

Diachronic Dutch Books & Evidential Import

J. Dmitri Gallow [†]

ABSTRACT

A handful of well-known arguments (the ‘diachronic Dutch book arguments’) rely upon theorems establishing that, in certain circumstances, you are immune from sure monetary loss (you are not ‘diachronically Dutch bookable’) if and only if you adopt the strategy of conditionalizing (or Jeffrey conditionalizing) on whatever evidence you happen to receive. These theorems require non-trivial assumptions about which evidence you might acquire—in the case of conditionalization, the assumption is that, if you might learn that e , then it is not the case that you might learn something else that is consistent with e . These assumptions may not be relaxed. When they are, not only will non-(Jeffrey) conditionalizers be immune from diachronic Dutch bookability, but (Jeffrey) conditionalizers will themselves be diachronically Dutch bookable. I argue: 1) that there are epistemic situations in which these assumptions are violated; 2) that this reveals a conflict between the premise that susceptibility to sure monetary loss is irrational, on the one hand, and the view that rational belief revision is a function of your prior beliefs and the acquired evidence alone, on the other; and 3) that this inconsistency demonstrates that diachronic Dutch book arguments for (Jeffrey) conditionalization are invalid.

I INTRODUCTION

THE Bayesian belief-revision norm *conditionalization* says that, when experience rationalizes absolute certainty that e , nothing stronger, and nothing else,¹ you should be disposed to adopt a credence function, C_e , which is your prior credence function conditionalized on e .

CONDI

Upon undergoing an experience which rationalizes absolute certainty that e , nothing stronger, and nothing else, you should be disposed to conditionalize on e .²

$$C_e(\cdot) \stackrel{!}{=} C(\cdot | e)$$

Final Draft; forthcoming in *Philosophy and Phenomenological Research*

✉: jdmitrigallow@pitt.edu

[†] I am indebted to Gordon Belot, Rachael Briggs, Irena Cronin, Kenny Easwaran, Branden Fitelson, Alan Hájek, Daniël Hoek, Jim Joyce, Harvey Lederman, Richard Pettigrew, Bernhard Salow, Charles Sebens, Rohan Sud, Jonathan Weisberg, and anonymous reviewers for helpful feedback on this material. Thanks also to audiences at USC, UCSB, and the University of Pittsburgh.

¹ By ‘nothing else’, I mean that your experience has *only* rationalized absolute certainty that e ; it has not additionally, *e.g.*, rationalized becoming more confident in $a \cap e$ than $\neg a \cap e$.

Many Bayesians have thought that experience can rationalize belief revisions other than absolute certainty in some proposition. In that case, CONDI does not apply. JEFFREY (1965), therefore, provides a generalization of that norm, now known eponymously as *Jeffrey conditionalization*. This belief revision norm applies when experience rationalizes shifting your credence in each cell e_i of a partition \mathcal{E} to a new posterior credence ϵ_i .³ We can represent this kind of evidence with a set of ordered pairs $E = \{\langle e_1, \epsilon_1 \rangle, \dots, \langle e_N, \epsilon_N \rangle\}$, and we can call evidence like this a *Jeffrey shift*. Jeffrey conditionalization says that, when you undergo a Jeffrey shift E , you should be disposed to adopt a posterior credence function $C_E(\cdot)$ which is equal to $\sum_i C(\cdot | e_i) \cdot \epsilon_i$.

JCONDI

Upon undergoing an experience which rationalizes credence ϵ_i in e_i , for every e_i in a partition \mathcal{E} , and nothing else, you should be disposed to Jeffrey conditionalize on $\{\langle e_1, \epsilon_1 \rangle, \dots, \langle e_N, \epsilon_N \rangle\}$.

$$C_E(\cdot) \stackrel{!}{=} \sum_i C(\cdot | e_i) \cdot \epsilon_i$$

There is a popular way of justifying CONDI and JCONDI which appeals to the fact that, roughly, failure to follow the dictates of either CONDI or JCONDI will lead you into diachronic practical irrationality. More carefully, when you stand to learn that one of a partition of propositions is true, if you are disposed to revise your degrees-of-belief in any way other than the way that CONDI tells you to, then a clever bookie could concoct a strategy which exposes you to the risk of losing money with no hope of winning any money (*i.e.*, you are ‘diachronically Dutch book-able’).⁴ And, when you stand to undergo any Jeffrey shift on a partition, if you are disposed to revise your degrees-of-belief in any way other than the way that JCONDI tells you to, then you are diachronically Dutch book-able.

Many have objected to these diachronic Dutch book arguments (‘*DDBAS*’) for CONDI and JCONDI on the grounds that they offer a merely *pragmatic*, rather than an *epistemic*, reason to revise your degrees-of-belief in accordance with CONDI or JCONDI, and/or that diachronic pragmatic inconsistency need not be a symptom of any underlying epistemic irrationality.⁵ Put these objections aside. There is a deeper problem with using a *DDBA* to justify an update rule like CONDI or JCONDI. As I will explain in further depth below, in order to concoct a diachronic Dutch book strategy, you must know ahead of

² Two words on notation: 1) I will use ‘ C_e ’ to refer to the credence function which you are disposed to adopt, upon undergoing an experience that rationalizes absolute certainty that e ; and 2) throughout, I’ll place an exclamation mark ‘!’ over an equality to indicate that the equality has normative, and not descriptive, force. That is: the claim is not that $C_e(\cdot)$ will be $C(\cdot | e)$, but rather merely that it *should* be.

³ A partition \mathcal{E} of your credal state is a set of propositions (sets of doxastic possibilities) such that exactly one of the propositions in the set is true at every doxastically possible world.

⁴ Throughout, I’m going to focus on what HÁJEK (2008) calls *semi-Dutch books*—*i.e.*, Dutch books which at best break even and at worst lose money. This is entirely for ease of presentation; Dutch books which lose money in every possibility are available for both CONDI (in TELLER, 1976) and JCONDI (in ARMENDT, 1980 and SKYRMS, 1987).

