fail without adopting further presumptions.[27] For the time being, we regard the status of the probabilistic Limit Assumption as undecided.

One additional thought on the Limit Assumption and the representation theorem above is in order though: if Popper functions are assumed to be counterfactually deterministic at all worlds (recall Section 1.1), then by Theorem 2.3 from above there exist corresponding linearly preordered families of absolute probability measures which also satisfy the Limit Assumption and which determine conditional probabilities of B given A (where defined) which converge to either 1 or 0, in correspondence with the truth value of $A \mathrel{\Box\!\!\rightarrow} B$ at the

[27] Does this point to a feature of the semantics that is problematic more generally, as truth should be preserved by taking conjunctions of *arbitrary* cardinality? Not quite:

$$A \mathrel{\Box\!\!\rightarrow} B_1, A \mathrel{\Box\!\!\rightarrow} B_2, \ldots, A \mathrel{\Box\!\!\rightarrow} B_n, \ldots \vdash (A \mathrel{\Box\!\!\rightarrow} B_1) \wedge (A \mathrel{\Box\!\!\rightarrow} B_2) \wedge \ldots \wedge (A \mathrel{\Box\!\!\rightarrow} B_n) \wedge \ldots$$

is valid for arbitrary cardinalities of premises even in our semantics (if such infinite conjunctions were expressible at all), but it is not clear—though tempting— that this should entail the unrestricted validity of the conjunction rule *within the context of an arbitrary antecedent* as it would be the case with an unrestricted Agglomeration rule for counterfactuals. Should the rule of Agglomeration of counterfactuals be regarded logically valid even for infinitely many premises and indeed for sets of premises of arbitrary cardinality? We doubt that there is a straightforward answer to this question. In any case, in the original Popper function semantics from section 3 of part A, Agglomeration for infinitely premises is invalid, as it is according to Lewis' official semantics in which the Limit Assumption is rejected for philosophical reasons. In our semantics even infinite Agglomeration for conditionals with one and the same *tautological* antecedent is invalid, since in contrast with Lewis' semantics conditionals with tautological antecedents are not logically equivalent to their consequents. (Discussions with Tim Williamson on this topic have been very helpful.)

world in question. Indeed, for every w and every A there will be a minimal index with respect to the ordering of w at which the absolute probability of A is greater than 0, and for all B whatsoever the conditional probability of B given A relative to the measure at that index will be either 1 or 0. But that means that the absolute probability measure for that index may be regarded as nothing else but an indexed truth evaluation or a possible world, and indeed a world that makes A true (since the conditional probability of A given A must be 1). If there are several indices of minimal rank which assign a positive probability to A, then by the representation property the conditional probabilities on A that they determine must all coincide with each other, so we may just as well focus on just one of them, say, i. One could thus view of "that" indexed truth evaluation or world i as the uniquely defined A-world that is maximally similar to w. And if all indices other than such "uniquely determined" indices that are the minimizing indices with respect to some proposition $[\![A]\!]$ are eliminated from the ordered family of absolute probability measures that represents the given Popper function $\mathfrak{P}_w$, then one ends up with a linear partial order of indexed truth evaluations amongst which for every formula A there is a least one in the ordering that satisfies A (if there is one that satisfies A at all). In other words: *the Stalnakerian picture of world-relative selection functions which pick a unique maximally similar A-world w_A for each possible A relative to any given world w in order to determine the truth value of $A \;\Box\!\!\rightarrow\; B$ in w by identifying it with the truth value of B in w_A is then restored.* And of course $A \;\Box\!\!\rightarrow\; B$ is true in w in the thus constructed Stalnaker model if and only if it is true in w in the Probabilistic Limit model, given the Limit Assumption and Counterfactual Determinism.

Vice versa, every Stalnaker model can be turned easily into a Limit model that satisfies the additional constraints of the Limit Assumption and Counterfactual Determinism, by simply regarding the truth evaluation function at each world as an absolute probability measure that is indexed by that world. All the relevant formal properties carry over again, and $A \;\Box\!\!\rightarrow\; B$ is true in a world w in the original Stalnaker model if and only if it is true in w in the corresponding Limit model.

In this sense, Stalnaker's semantics is nothing but the deterministic case of the Limit Semantics—and hence, via representation, of the Popper function semantics—subject to these additional requirements. Taken together with the observation that we made back in section 2 of part A on Skyrms' update operation collapsing into Lewis' imaging operation if given that all worldly conditional chance functions are deterministic, we find that just as the suppositional theory of counterfactuals generalizes the *acceptability pragmatics* of Stalnaker's theory to the indeterministic case, our semantics generalizes the *truth semantics* of Stalnaker's theory to the indeterministic case.[28]

2.5. Representing $\Box\!\!\rightarrow$ in terms of $\Box\!\!\rightarrow_{Lewis}$. In the models of the Limit semantics, the absolute probability measures P_i^w around w are "free-floating" in the sense that they are not necessarily actual absolute probability measures *of worlds in the model* themselves. But suppose the semantics is modified in the way that every probability measure P_i^w itself is the actual absolute probability measure of some world $w' \in W$—so that P_i^w is the least absolute probability measure by the lights of $\leq^{w'}$—and that the linear preorders $\leq^w$ are no longer taken to be orderings of the indices of probability measures but rather of possible

[28] Section 5 of McGee (1994) notes the correspondence between, as he calls them, "decisive" Popper functions (our counterfactual deterministic ones) and Stalnaker models for the case of an Adams-type language with only simple conditionals and with factual formulas without any nonmaterial conditional sign.

worlds. Then this would make the so-amended Limit semantics look even more similar to Lewis' semantics. In fact, one could introduce another condition operator $\Box\!\!\rightarrow_{Lewis}$ into $\mathcal{L}$ and state the truth condition for this operator in standard Lewisian terms on the basis of these world-relative linear preorders for worlds. Moreover, suppose that this extended language even contained object-linguistic operators of the forms

$$P(A) > 0$$

and

$$P(B|A) = 1$$

by which the corresponding properties of absolute and conditional probabilities of statements at worlds relative to the absolute probability measures of those worlds could be expressed.

Given all of that, and given also the Limit Assumption (and of course the Convergence Assumption), it would thus become possible to represent our $A \Box\!\!\rightarrow B$ solely in terms of the Lewisian $\Box\!\!\rightarrow_{Lewis}$ and the P-operators, by means of:

$$A \Box\!\!\rightarrow B =_{df} (P(A) > 0 \Box\!\!\rightarrow_{Lewis} P(B|A) = 1)\,.$$

This confirms the thesis that our semantics is not far at all from Lewis'. It also shows that if Lewis' semantics is extended by semantic clauses for probabilistic operators which track the absolute probability measures that have been assumed to be assigned to worlds, then the resulting Lewisian semantics is structurally richer than ours. For one can define $\Box\!\!\rightarrow$ in terms of $\Box\!\!\rightarrow_{Lewis}$ but not the other way round, since $\Box\!\!\rightarrow$ is implicitly probabilistic and this implicit component cannot be removed anymore such that an analysis of the qualitative $\Box\!\!\rightarrow_{Lewis}$ would be achieved. However, for the same reason, one also needs to tell a more substantial philosophical story of what the intended interpretation of this Lewisian semantic structure is meant to be, especially, the surplus that goes beyond our original probabilistic models, and how we should be able to gain epistemic access to it. It is the task of the following section to point out why we have reasons to believe that our probabilistic semantics stands on firmer grounds in this respect than Lewis'. Furthermore, if the probabilistic Limit Assumption is dropped—maybe because it turns out to be unacceptable ultimately—then one can show that $\Box\!\!\rightarrow$ is in fact not representable on the basis of $\Box\!\!\rightarrow_{Lewis}$. The "Lewisian" representation of our $\Box\!\!\rightarrow$ is also adequate at @ only if the similarity relation at @ that is underlying $\Box\!\!\rightarrow_{Lewis}$ is representing the Popper function at @ in the sense of the representation theorems above; but as pointed out before, similarity in that sense does not necessarily coincide with similarity in Lewis' sense in each and every respect. Finally, it is not the case that using the equivalent Lewisian representation would simplify the semantics of our original $\Box\!\!\rightarrow$ operator: after all, the original metalinguistic translation of $\Box\!\!\rightarrow$ was just in terms of a conditional chance ascription, whereas the Lewisian representation adds to that conditional chance ascription (which now figures in the consequent) an absolute probability ascription in the antecedent and Lewis' $\Box\!\!\rightarrow_{Lewis}$ itself (although it is true that the conditional chance ascription in the consequent is simple now in the sense that it may be understood in terms of the Ratio formula rather than in terms of the Popper function axioms). For these reasons, we prefer not to analyze our $\Box\!\!\rightarrow$ on the basis of some combination of a Lewisian modality and some probabilistic expressions, but we rather leave $\Box\!\!\rightarrow$ as an object language primitive in the sense explained by part A of this article.

One final remark on involving additional probabilistic operators: It is clear that either of our probabilistic semantics may be extended easily to a semantics for expressions of the form $P(B|A) = r$, $P(B|A) > r$, and the like, which are evaluated at worlds in an analogous way to $\Box\!\!\rightarrow$, and for which therefore iterations of P and applications of propositional operators to P-formulas are unproblematic. For instance, in the Limit semantics in the nontrivial case, the truth of $P(B|A) = r$ at w would amount to the convergence of $\left(\frac{P_i^w([\![B \wedge A]\!])}{P_i^w([\![A]\!])}\right)_{i \in I_w}$ towards r; that is: for all $\epsilon > 0$ there is an index $j \in I_w$ with $P_j^w([\![A]\!]) > 0$, such that for all $i \leq^w j$ with $P_i^w([\![A]\!]) > 0$ it holds that $\left|\frac{P_i^w([\![B \wedge A]\!])}{P_i^w([\![A]\!])} - r\right| < \epsilon$. This defining condition is not so far from Lewis' suggestion of how to formally represent statements of the form 'the probability that C, if it were the case that A, would be r' in Lewis (1973a, p. 22) under the assumption that the absolute probability measure at a world tracks the sphere system of that world: Lewis says that such a probabilistic statement "is true at a world i (with respect to a given probability measure) iff for any positive ϵ there exists an A-permitting sphere T around i such that for any A-permitting sphere S around i within T, $Prob(C|AS)$, unless undefined, is within ϵ of r."[29] However, Lewis is simply wrong when he also claims there that 'the probability that C, if it were the case that A, would be r' *cannot* be analyzed in terms of an expression of the form

$$P(C|A) = r,$$

as should be clear from the rest of our paper.[30]

2.6. Lewis' semantics naturalized. We conclude this section with some tentative thoughts on the philosophical status of our semantics vis-à-vis Lewis', when one takes a naturalistic viewpoint—the viewpoint of *philosophy as being continuous with science*—on the semantics and metaphysics of counterfactual conditionals.

Even though Lewis' semantics for counterfactuals may be regarded as being part of a grand overall metaphysical theory which comes in touch with science at various places, it does not seem as if science were able to tell us a lot about what Lewis' intended models for counterfactuals ought to be like. Indeed, Lewis' explanations of comparative similarity for possible worlds seem to be based on a priori intuition only. One cannot help wondering to what extent the prioritized list of features by which he claims one world to be closer to the actual world than another is arbitrary and merely the result of overfitting some more or less plausible examples from natural language. How would a scientist be able to integrate Lewis' intended model(s) into his or her scientific theories in any systematic manner?

This is where our probabilistic semantics for counterfactuals works to our advantage. The Popper function semantics is based solely on the notion of objective conditional chance—something modern science presupposes anyway, in one way or another. Of course, we do admit that objective Popper functions in our sense are not quite the immediate objects of empirical study themselves: if Popper functions are to be based on the set of metaphysically possible worlds, then this goes probably beyond science (though not so if they are meant to be defined on the set of physically possible worlds); scientific theories normally do not derive chance statements of the form 'the probability that B would have

29 By '$Prob(C|AS)$' Lewis means the probability of C conditional on the intersection of S with the proposition expressed by A.

30 See Barker (1999) for an analysis and criticism of some of Lewis' suggestions for how to treat probabilistic counterfactuals.

happened given that *A* had happened' where *A* describes some remote event in the past (though it certainly does so for chance statements of the form 'the probability that *B* would happen given that *A* were to happen' where *A* describes some physically possible event in the future); scientific theories normally do not predict conditional chances of everyday events such as a student coming to a lecture given the student met some particular other student on the way to the lecture hall (though it does so for chances of laboratory events such as a particle showing some particular kind of behavior given the particle hit some particular other particle on its way to the measurement device); in scientific theories conditional probability measures get applied to propositions which describe events at precisely determined points of time or within precisely determined intervals of time, not propositions which involve complete world histories "close to, but not including" some point of time. And so forth. And of course we do not claim that conditional chances in our sense could be measured empirically in any immediate manner (though scientific theories that involve the concept of chance can be confirmed empirically). But even if all these problems are taken into account, the Popper function semantics still seems to be closer to science than Lewis' semantics could ever be, for two reasons: according to the former—but not according to the latter—counterfactuals track structures which are "out there" even in the eyes of the scientist; and the Popper function semantics is based on a concept of conditional probability that is very much like the scientist's own concept of conditional probability whereas there does not seem to be any clear scientific counterpart to Lewisian similarity or closeness.

If so, the following naturalization of Lewis' project ought to be advisable and possible in principle: Let science—and substantial philosophical extrapolations from science[31]—determine the conditional probability measures in an intended Popper function model a posteriori. Apply Theorem 2.2 in order to determine a Limit model in which these Popper functions are represented by ordered families of absolute probability measures. Consider the resulting Limit semantics as the closest possible approximation of the Lewisian semantics within the limits of science. Perhaps a Popper function semantics that is supported by science in such a way can thus be used *in principle* to naturalize the semantics and metaphysics of a quasi-Lewisian semantics for counterfactuals? And then the conceptual framework of Lewis' theory will allow us *in principle* to interpret scientific facts in a manner akin to the manifest image?[32]

§3. An application: are most ordinary counterfactuals false? Finally, we turn to an application of our probabilistic semantics to a concrete philosophical question. This application will highlight at the same time some of the merits and some of the limitations of this semantics.

The following question on counterfactuals has received a lot of attention lately: *Do most of the counterfactuals that we use in ordinary conversation happen to be false?* This is the

---

[31] Maybe even some "intuition-driven fine-tuning", using a phrase from Yablo (2009).

[32] Williams (2008) may be read as a different attempt to give Lewisian similarity orderings a more naturalistically acceptable interpretation: Williams keeps the thought of ordering worlds in terms of a Lewis-style heuristics for similarity to the actual world, but he suggests to analyze Lewis' avoidance criterion of lack of "fit with laws of nature" in terms of a probabilistic notion of randomness or typicality which he derives from Elga (2000) (who in turn derives it from formal work by Gaifman and Snir). Williams also argues that Hawthorne's (2005) arguments against Lewis' semantics (see next section) no longer apply to the so-amended semantics.

central topic of Hájek (unpublished) and partially the topic of Hawthorne (2005);[33] we will discuss their arguments simultaneously, and we are going to react to them in stages.

3.1. Stage 0 of the dialectic. At least as far as Hájek is concerned, the answer to the question above is an accentuated YES![34] Here is the reason given for this answer: Quantum mechanics tells us (see Hawthorne, 2005; Hájek, unpublished, sec. 2),

> If I had dropped the plate, it might have flown off sideways

since according to quantum physics our world is indeterministic, so that there was a small probability for the plate to fly off sideways given it had been dropped (even when the plate actually fell to the floor). But then the following subjunctive conditional must be false:

> If I had dropped the plate, it would have fallen to the floor.