⁵ E.g., CHRISTENSEN (1991).

time which possible learning experiences your target might undergo—which propositions or Jeffrey shifts they might learn. The DDBA for CONDI assumes that their potential total evidence propositions form a partition; and the DDBA for JCONDI assumes that the potential Jeffrey shifts are all shifts on a given partition. These assumptions may not be relaxed. If we allow some of the potential evidence propositions to overlap,⁶ then conditionalizers will themselves be diachronically Dutch book-able (§2.2). And if we allow that all potential Jeffrey shifts confirm overlapping propositions, Jeffrey conditionalizers will themselves be diachronically Dutch book-able (§2.4).

In general, whether you are diachronically Dutch book-able depends not merely on the learning experience you *actually* undergo, but additionally upon what other learning experiences you *might* have undergone instead. Therefore, on the assumption that there are learning scenarios in which you might learn either of two consistent propositions, the premises of the DDBA entail a contrastivist theory of evidential import, according to which the rational response to a learning experience depends not just upon your prior credences and the total acquired evidence, but additionally upon what other evidence might have been acquired instead. On the other hand, as they are standardly understood, CONDI and JCONDI are not sensitive to which other learning experiences you might have undergone. CONDI and JCONDI are only sensitive to the learning experiences you actually undergo. As standardly understood, these rules entail a noncontrastivist theory of evidential import, according to which the rational response to a learning experience depends upon your prior credences and the acquired evidence alone. We must therefore either abandon DDBAs or else emend our understanding of CONDI and JCONDI.

While the DDBA entails contrastivism about evidential import, many other vindications of CONDI and JCONDI (as well as the standard understandings of those rules themselves) entail non-contrastivism about evidential import. This distinction has been largely papered over in philosophical work at the foundations of Bayesian epistemology. More generally, then, this article serves as a brief on the importance of adjudicating whether evidential import is contrastive, as suggested by the DDBA, or non-contrastive, as suggested by several other justifications of CONDI and JCONDI.

2 DIACHRONIC DUTCH BOOK ARGUMENTS

The diachronic Dutch book arguments for CONDI and JCONDI presuppose that at least one part of your doxastic state—what we can call your *credal state*—is representable as a triple $\langle \mathcal{W}, \mathcal{A}, C \rangle$ of a set of *worlds*, \mathcal{W} , a set of what I will call *propositions* (sets of worlds) $\mathcal{A} \subseteq \mathcal{P}(\mathcal{W})$, and a *credence function*, C , from the members of \mathcal{A} to the unit interval $[0, 1]$. \mathcal{W} contains all the worlds which you take to be possible. \mathcal{A} is known as your *algebra*. It contains all those propositions about which you have credal opinions—those propositions over which your credence function is defined. The credence function C represents how confident you are in each proposition $p \in \mathcal{A}$. If $C(p) = 0$, then you are certain that p is false. If $C(p) = 1$, then you are certain that p is true. If $C(p) = 0.5$, then you have a middling degree of confidence that p .

⁶ Throughout, I will say that the propositions p_1, p_2, \dots, p_N *overlap* just in case $\bigcap_i p_i \neq \emptyset$.

I'll assume throughout that, at any given time, your credal state ought to satisfy the axioms of probability. How should your credal state evolve over time? A standard assumption—for some, a simplifying assumption, for others,⁷ a substantive theoretical commitment—is that evidence comes in the form of propositions about which your learning experience rationalizes absolute certainty. It is also standardly supposed—though perhaps only as a simplifying assumption—that \mathcal{W} remains fixed over time, and that no propositions are ever added to or subtracted from \mathcal{A} . You never recognize any new possibilities, you never start being opinionated about any new propositions, nor do you stop being opinionated about any old propositions. Given these stipulations, CONDI says that, if e is the strongest proposition about which experience rationalizes absolute certainty, then you ought to transition from the prior credal state $\langle \mathcal{W}, \mathcal{A}, C \rangle$ to the posterior credal state $\langle \mathcal{W}, \mathcal{A}, C_e \rangle$, where $C_e(\cdot) = C(\cdot | e)$.

Let's say that, when the sole effect of a learning experience is to rationalize absolute certainty about a proposition, you have acquired *propositional evidence*. Perhaps not all evidence is propositional. JEFFREY (1965) believed that experience could rationalize shifting your credence across the cells of some partition without rationalizing absolute certainty about any particular proposition—what I called earlier a *Jeffrey shift*. JCONDI says that, if you undergo a Jeffrey shift rationalizing posterior credences of ϵ_i in the cells e_i of a partition \mathcal{E} , then you ought to transition from the prior credal state $\langle \mathcal{W}, \mathcal{A}, C \rangle$ to the posterior credal state $\langle \mathcal{W}, \mathcal{A}, C_E \rangle$, where $C_E(\cdot) = \sum_i C(\cdot | e_i) \cdot \epsilon_i$.

In order to construct a DDBA for or against any belief-updating rule, we must specify not merely which proposition is actually learned, but also which other propositions *could* have been learned instead. To show that an update rule is diachronically Dutch book-able in any objectionable sense, it is not enough to show that an agent conforming to that rule *could* be sold a series of bets, the combination of which is guaranteed to lose money, no matter what. Even a perfectly rational agent could have a credence of 0.6 in p —and therefore buy a dollar bet on p for \$0.6⁸—and later acquire evidence the effect of which is to change their credence in p to 0.4. They could at that point buy a dollar bet on $\neg p$ for \$0.6. They will then have purchased a combination of bets which is guaranteed to lose \$0.2 no matter what. This does not show that you ought never change your credence in any proposition from 0.6 to 0.4, no matter what evidence comes in. Even though this agent has purchased a series of bets which is guaranteed to lose money come what may, it needn't be the case that this could have been predicted ahead of time by a bookie who knew no more than the agent. Perhaps the agent might have acquired evidence that raised, rather than lowered, their credence that p . If they had acquired this evidence, then they would not have purchased the second bet, and

⁷ For instance, LEWIS (1996, 1999).