But most of the typical subjunctive conditionals in everyday conversation—and, for that matter, in philosophy—are like that; so most ordinary counterfactuals are false.

Hájek (section 4) also gives a corresponding argument for the deterministic case based on statistical mechanics and imprecise or underspecified antecedents:

> If I were to drop the plate, it would fall to the floor

is imprecise in the sense that the manner in which I were to drop the plate is left open. Consequently, there is a huge range of initial conditions which all satisfy the qualification 'dropping the plate', among which there will be some initial conditions which deterministically give rise to anomalous trajectories that lead to the plate doing anything but falling to the floor. Hence, the subjunctive conditional is false again, even if our world happens to be deterministic, and our language is likely not to be fine-grained enough semantically in order to exclude the anomalous trajectories by adding to the antecedents of those conditionals.

Why is it that a supporter of Lewis' semantics would be in trouble to fend off examples like these—examples which, if successful, would threaten to trivialize the evaluation of most standard counterfactuals? What a Lewisian would like to say, presumably, is that in *sufficiently close* worlds in which the plate had been dropped (were to be dropped), it *did* (*would*) fall to the floor. But then the corresponding 'might'-conditionals would have to be false, which they do not seem to be. Accordingly, assume with Lewis that worlds with bizarre low-probability events are by and large more distant from the actual world than worlds without such events: But bizarre events are possible to happen, sometimes even with nonnegligible probability. Using an example by Hawthorne (2005, p. 399): Assume a monkey at a typewriter is so configured that there would be a chance of 0.20 of it to type a novel. You take away the typewriter, nothing remarkable happens; the counterfactual

> If you had not taken away the typewriter, the monkey would not have typed a novel

is false. However, each *particular way* for a bizarre thing to happen—each particular way for the monkey to write a novel—will correspond to a bizarre *low-probability* event that may only take place in remote worlds, in line with the Lewisian proposal. But if all the

---

[33] Some of their points have been made before by Edgington (1995, pp. 250, 252, 258, 299, 321), Edgington (2004, p. 14f), and Bennett (2003, sec. 16).

[34] The conclusion section of Hawthorne (2005) does in fact hint at ways of how to reject the thesis that most ordinary counterfactuals are false from the perspective of either a contextualist variant of Lewis' semantics or of a non-Lewisian, Stalnakerian semantics for counterfactuals; but neither suggestion is worked out in any detail.

particular unlikely ways for a bizarre thing to happen can only be instantiated in distant worlds, then also the bizarre events of reasonably high probability themselves can only inhabit distant worlds, and the monkey counterfactual above should be *true*, which it is not. As both Hawthorne and Hájek note, this is not just reasoning about possible but ultimately unrealistic circumstances: after all, bizarre things happen all the time! Indeed, maybe for every possible world there is some respect in which that world is strangely remarkable. Hence, on the basis of Lewis' semantics one is in serious trouble to prove the arguments from above unsound, and it remains quite plausible to believe that most ordinary counterfactuals are false. Conversation and philosophical theorizing in terms of subjunctive conditionals is doomed, or so it seems. Call this stage 0 of the dialectic.[35]

3.2. Stage 1 of the dialectic. Now we start to invoke the resources of our probabilistic semantics. Are most ordinary counterfactuals false? NO!

First of all, reconsider

If I had dropped the plate, it might have flown off sideways.

We have not dealt with might-conditionals in this paper at all. And this was for good reasons: the conditional above seems to refer to a possibility (here: one in the past), that is, the possibility of the plate flying sideways if dropped; but our theory does not have to say much about possibility over and above what standard modal semantics teaches us. In particular, it is clear that might-conditionals should *not* be represented as the duals of the would-conditionals that obey our semantics, that is, in terms of $\neg(A \mathrel{\Box\!\!\rightarrow} \neg B)$ where $\Box\!\!\rightarrow$ is understood probabilistically, as this would not enable us to distinguish between the unlikely possibility of B given A and its plain impossibility. At best, one may turn to

The plate is dropped (at t) $\Box\!\!\rightarrow$ $\Diamond$ it flies off sideways (shortly after t),

thus understanding the might-conditional from before as expressing that dropping the plate would have led to a true possibility with very high probability. As Hawthorne (2005, p. 398f) points out, a similar move can be made by a proponent of Lewis's semantics on the basis of the Lewisian $\Box\!\!\rightarrow$. (Lewis himself considered this an option for 'might' conditionals.) But while this would yield the truth of the might-conditional above as intended, it is still not clear at all that this step just by itself would help the Lewisian with saving

If I had dropped the plate, it would have fallen to the floor

from being false, since it is still unclear why and how all closest antecedent-worlds should be consequent-worlds if at the same time the chance of the consequent given the antecedent is close to, but less than, 1. What magic trick would prohibit the odd exceptional $A \wedge \neg B$-world to be among the closest A-worlds?

It is the latter 'would'-counterfactual where our probabilistic semantics has bite (at least at first glance, for the dialectic will continue): In order for this counterfactual to be true according to our probabilistic truth condition, it suffices for there to be a *very high* conditional probability of the plate falling to the floor if dropped. But that conditional probability *is* in fact high—exceptions to it having a probability close to 0—even by the lights of quantum physics. So the counterfactual comes out true in our semantics according to our second interpretation of '$\mathfrak{P}([\![B]\!] | [\![A]\!]) = 1$' in terms of 'the probability of B given

[35] Hawthorne (2005) presents two additional related problems for Lewis' semantics, and Hájek (unpublished) criticizes Lewis' semantics not just on basis of the presented arguments but also on independent grounds.

A is very high, or close to 1', which we justified in Section 2.3. By generalization, it is not the case that most ordinary counterfactuals may be expected to be false. And for the approximate truth version of our semantics, the worries should disappear even more straightforwardly, since then conditional chances greater than or equal to $1 - \alpha$ with some fixed real $\frac{1}{2} < 1 - \alpha \leq 1$ are sufficient for the approximate truth of $A \mathrel{\Box\!\!\rightarrow} B$ (to degree $1 - \alpha$).

Indeed, Hawthorne (2005, p. 397f) considers replacing Lewis' truth condition by

$A \mathrel{\Box\!\!\rightarrow} B$ is true iff most of the closest A-worlds are B-worlds

as one possible solution strategy (the other one being Lewis' "bizarre low-probability worlds are far off" strategy, which Hawthorne rules out in detail). Obviously, this is close enough to our semantic rule for counterfactuals. He grants that the threshold that corresponds to 'most' will be vague and that this should not cause any concerns. But then he criticizes this strategy on the basis of the claim that Agglomeration, that is,

$$A \mathrel{\Box\!\!\rightarrow} B, A \mathrel{\Box\!\!\rightarrow} C \vdash A \mathrel{\Box\!\!\rightarrow} B \wedge C$$

would be invalid in the resulting semantics. As he says, "Agglomeration is overwhelmingly intuitive" (p. 397), so this would be no plausible way out. However, as we have seen in Section 2.4, Agglomeration is indeed *valid* in our truth semantics, and this is so precisely because the threshold that corresponds to 'high probability' or 'probability close to 1' is vague (the vagueness being explained in Section 2.3). On the other hand, according to the approximate truth part of our semantics Agglomeration *does* fail to be valid, as suggested by Hawthorne; but that is for reasons very similar to those for which logical conclusions from sets of approximately true scientific hypotheses may fail to be approximately true even independently of reasoning in terms of counterfactuals. For instance: say we have good empirical reasons for believing that (i) $\forall x(P(x) \rightarrow f(x) = r)$ is true, and we have equally good independent reasons for believing also (ii) $\forall x(g(x) = f(x) + s)$ to be true. Assume that in fact these empirical hypotheses are only true up to an α, that is, really (i') $\forall x(P(x) \rightarrow f(x) \in [r - \alpha, r + \alpha])$ and (ii') $\forall x(g(x) \in [f(x) + s - \alpha, f(x) + s + \alpha])$ are true, while (i) and (ii) are actually not. Being good scientists we draw the conclusion (iii) $\forall x(P(x) \rightarrow g(x) = r + s)$ from (i) and (ii) by logically valid inference, even though the actual truth might be not so much that (iii') $\forall x(P(x) \rightarrow g(x) \in [r + s - \alpha, r + s + \alpha])$ but rather (iii') $\forall x(P(x) \rightarrow g(x) \in [r + s - 2\alpha, r + s + 2\alpha])$. So (iii) might not even be approximately true anymore, or at least not with respect to the same threshold value as (i) and (ii), even though (iii) did follow logically from approximately true premises. Similarly, Agglomeration does not necessarily preserve approximate truth with respect to a fixed threshold, even though its conclusion, which is a counterfactual, follows logically from its counterfactual premises. Summing up: Hawthorne's argument by failed Agglomeration against the semantic "most worlds" strategy of dealing with ordinary counterfactuals from above—which is close to our strategy—does not apply to our truth semantics, and while it does apply to our approximate truth semantics it is unclear why this should be more worrisome than the failed preservation of the approximate truth of noncounterfactual scientific statements by logically valid inferences.[36]

[36] Kvart (2001, footnote 91) also mentions that Agglomeration is invalid in a "high conditional probability" semantics for counterfactuals that is based on presupposing a precise real number threshold less than 1 again. As noted in Section 2.4, *infinite* versions of Agglomeration are indeed invalid even in our *truth* semantics, as long as no additional assumptions such as the

Hájek (unpublished, sec. 8) acknowledges that "In the neighborhood of the ordinary but false counterfactuals that we utter, there are closely related counterfactuals that are true but not ordinary. They are counterfactuals with appropriate probabilistic or vague consequents." In our Section 2.5 we have shown that under certain conditions one can in fact view our $\Box\!\!\rightarrow$-conditionals as Lewisian counterfactuals with probabilistic consequents (and antecedents). So subjunctive conditionals in our sense are not far from Hájek's "closely related counterfactuals that are true but not ordinary," the only differences being: we suggest to leave the reference to probabilities implicit in the semantics of $\Box\!\!\rightarrow$, whereas Hájek thinks of them as being made explicit; the conditional probabilities that are referred to implicitly by our $\Box\!\!\rightarrow$ are described to lie above a vaguely determined threshold in our truth semantics, whereas Hájek mostly considers the case of probabilities which are said to lie above an exact threshold just as in our semantics for approximate truth. Indeed, there is no real contradiction with Hájek's thesis anyway: While he wants to show that *from the conceptual analysis of ordinary counterfactuals* (and science) it follows that most ordinary counterfactuals are false, our aim (as explained in section 4.1 of part A) was to demonstrate that *on the basis of what the truth conditions of ordinary counterfactuals ought to be like* (and science) it does not follow that most ordinary counterfactuals are false. It does not seem that Hájek would deny this; so we offer our semantics as a way of saving the truth evaluation of ordinary counterfactuals from trivialization. Deviating from conceptual analysis in this way is a price worth being paid. Call this Stage 1 of the dialectic.

3.3. Stage 2 of the dialectic. But that is not quite the end of the story. Enter Hawthorne and Hájek again: The answer to the question above is still YES! This is because there is a problem at least with the new probabilistic truth (rather than approximate truth) semantics, as may be seen from considering lottery examples, as both Hawthorne and Hájek do.

We take the version of Hájek (unpublished, sec. 2). Assume there are 1000001 tickets; you do not play the lottery. It holds:

Lottery is played $\Box\!\!\rightarrow$ ticket #1 loses
Lottery is played $\Box\!\!\rightarrow$ ticket #2 loses
⋮
Lottery is played $\Box\!\!\rightarrow$ ticket #1000001 loses

all seem to be true, according to common sense as applied "in the midst of ordinary thought and talk" (as Hawthorne, p. 400, puts it), or according to an exact probabilistic threshold of 0.999999 above which conditional probability turns into truth (in Hájek's version). However, Agglomeration leads to plain falsity, since one of the lottery tickets must win. Hence, in a semantics like ours there are really only two options: either not all of the conditionals above are evaluated true, in contradiction with common sense and with what one would get from a reasonably high exact threshold; or we do evaluate all of them as true, which means that we will then have to evaluate their conclusion by Agglomeration as true even though it is clearly a falsity pretheoretically. Either option is unacceptable. Our new truth semantics, for which Agglomeration is valid, simply does not capture everyday talk, since the underlying transition from a quantitative concept, that is, probability, to a qualitative one such as $\Box\!\!\rightarrow$ does not harmonize with the quantitative

Limit Assumption and Countably Infinite Additivity are employed. But then again it is unclear whether infinite Agglomeration ought to count as logically valid; see the corresponding Footnote 27 in Section 2.4.

nature of the lottery example above, nor with the quantitative nature of quantum reality as in the original examples of false ordinary counterfactuals. Hence, not even our probabilistic truth semantics for counterfactuals can save typical counterfactuals from being false. Call this Stage 2 of the dialectic.

3.4. Stage 3 of the dialectic. Of course, there is a Stage 3, too. (This is philosophy after all.) The answer to the question above is: NO(T SO FAST)!

We admitted already that our approximate truth semantics for counterfactuals does not support Agglomeration; it would count the premises, but not the conclusion, of Hájek's argument in the last subsection as true (for the relevant $1-\alpha$ threshold). Our truth semantics does support Agglomeration, so it could not count all the premises of Hájek's argument as true. So where does this leave us? We still do not lose heart, based on three independent thoughts:

First of all, maybe evaluating each of the premise counterfactuals in the lottery case example as false—as predicted by our truth semantics—is actually what we ought to do, regardless of what common sense or approximate truth relative to an exact threshold of 0.999999 seem to recommend. After all, if one pushes people sufficiently, going through the premises one after the other, and asking for each of them 'Is that counterfactual really true?', one might in fact get them to admit eventually: *No, they are all false*. But then the example ceases to constitute a problem for a semantics such as ours. No vaguely determined threshold plays a role here; since the exact conditional probability of any premise consequent given the corresponding premise antecedent in the lottery example falls short of being precisely 1, the premises should all be false.

Secondly, perhaps finite lotteries—with their precisely determined probabilities—are not "ordinary" at all, and neither are events involving single quantum particles. If so, then none of the counterfactuals considered at the last stage are ordinary either, and it has not been established that most *ordinary* counterfactuals are false. Maybe the conditional probabilities for everyday macroscopic events are in fact only vaguely determined. Then describing such events by means of conditionals which are made true by "very high" probabilities, where 'very high' is a vague term, is precisely what one ought to do, since only they can express the vaguely determined probabilities of such events adequately. For instance, say, we would be considering a large finite number of ways in which the plate might fall to the floor, such that every possible maximally specific way of it falling to the floor is subsumed by one of these finitely many ways, and the conditional probability of falling to the floor if dropped is distributed uniformly over these ways of falling:[37] then it could not be the case that each of the

> If I had dropped the plate, it would have fallen to the floor but not in the way so-and-so

conditionals were true since their Agglomeration would lead to certain falsity. But now our claim would be that neither of these conditionals is ordinary anymore, unlike the original plate-dropping counterfactual which is ordinary and the conditional probability for which might actually be close to 1 without having a precise real number value.[38]

Is there reason to believe that the probability of a macroscopic event allows for borderline cases for which it is neither clear whether the event has that probability nor that

[37] We owe this example to Richard Pettigrew.