⁸ For the sake of simplicity, I'll be supposing throughout that the agents under consideration have utility functions which are linear in dollars, are not risk averse, that the truth of the proposition being bet upon does not depend upon whether the agent takes the bet, and so on and so forth. I will also assume that a rational agent will always be happy to either purchase or sell a *fair* bet; that is, a bet which has an expected value equal to its price. (This assumption makes the math a bit easier, but is eliminable. We could get by with the weaker assumption that a rational agent will always be happy to sell a bet with negative expected value and purchase a bet with positive expected value. Thanks to an anonymous reviewer for directing my attention to this point.)




FIGURE 1: The learning scenario presupposed by Lewis's diachronic Dutch book strategy. The agent will either become certain that e_1 is true, and therefore that e_2 through e_N are false, or that e_2 is true, and therefore that e_1, e_3, \dots, e_N are false, or that ... or that e_N is true, and therefore that $e_1 \dots e_{N-1}$ are false.

they would not have bought into a sure loss. What a *DDBA* requires is a *strategy* which will in some possibilities lead to the agent losing money and in no possibilities lead to the agent winning any money. And a strategy like this requires a specification of which propositions the agent might end up learning.

I'll call such a specification an *learning scenario*, \mathcal{S} , and I'll model a learning scenario with a set of propositions $\{e_1, e_2, \dots, e_N\}$. In order to be in this learning scenario, it must be epistemically possible that your total evidence will be e_1 , it must be epistemically possible that your total evidence will be e_2, \dots , it must be epistemically possible that your total evidence will be e_N , and it must be epistemically necessary that your total evidence will be exactly one of the $e_i \in \{e_1, e_2, \dots, e_N\}$.⁹ Notice that nothing in the foregoing definition of a learning scenario implies that a learning scenario \mathcal{S} partitions the epistemic possibilities. However, it does imply that, if $\mathcal{S} = \{e_1, e_2, \dots, e_N\}$ is your learning scenario, then $\{\mathbb{T}e_1, \mathbb{T}e_2, \dots, \mathbb{T}e_N\}$ partitions the epistemic possibilities—where $\mathbb{T}\phi$ says that ϕ is your *total* evidence. You are diachronically Dutch-bookable iff there is some strategy a clever bookie (who knows only that your total evidence will be exactly one of e_1, e_2, \dots , and e_N) could adopt which has the following property: for some $e_i \in \mathcal{S}$, if $\mathbb{T}e_i$, then you lose money, and, for no $e_i \in \mathcal{S}$ will you win money if $\mathbb{T}e_i$. If you are diachronically Dutch bookable, then a clever bookie who knew no more than you could be assured a chance of winning money off you without any chance of losing money to you.

2.1 THE *DDBA* FOR *CONDI*

LEWIS (1999)'s *DDBA* for conditionalization, which first appeared in print in TELLER

⁹ I say that it must be epistemically possible that your total evidence be e_i . What sort of modality is in play here? Is the claim that you don't know that you won't acquire total evidence e_i ? Is it that there's some $w \in \mathcal{W}$ at which you acquire total evidence e_i ? Is it that you assign non-zero credence to the proposition that you acquire total evidence e_i ? The proponent of the *DDBA* for *CONDI* ought to say something about this, though little has actually been said. I suspect that the best thing to say is that ϕ is necessary (in the relevant sense) iff ϕ is true at all possibilities in \mathcal{W} , but I don't believe that anything I have to say here hinges upon this question. To try to remain neutral on these questions, I will continue to use 'doxastic possibilities' for those possibilities in \mathcal{W} , and 'epistemic possibilities' for those possibilities in terms of which a learning scenario is defined.

(1973, 1976), presupposes that an agent finds themselves in a learning scenario \mathcal{S} which forms a partition. (See figure 1.) In a learning scenario like this, if the agent conforms to a belief-revision norm¹⁰ which tells them to set their posterior credence that p (for any proposition p) to $C_{e_i}(p) > C(p | e_i)$, for any of the $e_i \in \mathcal{S}$, then they could be sold bet 1 prior to undergoing the learning experience. (The dollar amount on the left is the agent's *net* gain from the bet if the proposition to the right turns out to be true.)

Bet 1	
$\$C(p e_i) - 1$	if $p \cap e_i$
$\$C(p e_i)$	if $\neg p \cap e_i$
$\$0$	if $\neg e_i$

Then, if the agent doesn't learn e_i , they will learn one of the other propositions in \mathcal{S} , which entails that $\neg e_i$, so they will break even on bet 1. If, on the other hand, they learn that e_i , then they can be sold bet 2.

Bet 2	
$\$1 - C_{e_i}(p)$	if p
$\$ - C_{e_i}(p)$	if $\neg p$

Given that e_i is true, the combination of bets 1 and 2 is guaranteed to lose the agent $\$C(p | e_i) - C_{e_i}(p)$. The agent has therefore opened themselves up to the risk of losing money without any hope of winning money.

If, on the other hand, the agent conforms to a belief-revision norm which tells them to set their posterior credence that p to $C_{e_i}(p) < C(p | e_i)$ for any proposition p and any $e_i \in \mathcal{S}$, then they will sell bet 1 to you prior to learning anything. If they don't learn that e_i , then they will break even with bet 1. If they learn that e_i , then they will sell bet 2 to you. At that point, the agent will have bought into a guaranteed loss of $\$C_{e_i}(p) - C(p | e_i)$. Their belief-revision rule will then have opened them up to the possibility of losing money without any chance of winning money.

So, if an agent in a learning scenario \mathcal{S} which forms a partition conforms to a belief revision norm which tells them to set their posterior credence in any proposition p to anything other than $C(p | e_i)$ when they learn that e_i , for any of the $e_i \in \mathcal{S}$, then they are diachronically Dutch book-able. SKYRMS (1987) shows that the converse is true as well; an agent who conforms to CONDI is *not* diachronically Dutch book-able in this kind of learning scenario. The DDBA for CONDI proceeds as follows.

- 1) Any belief revision norm other than CONDI is diachronically Dutch book-able in a learning scenario which forms a partition.
- 2) CONDI is not diachronically Dutch book-able in a learning scenario which forms a partition.

¹⁰ I will say that an agent conforms to a belief-revision norm iff they are disposed to adopt some particular posterior doxastic state after learning that e is true (for any e), or after undergoing any Jeffrey shift. If an agent is disposed, e.g., to pick a doxastic state at random, willy-nilly, upon learning that e (for some e), or undergoing some Jeffrey shift, they will not be diachronically Dutch book-able.

- 3) Ideally epistemically rational agents conform to belief-revision norms.
- 4) Ideally epistemically rational agents are not diachronically Dutch book-able.
- 5) So, in learning scenarios which form a partition, ideally epistemically rational agents conform to CONDI.

This argument is valid. Its conclusion, (5), however, is not that CONDI is always rationally ideal. *That* conclusion, (6), does not follow from (1)–(4).

- 6) So, ideally epistemically rational agents always conform to CONDI.

If we had an extra premise to the effect that ideally epistemically rational agents only acquire propositional evidence in learning scenarios which form a partition, then (6) would follow. However, if either i) it is possible to acquire propositional evidence outside of any learning scenario, or ii) it is possible to acquire propositional evidence in a learning scenario which is not a partition, then the DDBA for CONDI would not establish that ideally rational agents update conform to CONDI in those cases.

How could an agent acquire evidence outside of a learning scenario? One possibility is that, while their doxastically possible worlds specify such things as whether a rolled die lands odd or even, they do not specify what the agent themselves learn. For instance, if we think of doxastic possibilities as maximally consistent sets of propositions expressible in some underlying language, and the language lacks the resources to express propositions like “The agent learns that e ”, then doxastic possibilities could fail to specify what the agent learns. Assuming that doxastic possibility is epistemic possibility (see fn 9), the definition of a learning scenario provided above entails that the agent whose credal state is represented with such possibilities could never be in a learning scenario. Of course, we might have views about which kinds of entities can serve as doxastically possible worlds which guarantee that an agent will always be in some learning scenario or other—perhaps, for instance, a doxastic possibility must specify a truth-value for every proposition. I don’t wish to take a stand on such questions here; I simply wish to note that there are ways of understanding epistemic possibility which allow agents to acquire evidence outside of any learning scenario. If this is so, then the DDBA does not establish that such an agent should conform to CONDI.

In presenting their DDBA, did LEWIS and TELLER intend to establish (6), or rather the weaker (5)? I believe it is clear enough that *Lewis* intended the DDBA to establish that “a severely idealized, superhuman subject who runs no risk of mistaking his evidence”¹¹ should always conditionalize on their total evidence. As I will discuss in more depth below, LEWIS believed that this superhuman subject represented not just any idealization, but a *rational* idealization—that is, LEWIS held that an ideally rational agent would always find themselves in a learning scenario which forms a partition; and consequently, by the DDBA, they would always conditionalize on their total evidence. In any case, it is important to stress that, whether he thought the DDBA *established* (6), LEWIS certainly *accepted* (6).¹²

¹¹ LEWIS (1999, p. 404)

¹² “Conditionalising on one’s total evidence is a rational way to learn from experience.” (LEWIS, 1981, p. 6)




FIGURE 2: A learning scenario $\mathcal{S} = \{e_1, e_2\}$ in which the potential evidence propositions e_1 and e_2 overlap.

Whether **TELLER** intended the DDBA to establish the stronger (6), or rather the weaker (5), is a difficult exegetical question which is not, on my reading, settled by the text. However, it is worth stressing that, whether or not **TELLER** thought that the DDBA *establishes* that you should always conditionalize on propositional evidence—no matter whether or not your learning scenario forms a partition—the article does *accept* this conclusion. For **TELLER** (1976) contains at least three arguments for updating your degrees of belief by conditionalization, all of which are presented as arguments in favor of the very same update rule; and only one of these arguments—the DDBA—relies upon or contains any mention of the agent’s learning scenario. Therefore, those other arguments, if successful, establish the stronger conclusion (6). This is made explicit in §1.7 of **TELLER** (1976), when we receive a precise characterization of “the circumstances under which it seems that reasonable change of belief is described by conditionalization”.¹³ As the curious reader may verify for themselves, this characterization simply elucidates what it is to acquire what I’ve called *propositional evidence*, and it makes no mention whatsoever of the agent’s learning scenario. (I am presenting this evidence that **LEWIS** and **TELLER** accepted (6) because some have reacted to the arguments I am about to present by insisting that **LEWIS** and **TELLER** in fact, contrary to all appearances, *rejected* (6). This contention is not borne out by a careful reading of the relevant texts.)

2.2 THE DDBA AGAINST CONDI

If it is possible to acquire evidence in a learning scenario, some of whose members overlap in a particular way, then, not only does it not follow that any rule other than CONDI is diachronically Dutch book-able—it follows that CONDI itself is diachronically Dutch book-able.

For illustration, consider an agent who finds themselves in the learning scenario shown in figure 2. That agent might learn that e_1 (and nothing stronger), they might learn that e_2 (and nothing stronger), it is guaranteed that they will learn exactly one of these, and e_1 and e_2 overlap. If they conditionalize on e_1 , then their credence that $e_1 \cap e_2$ will go up. So long as $0 < C(e_1) < 1$, $C(e_1 \cap e_2 | e_1) > C(e_1 \cap e_2)$. Similarly, if they conditionalize on e_2 , their credence that $e_1 \cap e_2$ will go up. So long as $0 <$

¹³ **TELLER** (1976, p. 225–226).