[38] We thank Richard Pettigrew and Robbie Williams for very helpful and clarifying discussions on this matter.

it does not have it, and that certain macroscopic events have a probability that is high without there being any definite and exact cutoff point which would distinguish between high probabilities and ones that are not high? This is not out of the question: As Wallace summarizes the view he develops at the end of Wallace (2003), "Macroscopic objects are to be understood as structures and patterns in the universal quantum state ... We can tolerate some small amount of imprecision in the macroworld: a slightly noisy pattern is still the same pattern." So maybe from the viewpoint of our best empirical theory—quantum physics— macroscopic entities are in fact *ontically* vague (see also French & Krause, 2003, on "Quantum Vagueness"); but if these entities are vague, then maybe the probabilities of events in which they are involved are so as well.[39] Alternatively, one might argue that although quantum theory itself is a theory which involves only exact probabilities, its application to the macroscopic realm might not be so. Every application of theoretical laws, including probabilistic ones, to the real world is subject to an enormous number of *ceteris paribus* clauses (cf. Cartwright, 1983). Therefore, if quantum physics gets applied to dropped plates which might or might not fall to the floor, then no precise probabilities for such events can be derived at all—at best, the probabilities that one can predict by theory and practical experience are "smeared out" by relevant factors which cannot all be made explicit in the theory. In that case, the events themselves might be crisp, but the theoretical concept of probability as applied to macroscopic events would still be vague. In either case, the vague very high conditional probability statement that is expressed by

If I had dropped the plate, it would have fallen to the floor

might describe reality correctly after all. Since this argument relies on (particular interpretations of) quantum theory to be true, this will only amount to a defense of our probabilistic truth condition within the realm of physically possible worlds, since it is in these worlds, presumably, that the laws of quantum physics hold as laws of nature. The defense might not work more generally as applied to arbitrary metaphysically or even conceptually possible worlds.[40] In fact, the situation might even be "worse": the argument might only go through for counterfactuals as being evaluated in the actual world. It might be a contingent fact that our ordinary counterfactuals implicitly refer to vaguely determined probabilities in a way such that they are not all evaluated false by our probabilistic semantics. If our world or our language were different, then they might well all be false as highlighted by Hájek. But we would be satisfied already if our semantics did just a good job in the actual world, so that would be fine with us.

Finally, here is the third thought: Say, neither are macroscopic entities vague nor the concept of probability is. So the right scale level of descriptions of plates being dropped and falling to the floor is indeed the exact quantitative one. Still counterfactuals might have a role as *coarse-grained* qualitative descriptions of the actual underlying quantitative reality. The ultimate purpose of speaking in terms of counterfactuals would thus be to push down the complexity of quantitative statements such as 'the probability of B given A is 0.98' to that of qualitative statements of the form $A \mathrel{\Box\!\!\rightarrow} B$ by changing the scale

[39] We thank James Ladyman for pointing us to this literature.
[40] We thank Alan Hájek for making us aware of this.

type. If so, this would not mean that no quantitative aspect of reality whatsoever would get preserved in the transition. In particular, if

$$A_1 \Box\!\!\rightarrow B_1, \ldots, A_n \Box\!\!\rightarrow B_n \vdash C \Box\!\!\rightarrow D$$

is valid in the Popper function semantics, then one can show that $C \Box\!\!\rightarrow D$ is derivable from the premises by means of the axioms and rules of the flat fragment of the system V of conditional logic, and hence by Adams' results in Adams (1975) the *exact* uncertainty of the conclusion is bounded from above by the sum of the *exact* uncertainties of the premises, where the uncertainty of $\varphi \Box\!\!\rightarrow \psi$ is defined as $1 - \mathfrak{P}([\![\psi]\!] \mid [\![\varphi]\!])$ (and where it is assumed that the exact absolute probabilities of all relevant antecedents are greater than 0). Hence, an inference that is valid for qualitative counterfactuals according to our semantics would be inferentially *reliable* even when the conditional probabilities of the premises in question are exact. It would be reliable in the sense that it comes with a guarantee: if the exact uncertainties of the premises are so and so, then the exact uncertainty of the conclusion lies within the exact boundaries of so and so. This does not mean, of course, that the conditional probability that goes with the conclusion could not go down drastically, but it means that one can control how much it can go down; one knows with certainty that given those premise boundaries and that number of premises, the uncertainty of the conclusion cannot be worse than a certain precise number.[41] It would be up to the reasoner to apply such valid inferences sensibly, but at least these valid inferences would allow the reasoner to apply them sensibly at all, as they have a quality seal attached to them that can be made explicit in the form of Adams' sum rule. Accordingly, counterfactuals would remain to be useful vehicles of communication whenever one does not have complete information about the actual precise probabilities, as it is the case both in everyday life and in philosophical theorizing. (Or what is the precise conditional probability of a plate falling to the floor if dropped?)

If this third thought—rather than the first or the second one from before—was right, then Hájek (unpublished, sec. 9) would also be right when he concludes that there are in fact various counterfactuals which we can legitimately *assert*, but which nevertheless are false as in fact they would be then also according to our probabilistic truth semantics. On the other hand, they would also still be true *approximately* according to the second part of our semantics. Ordinary counterfactuals would thus be in the very same ballpark as many of our scientific hypotheses: not true but approximately true. One might think, however, that there would still be a difference between proper scientific statements and such counterfactuals: once scientists have realized that their hypotheses are "merely" approximately true, they change their hypotheses, and maybe even the theoretical language in which they were formulated, in favor of hypotheses which they hope to get "closer to the truth." On the other hand, we seem to be stuck in natural language with our ordinary counterfactuals, and no passage in this paper seemed to recommend to replace them by something else in everyday contexts. But if the current thought is right, shouldn't we rather stop talking in terms of ordinary counterfactuals, accordingly, and make precise probabilistic statements in their place—statements which would be *literally* true then rather than merely approximately true? As we argued before, this move would be unwise: ordinary subjunctive conditionals are not just approximately true, they are also very useful, and they are so for exactly the reason that they do not syntactically involve any numerical value ascriptions. (For the same reason, their truth semantics is governed by simple, elegant, and computationally tractable

[41] See Schurz (2005) for such an interpretation of Adams' logic in terms of probabilistic reliability.

logical axioms and rules.) Scientists very often do not replace theories that are merely approximately true either if they prove to be very useful in many different contexts: For instance, Newtonian mechanics is false for very principled reasons, namely, because the world is both relativistic and indeterministic, but: most of its applications are still approximately true, no scientist would want to apply relativity theory or quantum mechanics to ordinary everyday physics contexts instead, and whenever they apply Newtonian mechanics in these contexts they usually don't distinguish between its truth and approximate truth anymore and simply proceed as if their Newtonian claims were literally true. Following the scientists' lead on this, we should not care too much then either and still speak of counterfactuals as being true when we know they are strictly speaking not, as long as we do not get ourselves in a position in which the difference between truth and *mere* approximate truth starts to matter (as it does for Newtonian mechanics, for instance, on the astrophysical or on the atomic level). Adams' methods of uncertainty estimation will then hopefully facilitate a controlled retreat. In brief: if asserting ordinary counterfactuals in everyday contexts ends up being just as good or bad as making Newtonian predictions in everyday scientific contexts, we should be on the safe side at last.

§4. Summary. We have given a theory of subjunctive conditionals. We proposed a semantics of world-relative primitive conditional chance measures from which the truth or the approximate truth of counterfactuals at possible worlds could be determined; by a representation theorem we showed that there was an equivalent semantics which one might regard as the probabilistic version of David Lewis' semantics for counterfactuals. We concluded that our theory amounted to a naturalization of Lewis' theory. The resulting logic for our probabilistic truth semantics turned out to be a well-known system of conditional logic which does not include Lewis' Centering axioms; we explained why this was so, how it would be possible to get these axioms back, and why it would not be so worrisome if one refrained from doing so. According to our probabilistic semantics, subjunctive conditionals contain implicit *ceteris paribus* clauses. That was put to use when we answered the question 'Are most ordinary counterfactuals false?' by a qualified 'No', at least as far as the approximate truth of such counterfactuals was concerned. We dispelled various formal and philosophical concerns that a semantics like this might cause; and we considered ways of extending or modifying the semantics for various purposes. On the pragmatic side, we suggested that counterfactuals come with two kinds of pragmatic meanings which correspond to two types of degrees of acceptability or belief, one of which was in line with the suppositional theory of counterfactuals, whereas the other one was derivative from our truth-conditional semantics. Although these two types of degrees could not generally coincide in value in light of a new triviality result on expected chance, as far as the mere assertability of counterfactuals goes, the difference between the two kinds of pragmatic meanings of counterfactuals was unlikely to show up.

Counterfactuals are not to be eliminated from our manifest image of the world in favor of purely scientific descriptions, rather our probabilistic semantics is meant to build a bridge between the manifest and the scientific image. If our theory is a good one up to exceptions, then that is all we can ask for.

§5. Acknowledgements. For their invaluable comments we want to thank Horacio Arló-Costa, Johan van Benthem, Dorothy Edgington, David Makinson, Brian Skyrms, Tim Williamson, Steve Yablo, as well as Frank Arntzenius, Alexander Bird, Franz Dietrich, Richard Dietz, Igor Douven, Paul Egré, David Etlin, André Fuhrmann, Alan Hájek, Vincent Hendricks, Leon Horsten, James Ladyman, Barry Loewer, Thomas Müller, Daniel Nolan, Rohit Parikh, Richard Pettigrew, Richard Bradley, Jan-Willem Romeijn, Daniel Rothschild, Hans Rott, Gerhard Schurz, Wolfgang Spohn, Robbie Williams, and members of the audiences at talks in Bristol, Copenhagen, Rotterdam, Oxford, Konstanz, Irvine, Manchester, Utrecht, Frankfurt, Düsseldorf, Princeton, Regensburg, London, Sheffield. We would like to thank the Leverhulme Trust and the Alexander von Humboldt Foundation for their generous support of this work.

§6. Appendix: proofs. The following four proofs relate to results stated in Section 1.1; we only give sketches:

Proof of: Actual World Determinism *corresponds to* Centering. Soundness can be proven directly. Completeness follows from Lewis' completeness proof for VC in Lewis (1973b) and the construction in the proof of theorem 4 in part A. □

Proof of: 'For all $w \in W$*:* $\mathfrak{P}_w([\![s_w]\!]|W) = 1$*' corresponds to* Centering. Soundness can be proven directly again. Completeness also follows from Lewis' completeness proof for VC in Lewis (1973b) and the construction in the proof of theorem 4 in part A. □

Proof of: 'For all $w \in W$*:* $\mathfrak{P}_w([\![s_w]\!]|W) > 0$*' corresponds to* Weak Centering. Soundness can be proven directly again. Completeness also follows from Lewis' completeness proof for VW in Lewis (1973b) and the construction in the proof of theorem 4 in part A. □

Proof of: Counterfactual Determinism *corresponds to* Conditional Excluded Middle. Soundness may be proven directly. Completeness follows from Lewis' completeness proof for V+CEM in Lewis (1973b) and the construction in the proof of theorem 4 in part A. □

These are the proofs of theorems in Section §2:

Proof of Theorem 2.2.

- Let $(P_i)_{i\in I}$ be given, and let $\mathfrak{P}$ be defined as explained in the theorem.
  We prove the Multiplication Axiom 4:
  If there is no $i \in I$, such that $P_i(Z) > 0$, then there cannot be an $i \in I$, such that $P_i(X\cap Z) > 0$ either; hence, in that case, $\mathfrak{P}(X\cap Y|Z) = 1 = \mathfrak{P}(X|Z)\,\mathfrak{P}(Y|X\cap Z)$. On the other hand, assume that there is an $i \in I$, such that $P_i(Z) > 0$: If there is no $i \in I$, such that $P_i(X \cap Z) > 0$, then $\mathfrak{P}(X \cap Y|Z) = \lim_{i\in I}\left(\frac{P_i(X\cap Y\cap Z)}{P_i(Z)}\right) = 0 = \lim_{i\in I}\left(\frac{P_i(X\cap Z)}{P_i(Z)}\right) = \mathfrak{P}(X|Z)$ and the axiom must hold again. So suppose there is an $i \in I$, such that $P_i(X \cap Z) > 0$: then $\mathfrak{P}(X \cap Y|Z) = \lim_{i\in I}\left(\frac{P_i(X\cap Y\cap Z)}{P_i(Z)}\right) = \lim_{i\in I}\left(\frac{P_i(X\cap Z)}{P_i(Z)} \cdot \frac{P_i(Y\cap X\cap Z)}{P_i(X\cap Z)}\right)$. By the Convergence Assumption and the case we are considering, the limits of the two factors must exist, hence by general properties of limits: $\lim_{i\in I}\left(\frac{P_i(X\cap Z)}{P_i(Z)} \cdot \frac{P_i(Y\cap X\cap Z)}{P_i(X\cap Z)}\right) = \lim_{i\in I}\left(\frac{P_i(X\cap Z)}{P_i(Z)}\right) \cdot \lim_{i\in I}\left(\frac{P_i(Y\cap X\cap Z)}{P_i(X\cap Z)}\right) = \mathfrak{P}(X|Z)\,\mathfrak{P}(Y|X \cap Z)$, and we are done.
  Secondly, we prove Axiom 5: Assume that $\mathfrak{P}(X|Y) = \mathfrak{P}(Y|X) = 1$. There are four conceivable cases: (i) There is no $i \in I$, such that $P_i(Y) > 0$, and there is no $i \in I$, such that $P_i(X) > 0$: then for all Z, $\mathfrak{P}(Z|X) = 1 = \mathfrak{P}(Z|Y)$. (ii) There is no $i \in I$, such that $P_i(Y) > 0$, but there is an $i \in I$, such that $P_i(X) > 0$: but in this case $\mathfrak{P}(Y|X) = \lim_{i\in I}\left(\frac{P_i(Y\cap X)}{P_i(X)}\right)$ would have to be 0, contradicting our assumption.