$C(e_2) < 1$, $C(e_1 \cap e_2 | e_2) > C(e_1 \cap e_2)$. So, no matter what evidence they acquire, if they conditionalize on that evidence, their credence that $e_1 \cap e_2$ will go up. This means that, prior to the learning experience, they will accept bet 3.

Bet 3	
$\$C(e_1 \cap e_2) - 1$	if $e_1 \cap e_2$
$\$C(e_1 \cap e_2)$	if $\neg e_1 \cup \neg e_2$

And, after the learning experience, if they learn that e_i , they will accept bet 4.

Bet 4	
$\$1 - C(e_1 \cap e_2 e_i)$	if $e_1 \cap e_2$
$\$-C(e_1 \cap e_2 e_i)$	if $\neg e_1 \cup \neg e_2$

The agent's net gain from bets 3 and 4 will be $\$C(e_1 \cap e_2) - C(e_1 \cap e_2 | e_i)$. Since $C(e_1 \cap e_2 | e_i) > C(e_1 \cap e_2)$, for $i = 1, 2$, this means that an agent who conditionalizes on their evidence is diachronically Dutch book-able in this learning scenario. This diachronic Dutch book strategy may be generalized to some but not all non-partitioning learning scenarios; I'll leave the details in appendix A.

Is this really a diachronic Dutch book? Here is one reason for doubt: a crucial stipulation in any Dutch book argument is that the imagined Dutch bookie does not have, and does not receive, any more evidence than their target. (The victim of insider trading is not thereby guilty of epistemic irrationality.) But note that, though our imagined agent will either acquire the total evidence e_1 or the total evidence e_2 , they will definitely *not* get the evidence that e_1 is their total evidence, $\mathbb{T}e_1$, nor the evidence that e_2 is their total evidence, $\mathbb{T}e_2$. For, as we saw above, it follows from the stipulation that their learning scenario is $\mathcal{S} = \{e_1, e_2\}$ that $\{\mathbb{T}e_1, \mathbb{T}e_2\}$ is a partition. If they were assured to learn whether $\mathbb{T}e_1$ or $\mathbb{T}e_2$, then their learning scenario would form a partition after all. So it is important that our imagined Dutch bookie not come to learn either $\mathbb{T}e_1$ or $\mathbb{T}e_2$; it is important that they only come to learn either e_1 or e_2 . But, we may worry, if the Dutch bookie doesn't come to learn whether $\mathbb{T}e_1$ or $\mathbb{T}e_2$, then how will they carry out their Dutch book strategy?¹⁴

The answer is: the Dutch bookie will carry out their strategy in precisely the same way that the conditionalizer updates their credences. The conditionalizer is disposed, upon acquiring the total evidence e_i , to adopt a new credence function $C(\cdot | e_i)$. That is: in the condition in which $\mathbb{T}e_i$ is true (the 'triggering condition'), they are disposed to change their credences to $C(\cdot | e_i)$. The conditionalizer may have this disposition without being in a position to acquire the evidence $\mathbb{T}e_i$, just as they could be disposed to have a headache after consuming MSG without thereby being able to learn whether they have consumed MSG. In general, being disposed to react a certain way in a triggering condition does not require being in a position to learn that the triggering condition obtains. And, just as the conditionalizer is disposed to adopt the credence function $C(\cdot | e_i)$ in the triggering condition $\mathbb{T}e_i$, the Dutch bookie is disposed to

¹⁴ Thanks to an anonymous reviewer for calling my attention to this point.

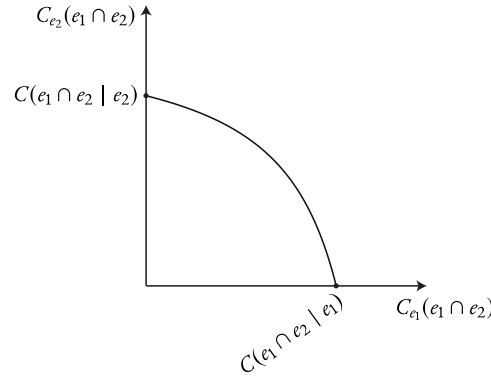


FIGURE 3: The horizontal axis is the credence in $e_1 \cap e_2$ the agent is disposed to adopt if they learn e_1 , and the vertical axis is the credence in $e_1 \cap e_2$ they are disposed to adopt if they learn e_2 . Thus, each point in the plane corresponds to a different update strategy for the proposition $e_1 \cap e_2$. None of the update strategies lying on the curve are diachronically Dutch book-able in a learning scenario $\mathcal{S} = \{e_1, e_2\}$, where e_1 and e_2 overlap. (This set will not include the conditionalization strategy $\langle C(e_2 | e_1), C(e_1 | e_2) \rangle$ unless $C(e_1 \cap e_2)$ is either 0 or 1.)

offer bet 4 in the triggering condition $\mathbb{T}e_i$. Such a disposition does not require the bookie to have any evidence which the conditionalizer lacks, nor does it require them to have dispositions any more finely-attuned than those of the conditionalizer. So we are treating the conditionalizer and the bookie even-handedly, and this is a genuine Dutch book strategy.

In conversation, some have wondered whether learning scenarios involving overlap are ones in which *any* update rule is diachronically Dutch book-able. In fact, there are continuum-many update strategies which are immune from diachronic Dutch book-ability in learning scenarios like these. For the special case in which $\mathcal{S} = \{e_1, e_2\}$, where e_1 and e_2 overlap and their intersection has non-zero credence, any of the strategies for updating your credence that $e_1 \cap e_2$ lying on the curve shown in figure 3 are immune from diachronic Dutch book-ability. Another particularly interesting example is shown in figure 4. There, p, e_1 , and e_2 are pairwise independent in the agent’s prior credence function C ; nevertheless, if that agent decides to revise their views about p in response to learning which of $\{e_1, e_2\}$ is true (by choosing any of the update strategies in the gray area in figure 4), such revisions will not render them diachronically Dutch-bookable. (See appendix B for the details.)