(iii) There is an $i \in I$, such that $P_i(Y) > 0$, but there is no $i \in I$, such that $P_i(X) > 0$: analogous to (ii). (iv) There is an $i \in I$, such that $P_i(Y) > 0$, and there is an $i \in I$, such that $P_i(X) > 0$: then there must be an $i \in I$, such that $P_i(Y \cap X) > 0$, for otherwise $\mathfrak{P}(Y|X) = \lim_{i \in I}\left(\frac{P_i(Y\cap X)}{P_i(X)}\right)$ could not be 1. Thus, $\mathfrak{P}(Z|X) = \lim_{i\in I}\left(\frac{P_i(Z\cap X)}{P_i(X)}\right) = \lim_{i \in I}\left(\frac{P_i(Z\cap X \cap Y)}{P_i(X)} + \frac{P_i(Z\cap X\cap[W\setminus Y])}{P_i(X)}\right) = \lim_{i\in I}\left(\frac{P_i(Z\cap X\cap Y)}{P_i(Y\cap X)} \cdot \frac{P_i(Y\cap X)}{P_i(X)} + \frac{P_i(Z\cap X\cap[W\setminus Y])}{P_i(X)}\right) =$, since the limits of all summands and factors exist by the Limit Assumption and our previous remarks, $= \lim_{i\in I}\left(\frac{P_i(Z\cap X\cap Y)}{P_i(Y\cap X)}\right) \cdot \lim_{i\in I}\left(\frac{P_i(Y\cap X)}{P_i(X)}\right) + \lim_{i\in I}\left(\frac{P_i(Z\cap X\cap[W\setminus Y])}{P_i(X)}\right) =$, by our assumption at the very beginning, $= \lim_{i\in I}\left(\frac{P_i(Z\cap X\cap Y)}{P_i(Y\cap X)}\right)\cdot \lim_{i\in I}\left(\frac{P_i(X\cap Y)}{P_i(Y)}\right) + 0 =$, for the same reason again, $= \lim_{i\in I}\left(\frac{P_i(Z\cap X\cap Y)}{P_i(Y\cap X)}\right) \cdot \lim_{i\in I}\left(\frac{P_i(X\cap Y)}{P_i(Y)}\right) + \lim_{i\in I}\left(\frac{P_i(Z\cap Y\cap[W\setminus X])}{P_i(Y)}\right) =$, since all single limits exist again as before, $= \lim_{i\in I}\left(\frac{P_i(Z\cap X\cap Y)}{P_i(Y\cap X)} \cdot \frac{P_i(Y\cap X)}{P_i(Y)} + \frac{P_i(Z\cap Y\cap[W\setminus X])}{P_i(Y)}\right) = \lim_{i\in I}\left(\frac{P_i(Z\cap X\cap Y)}{P_i(Y)} + \frac{P_i(Z\cap Y\cap[W\setminus X])}{P_i(Y)}\right) = \lim_{i\in I}\left(\frac{P_i(Z\cap Y)}{P_i(Y)}\right) = \mathfrak{P}(Z|Y)$, which was to be shown.
All other axioms follow straightforwardly from taking limits of conditional probabilities.

- Now let a Popper function $\mathfrak{P}$ be given on a countable algebra $\mathfrak{A}$.
For $X, Y \in \mathfrak{A}$, let $X \leq Y$ iff $\mathfrak{P}(X|X \cup Y) > 0$. By standard arguments (cf. van Fraassen, 1976, p. 426; Spohn, 1986, p. 72), it follows from the Popper function axioms for $\mathfrak{P}$ that $\leq$ is reflexive (by Axiom 1), transitive (by the Multiplication Axiom), and linear (that is, for all X, Y, $X \leq Y$ or $Y \leq X$, by Axiom 3 and by the fact that if $\mathfrak{P}(X|X \cup Y) = 0$ then $X \cup Y$ cannot be abnormal, that is, it does not hold that $\mathfrak{P}(W \setminus (X \cup Y)|X \cup Y) > 0$, which needs Axiom 5 for its proof). It follows from the properties of abnormal sets that if Y is abnormal then for all $X \in \mathfrak{A}$: $X \leq Y$. Let $X \equiv Y$ iff $X \leq Y$ and $Y \leq X$, and $X < Y$ iff $X \leq Y$ and $Y \not\leq X$; $\equiv$ is an equivalence relation. Since $\mathfrak{A}$ is countable, for each equivalence class $[X]_\equiv$ there is an enumeration, such that $[X]_\equiv = \{A_n^{[X]_\equiv} : n \in \mathbb{N}\}$. Finally, we observe that $[X]_\equiv$ is closed under finite unions: if $\mathfrak{P}(X|X \cup Y), \mathfrak{P}(Y|X \cup Y) > 0$, then $\mathfrak{P}(X|X \cup (X \cup Y)), \mathfrak{P}((X \cup Y)|X \cup (X \cup Y)) > 0$. Now we are in the position to define the family of absolute probability measure that is going to represent $\mathfrak{P}$ as intended: Let $I = \left\{\left\langle [X]_\equiv, A_1^{[X]_\equiv} \cup \ldots \cup A_n^{[X]_\equiv}\right\rangle : X \in \mathfrak{A}, \mathfrak{P}(W \setminus X|X) < 1, n \in \mathbb{N}\right\}$ be the index set of the family. For every $i = \langle [X]_\equiv, A_1^{[X]_\equiv} \cup \ldots \cup A_n^{[X]_\equiv}\rangle \in I$, let P_i be defined by: $P_i(Y) = \mathfrak{P}\left(Y|A_1^{[X]_\equiv} \cup \ldots \cup A_n^{[X]_\equiv}\right)$. Note that finite unions of members of any equivalence class are in $\mathfrak{A}$ again, for $\mathfrak{A}$ is generated by sentences of a language that is closed under finite conjunctions. Furthermore, since the enumeration of $[X]_\equiv$ is total, for every member Y of $[X]_\equiv$ there is an m, such that $Y = A_m^{[X]_\equiv}$. Every such P_i is a finitely additive probability measure as follows from Axiom 3, $\mathfrak{P}(W \setminus X|X) < 1$ and $[X]_\equiv$ being closed under finite unions (which entails that $A_1^{[X]_\equiv} \cup \ldots \cup A_n^{[X]_\equiv}$ is normal for all n). Now order I by means of: $i = \langle [X]_\equiv, A_1^{[X]_\equiv} \cup \ldots \cup A_m^{[X]_\equiv}\rangle \leq j = \langle [Y]_\equiv, A_1^{[Y]_\equiv} \cup \ldots \cup A_n^{[Y]_\equiv}\rangle$ iff (i) $X < Y$ or (ii) $X \equiv Y$ and $A_1^{[X]_\equiv} \cup \ldots \cup A_m^{[X]_\equiv} \supseteq A_1^{[X]_\equiv} \cup \ldots \cup A_n^{[X]_\equiv}$. It is easy to see that this definition is independent of the choice of representatives of equivalence classes, and I is totally ordered by the so-defined $\leq$.

At first we show that if there is an $i \in I$, such that $P_i(X) > 0$, then $\mathfrak{P}(Y|X) = \lim_{i \in I} \left(\frac{P_i(Y \cap X)}{P_i(X)}\right)$; otherwise, $\mathfrak{P}(Y|X) = 1$: For suppose that there is no $i \in I$, such that $P_i(X) > 0$; this is only possible if X is abnormal, that is, $\mathfrak{P}(W \setminus X|X) < 1$, since otherwise $\mathfrak{P}(X|A_1^{[X]_\equiv} \cup \ldots \cup A_n^{[X]_\equiv})$ would be greater than 0 for some n, because $[X]_\equiv$ is closed under finite unions and for some n, $X = A_n^{[X]_\equiv}$. But since X is abnormal, it follows that $\mathfrak{P}(Y|X) = 1$ for any $Y \in \mathfrak{A}$ (see van Fraassen, 1976, p. 419). So we may focus on the case in which there is an $i \in I$, such that $P_i(X) > 0$, in which case X cannot be abnormal by the definition of I and by the fact that if X is abnormal and A is normal, then $\mathfrak{P}(X|A) = 0$ (see van Fraassen, 1976, p. 420). We find: (a) For every $j = \left\langle [Z]_\equiv, A_1^{[Z]_\equiv} \cup \ldots \cup A_n^{[Z]_\equiv} \right\rangle$ with $[Z]_\equiv < [X]_\equiv$ it holds that $P_j(X) = 0$: since then $\mathfrak{P}(Z|Z \cup X) > 0$, $\mathfrak{P}(X|Z \cup X) = 0$, hence $\mathfrak{P}(X|Z) =$, by the Multiplication Axiom, $= \frac{\mathfrak{P}(X \cap Z|Z \cup X)}{\mathfrak{P}(Z|Z \cup X)} = 0$; since that holds for all members Z of $[Z]_\equiv$, it is also true for any $A_1^{[Z]_\equiv} \cup \ldots \cup A_n^{[Z]_\equiv} \in [Z]_\equiv$. (b) There is an m, such that for all $n \geq m$, for all $i = \langle [X]_\equiv, A_1^{[X]_\equiv} \cup \ldots \cup A_n^{[X]_\equiv} \rangle$, $\mathfrak{P}(Y|X) = \frac{P_i(Y \cap X)}{P_i(X)}$: for there is an m, such that $X = A_m^{[X]_\equiv}$, and $[X]_\equiv$ is closed under finite unions, as mentioned before; therefore for all $n \geq m$, $\mathfrak{P}(Y|X) =$, by the Multiplication Axiom again, $= \frac{\mathfrak{P}(Y \cap X|A_1^{[X]_\equiv} \cup \ldots \cup A_n^{[X]_\equiv})}{\mathfrak{P}(X|A_1^{[X]_\equiv} \cup \ldots \cup A_n^{[X]_\equiv})} = \frac{P_i(Y \cap X)}{P_i(X)}$. (c) By (a) and (b) taken together, $\lim_{i \in I} \left(\frac{P_i(Y \cap X)}{P_i(X)}\right)$ exists and $\mathfrak{P}(Y|X) = \lim_{i \in I} \left(\frac{P_i(Y \cap X)}{P_i(X)}\right)$. □

Proof of Theorem 2.3. We only give a sketch. The first part simply follows from the first part of Theorem 2.2. The second part follows directly from the construction in the proof of the second part of Theorem 2.2, with the following addition: for every equivalence class $[X]_\equiv$, define $P_\infty^{[X]_\equiv}(.) = \lim_{n \to \infty} \mathfrak{P}(.|A_1^{[X]_\equiv} \cup \ldots \cup A_n^{[X]_\equiv})$; add indices for all such limit measures to the index set I; for the ordering of the resulting index set, keep the order on I, but for every $[X]_\equiv$ position a unique index k in the ordering, such that k is less than any index of the form $i = \langle [X]_\equiv, A_1^{[X]_\equiv} \cup \ldots \cup A_m^{[X]_\equiv} \rangle$ but greater than any index of the form $j = \langle [Y]_\equiv, A_1^{[Y]_\equiv} \cup \ldots \cup A_m^{[Y]_\equiv} \rangle$ with $Y < X$ (this position is uniquely defined); let $P_k(.) = P_\infty^{[X]_\equiv}(.)$. Then the Limit Assumption is satisfied, and part 2 of Theorem 2.3 follows. □

BIBLIOGRAPHY

Adams, E. W. (1975). *The Logic of Conditionals: An Application of Probability to Deductive Logic*. Synthese Library 86. Dordrecht, The Netherlands: Reidel.

Adams, E. W. (1986). On the logic of high probability. *Journal of Philosophical Logic*, **15**, 255–279.

Arló-Costa, H., & Parikh, R. (2005). Conditional probability and defeasible inference. *Journal of Philosophical Logic*, **34**, 97–119.

Barker, S. (1999). Counterfactuals, probabilistic counterfactuals and causation. *Mind*, **108**, 427–469.

Bennett, J. (2003). *A Philosophical Guide to Conditionals*. Oxford, UK: Clarendon Press.

Brogard, B., & Salerno, J. (2007). Remarks on counterpossibles. *Paper presented at the first annual Synthese conference Between Logic and Intuition: David Lewis and the Future of Formal Methods in Philosophy, October 3rd 2007*.

Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford, UK: Clarendon Press.

Csaszar, A. (1955). Sur la Structure des Espaces de Probabilitè Conditionelle. *Acta Mathematica Hungarica*, **6**, 337–361.

Edgington, D. (1995). On conditionals. *Mind*, **104**, 235–329.

Edgington, D. (1996). Vagueness by degrees. In Keefe, R., and Smith, P., editors. *Vagueness. A Reader*. Cambridge, MA: The MIT Press, pp. 294–316.
Edgington, D. (2004). Counterfactuals and the benefit of hindsight. In Dowe, P., and Noordhof, P., editors. *Cause and Chance: Causation in an Indeterministic World*. London: Routledge, pp. 12–27.
Eells, E., & Skyrms, B., editors. (1994). *Probability and Conditionals. Belief Revision and Rational Decision*. Cambridge, UK: Cambridge University Press.
Elga, A. (2000) Statistical mechanics and the asymmetry of counterfactuals. *Philosophy of Science*, **68**(Suppl.), 313–324.
French, S., & Krause, D. (2003). Quantum vagueness. *Erkenntnis*, **59**, 97–124.
Gundersen, L. B. (2004). Outline of a new semantics for counterfactuals. *Pacific Philosophical Quarterly*, **85**, 1–20.
Hájek, A. (unpublished). Most counterfactuals are false. Unpublished draft.
Halpern, J. Y. (2001). Lexicographic probability, conditional probability, and nonstandard probability. In *Proceedings of the Eighth Conference on Theoretical Aspects of Rationality and Knowledge*. Ithaca, NY: Morgan Kaufmann, pp. 17–30.
Hawthorne, J. (2005). Chance and counterfactuals. *Philosophy and Phenomenological Research*, **70**, 396–405.
Kvart, I. (2001). The counterfactual analysis of cause. *Synthese*, **127**, 389–427.
Lehmann, D., & Magidor, M. (1992). What does a conditional knowledge base entail? *Artificial Intelligence*, **55**, 1–60.
Leitgeb, H. (2004). *Inference on the Low Level. An Investigation into Deduction, Nonmonotonic Reasoning, and the Philosophy of Cognition*. Applied Logic Series. Dordrecht, The Netherlands: Kluwer.
Lewis, D. K. (1973a). Counterfactuals and comparative possibility. *Journal of Philosophical Logic*, **2**, 418–446. Reprinted in Lewis (1986, pp. 3–31).
Lewis, D. K. (1973b). *Counterfactuals*. Oxford, UK: Blackwell.
Lewis, D. K. (1979). Counterfactual dependence and time's arrow. *Noûs*, **13**, 455–476. Reprinted in Lewis (1986, pp. 32–52).
Lewis, D. K. (1980). A subjectivist's guide to objective chance. In Jeffrey, R. C., editors. *Studies in Inductive Logic and Probability*, Vol. 2. Berkeley, CA: University of California Press, pp. 263–293. Reprinted in Lewis (1986, pp. 83–113).
Lewis, D. K. (1986). *Philosophical Papers*, Vol. 2. Oxford, UK: Oxford University Press.
McGee, V. (1985). A counterexample to modus ponens. *The Journal of Philosophy*, **82**, 462–471.
McGee, V. (1994). Learning the impossible. In Eells, E., & Skyrms, B., editors. *Probability and Conditionals. Belief Revision and Rational Decision*. Cambridge, UK: Cambridge University Press, pp. 177–199.
Menzies, P. (2004). Causal models, token causation and processes. *Philosophy of Science*, **71**, 820–832.
Rényi, A. (1955). On a new axiomatic theory of probability. *Acta Mathematica Hungarica*, **6**, 285–333.
Schurz, G. (2001). What is 'normal'? An evolution-theoretic foundation of normic laws and their relation to statistical normality. *Philosophy of Science*, **68**, 476–497.
Schurz, G. (2005). Non-monotonic reasoning from an evolution-theoretic perspective: ontic, logical and cognitive foundations. *Synthese*, **146**, 37–51.
Schurz, G., & Leitgeb, H. (2008). Finitistic and frequentistic approximation of probability measures with or without σ-additivity. *Studia Logica*, **89**, 257–283.
Spohn, W. (1986). The representation of Popper measures. *Topoi*, **5**, 69–74.
van Fraassen, B. C. (1976). Representation of conditional probabilities. *Journal of Philosophical Logic*, **5**, 417–430.
Wallace, D. (2003). Everett and structure. *Studies in the History and Philosophy of Modern Physics*, **34**, 86–105.
Williams, J. R. G. (2008). Chances, counterfactuals, and similarity. *Philosophy and Phenomenological Research*, **77**, 385–420.
Yablo, S. (2009). Comments on Hannes Leitgeb, 'A Probabilistic Semantics for Counterfactuals'. *Commentary at the Conference on Philosophical Logic, Princeton University, May 24th 2009*.