HILD (1998b) noted the existence of diachronic Dutch book strategies in learning scenarios like these, and argued that these learning scenarios point to an inconsistency between CONDI and VAN FRAASSEN (1984, 1995)’s principle of REFLECTION, according to which your credence that p ought to be equal to your expectation of your future credence that p .¹⁵ To illustrate: imagine that a conditionalizer who obeys REFLECTION finds themselves in the learning scenario shown in figure 2. Suppose that their prior credence that $\neg e_2$ is $1/3$, their prior credence that $e_1 \cap e_2$ is $1/3$, and their prior credence

¹⁵ More carefully, if C_t is your time t credence function, and $C_{t'}(p) = x$ is the proposition that your time t' credence that p is x , then the principle of reflection is that, for every $t' > t$, and every proposition p , $C_t(p | C_{t'}(p) = x)$ ought to be x , and therefore, $C_t(p)$ ought to be $\sum_x x \cdot C_t(C_{t'}(p) = x)$.




FIGURE 4: The horizontal axis is the credence in p that the agent is disposed to adopt if they learn e_1 , and the vertical axis is the credence in p that they are disposed to adopt if they learn e_2 . Thus, each point in the plane corresponds to a different update strategy for the proposition p . In a learning scenario in which the agent will learn either that e_1 or that e_2 , where e_1 and e_2 overlap, none of the update strategies $\langle C_{e_1}(p), C_{e_2}(p) \rangle$ in the gray area are diachronically Dutch book-able. (To generate the region shown above, I utilized a probability function within which p , e_1 , and e_2 are pairwise independent, $C(p) = 1/2$, and $C(e_1) = C(e_2) = 2/3$.)

that $\neg e_1$ is $1/3$. We may dramatize their reasoning as follows: *whatever I learn, my credence that $e_1 \cap e_2$ will rise to $1/2$. So, I should have credence $1/2$ in $e_1 \cap e_2$ now (and credence $1/4$ in both $\neg e_2$ and $\neg e_1$). Let it be so. But now, no matter what I learn, if I conditionalize upon it, then my credence that $e_1 \cap e_2$ will rise to $2/3$. So I should have credence $2/3$ in $e_1 \cap e_2$ now (and credence $1/6$ in both $\neg e_2$ and $\neg e_1$). Let it be so. But now, no matter what I learn, if I conditionalize upon it, then my credence that $e_1 \cap e_2$ will rise to $4/5$. So I should have credence $4/5$ in $e_1 \cap e_2$ now (and credence $1/10$ in both $\neg e_2$ and $\neg e_1$). Let it be so. But now...* This reasoning iterates indefinitely, for any credence that $e_1 \cap e_2$ between 0 and 1. Nevertheless, there is a way to abide by both CONDI and REFLECTION when you learn that you will either learn e_1 or e_2 : you may either become absolutely certain of *both* e_1 and e_2 , or become absolutely certain that *neither* e_1 nor e_2 . This will steer you clear of diachronic Dutch book-ability, but it could hardly be counted as epistemically rational. HILD advises rejecting CONDI in these kinds of learning scenarios, opting instead for a rule which says that, when your total evidence is e_i , you should conditionalize on $\mathbb{T}e_i$ (even if you didn't learn $\mathbb{T}e_i$). SCHOENFIELD (forthcoming) also endorses this update rule, and provides an expected accuracy-maximization argument for it.¹⁶

¹⁶ As an aside, *pace* HILD and SCHOENFIELD, this rule is incompatible with learning scenarios like the one shown in figure 2. For, in that learning scenario, $\mathbb{T}e$ is a stronger proposition than e is (you learn

WILLIAMSON (2000) argues at length that learning scenarios like the one shown in figure 2 are possible in cases where your perceptual experience is inexact. In §10.6, he considers a diachronic Dutch book argument like the ones sketched above as an objection to this view. WILLIAMSON is unmoved by the objection, as he follows CHRISTENSEN (1991) in thinking that diachronic pragmatic irrationality is consistent with ideal epistemic rationality. BRONFMAN (2014) considers a variant of the three-prisoners problem (equivalent to the famous Monty Hall problem from GARDNER) which has the same structure as the learning scenario from figure 2, and argues that, in such a learning scenario, we should violate CONDI.

GALLOW (2014) argues that learning scenarios like these arise in cases involving theory-dependent evidence. In such cases, you could be in a position in which, if the background theory t is true, then experience will deliver either the evidence e or the evidence $\neg e$. If, however, the background theory t is false, then experience will not deliver any evidence at all. In these cases, experience will either rationalize absolute certainty that $t \rightarrow e$ (and nothing stronger) or it will rationalize absolute certainty that $t \rightarrow \neg e$ (and nothing stronger). And these propositions overlap.

These kinds of learning scenarios also plausibly arise when we discover the outcome of so-called ‘natural experiments’. We learn the outcome of an experiment, but since we didn’t ourselves design the experiment, it was an open possibility for us that we would instead have learned nothing at all (which is consistent with the outcome of the natural experiment).

In all of these examples, it is crucial that the overlapping propositions are the *strongest* propositions learned. Below, I’ll consider some arguments to the effect that, once you look at the *strongest* proposition learned, every learning scenario forms a partition. Here, I just want to make it seem plausible that we can learn things in learning scenarios some of whose members overlap.

In the following section, I will review the DDBA for JCONDI. The impatient reader should feel free to skip ahead to §3.

2.3 THE DDBA FOR JCONDI

SKYRMS (1987) presents a diachronic Dutch book argument for JCONDI. For this argument, he envisages a learning scenario like this: The agent begins with a prior credence function C . Then, they receive some Jeffrey shift, E_i , on the partition \mathcal{E} . Importantly, they might receive *any* Jeffrey shift on this partition. So their learning scenario at this stage consists of all possible Jeffrey shifts on the partition \mathcal{E} . In response to the Jeffrey shift E_i , they update their credence function to C_i . Then, they are told which of the $e_j \in \mathcal{E}$ is true; and they update their credence function to C_{ij} . (See figure 5.) The DDBA for CONDI establishes that if $C_{ij}(p) \neq C_i(p | e_j)$, for any $p \in \mathcal{A}$ and any $e_j \in \mathcal{E}$, then the agent is diachronically Dutch book-able. So assume that $C_{ij}(p) = C_i(p | e_j)$, for every $p \in \mathcal{A}$ and every $e_j \in \mathcal{E}$. Now, if there is any potential Jeffrey shift E_i , any

the latter but not the former). As we are using the term ‘evidence’, what it is for e to be your total evidence is for e to be the strongest propositions about which experience has rationalized absolute certainty. However, if $\mathbb{T}e$ is stronger than e , and if it is rational to conditionalize on $\mathbb{T}e$, then e is not the strongest proposition about which experience has rationalized absolute certainty.




FIGURE 5: The learning scenario presupposed by SKYRMS (1987)’s DDBA for JCONDI. First, the agent will undergo some Jeffrey shift on the partition $\mathcal{E} = \{e_1, e_2, \dots, e_N\}$. Then, they will learn which member of \mathcal{E} is true.

$e_j \in \mathcal{E}$, and any $p \in \mathcal{A}$ such that $C_i(p | e_j) > C(p | e_j)$, then the agent may be sold bet 5 prior to undergoing the Jeffrey shift.

Bet 5	
$\$C(p e_j) - 1$	if $p \cap e_j$
$\$C(p e_j)$	if $\neg p \cap e_j$
$\$0$	if $\neg e_j$

If they don’t learn that e_j , then the agent will break even on bet 5. If, however, they learn that e_j , then they may be sold bet 6.

Bet 6	
$\$1 - C_{ij}(p)$	if p
$\$-C_{ij}(p)$	if $\neg p$

The combination of bets 5 and 6 is guaranteed to lose the agent $\$C(p | e_j) - C_{ij}(p)$. (A symmetric argument shows that the agent is diachronically Dutch book-able in this learning scenario if they adopt a strategy where $C_i(p | e_j) > C(p | e_j)$ for any $p \in \mathcal{A}$, any Jeffrey shift E_j , and any $e_j \in \mathcal{E}$.)

SKYRMS (1987) also shows the converse: that JCONDI is not diachronically Dutch book-able in this two-stage learning scenario. SKYRMS’s DDBA for JCONDI proceeds as follows:

- 7) Any belief revision norm other than JCONDI is diachronically Dutch book-able in this two-stage learning scenario.
- 8) JCONDI is not diachronically Dutch book-able in a this learning scenario.
- 9) Ideally epistemically rational agents conform to belief-revision norms.

- io) Ideally epistemically rational agents are not diachronically Dutch book-able.
- ii) So, in this two-stage learning scenario, ideally epistemically rational agents conform to JCONDI.

As with the DDBA for CONDI, without extra premises, this argument does not establish JCONDI. Here, we would not only need an extra premise to the effect that every Jeffrey shift is acquired in a learning scenario consisting of all Jeffrey shifts all on the same partition; we would also need the extra premise that, if an update rule is diachronically Dutch book-able in this two-stage learning scenario, then an ideally rational agent would not conform to that update rule when they were undergoing only the first stage of that learning scenario.¹⁷

2.4 THE DDBA AGAINST JCONDI

In contrast to the complicated two-stage learning scenario in the DDBAs *for* JCONDI, we can easily provide a DDBA *against* JCONDI by looking to a learning scenario in which all potential Jeffrey shifts confirm overlapping propositions. For instance, suppose that an agent will either undergo a Jeffrey shift raising their credence that e_1 or a Jeffrey shift raising their credence that e_2 , where e_1 and e_2 overlap.¹⁸

In that case, if the agent conforms to JCONDI, then their credence that $e_1 \cap e_2$ will go up no matter which Jeffrey shift they undergo. Then, the Jeffrey conditionalizer may be sold bet 3 (reproduced below) prior to undergoing the learning experience.

Bet 3	
$\$C(e_1 \cap e_2) - 1$	if $e_1 \cap e_2$
$\$C(e_1 \cap e_2)$	if $\neg e_1 \cup \neg e_2$

And, after they undergo a Jeffrey shift raising their credence that e_i , they may be sold bet 4 (reproduced below).

Bet 4	
$\$1 - C_i(e_1 \cap e_2)$	if $e_1 \cap e_2$
$\$ - C_i(e_1 \cap e_2)$	if $\neg e_1 \cup \neg e_2$

Just as before, the combination of bets 3 and 4 will lose the Jeffrey conditionalizer $\$C(e_1 \cap e_2) - C_i(e_1 \cap e_2)$ come what may. So JCONDI is diachronically Dutch book-able in this learning scenario.¹⁹

¹⁷ ARMENDT (1980) presents a DDBA for JCONDI which I will not be discussing for the sake of brevity. ARMENDT's DDBA does without SKYRMS's two-stage learning scenario, but must suppose that the agent has credences in propositions describing the Jeffrey shift they have undergone.

¹⁸ When I say that the agent will undergo a Jeffrey shift 'raising their credence that e ', I mean that the shift will be representable with $\{\langle e, \epsilon \rangle, \langle \neg e, 1 - \epsilon \rangle\}$, with $\epsilon > C(e)$.

¹⁹ A generalization of this Dutch book strategy precisely analogous to the one provided in A is available.

3 DIACHRONIC DUTCH BOOKS & EVIDENTIAL IMPORT

In this section, I'm going to suggest that there's a tension between a norm of rationality which advises us to avoid diachronic Dutch books and a norm of rationality, like CONDI or JCONDI, according to which the rational response to a learning experience is determined entirely by your prior credal state and the evidence acquired in the learning experience. I will suggest that this tension shows us that DDBAS cannot be used to argue for update rules like CONDI and JCONDI, since if you accept these update rules, then you must reject one of the argument's premises—*viz.*, the premise that ideally epistemically rational agents are not diachronically Dutch book-able.

3.1 THE PROBLEM FOR THE DDBA FOR CONDI

Suppose that you are a participant in a psychologist's experiment. You are waiting in line to go into a room that contains a certain object which is either spherical or cubical, and is either green or blue; it is *not*, however, you have been assured, both blue and cubical. The room will either be pitch black, and you will be allowed to feel the object to ascertain its shape, or the room will be well lit, and you will be able to ascertain its color, though not enough of the object will be exposed for you to ascertain its shape. So you will either learn something about the object's shape, or something about its color, but not both. Every participant in the experiment will learn precisely the same thing when they enter the room. You are told, and are absolutely certain of, all of this ahead of time. Before you enter the room, a close friend of yours who was ahead of you in line and who you are absolutely confident is speaking truthfully says one of the following three sentences to you:

- α) You will either learn that it is green or that it is blue
- β) You will either learn that it is green or that it is spherical
- γ) You will either learn that it is cubical or that it is spherical.

Call the version of you that hears sentence α ' A ', call the version of you that hears sentence β ' B ', and call the version of you that hears sentence γ ' T '. Suppose further, just to cancel any implicatures that α , β , or γ might carry, that your friend is in the habit of uttering true disjunctions of this form, even when they know which of the disjuncts is true, that there's no interesting correlation between the order the disjuncts come in and which is known, and so on and so forth. That is to say: the only relevant information that your friend's utterance of α , β , or γ carries is that those disjunctions are true.

Since you are certain that it is not the case that the object is both cubical and blue, $\{c, g \cap s, b\}$ partitions your credal state. The information that your friend gives you, with α , β , or γ , rules out no cell of this partition. Let us suppose for the time being—though I will revisit this assumption later on—that learning α , β , or γ does not call for a change in the credences you assign to the cells of this partition, so that each of A , B , and T have the same prior credences in c , $g \cap s$, and b .




FIGURE 6: A 's, B 's, and Γ 's learning scenarios. A will either learn that b or that g . B will either learn that g or that s . Γ will either learn that s or that c .

Since A is certain that they will either learn that g or they will learn that b , their learning scenario is $\mathcal{S}_A = \{g, b\}$. Since B is certain that they will either learn that g or they will learn that s , their learning scenario is $\mathcal{S}_B = \{g, s\}$. Since Γ is certain that they will either learn that c or they will learn that s , their learning scenario is $\mathcal{S}_\Gamma = \{c, s\}$. These learning scenarios are shown in figure 6. This case brings out a tension between the following two theses.

Rational Updating is Determined by Priors and Evidence (PE) The rational way to update your credence that p when you acquire the evidence e is entirely determined by your prior credence that $p \cap e$ and your prior credence that e .

Rational Updating is not Diachronically Dutch Book-able (DB) If you rationally update your prior credences, then you will not be diachronically Dutch book-able.

The tension is that, if triads of learning scenarios like A 's, B 's and Γ 's are possible, then **PE** and **DB** are inconsistent. Suppose, for *reductio*, that learning scenarios like these are possible and that both **PE** and **DB** are true. \mathcal{S}_A is a partition; so, by **DB**, rationality demands that A 's posterior credence that $g \cap s$, if they learn g , be A 's prior credence that $g \cap s$ conditionalized on g . But—we are supposing for now— B has the same prior credences in $g \cap s$ and g that A does; so by **PE**, B 's posterior credence that $g \cap s$ after learning g must also be the prior credence that $g \cap s$ conditionalized on g . And again, \mathcal{S}_Γ is a partition, so by **DB**, rationality demands that Γ 's posterior credence that $g \cap s$, if they learn s , be Γ 's prior credence that $g \cap s$ conditionalized on s . But—we are supposing for now— B has the same prior credences in $g \cap s$ and s that Γ does; so by **PE**, B 's posterior credence that $g \cap s$ after learning s must also be their prior credence that $g \cap s$ conditionalized on s . Thus, rationality demands that B update their credence that $g \cap s$ by conditionalizing on whatever evidence they receive. But because \mathcal{S}_B is the overlapping $\{g, s\}$, if they do that, then they are diachronically

Dutch book-able (this is what we learned in §2.2). So rationality can't demand it of them, by **DB**. Contradiction.

CONDI presupposes **PE**. The **DDBA** for CONDI presupposes **DB**—it is just premise (4) of the argument outlined in §2.1. So, if triads of learning scenarios like A 's, B 's, and Γ 's are possible, then if you accept CONDI, then you must reject premise (4) of the **DDBA** for CONDI. This pushes against even the following, very minimal way of thinking about the **DDBA** for CONDI: while the **DDBA** does not establish that you ought to update by conditionalization in *all* circumstances, it does show that you should do when your learning scenario forms a partition; and surely this speaks in favor of CONDI. What we've just seen is that, if triads of learning scenarios like A 's, B 's, and Γ 's are possible, then even this way of thinking about the **DDBA** for CONDI cannot be correct. Far from speaking in CONDI's favor, the **DDBA** speaks decisively *against* CONDI, in the sense that, if you accept the premises of the **DDBA**, then you must reject CONDI. If triads of learning scenarios like A 's, B 's, and Γ 's are possible, then the **DDBA** for CONDI is not merely limited in scope; it is invalid.

Above, I made the assumption that your credences in $\{c, g \cap s, b\}$ should not change upon hearing α , β , or γ . However, this assumption is not necessary for generating the contradiction between **PE** and **DB**. The only assumption we need is that there is a triple $\langle C_A, C_B, C_\Gamma \rangle$ of prior credal states such that $C_{A,\alpha}(\phi) = C_{B,\beta}(\phi) = C_{\Gamma,\gamma}(\phi)$, for $\phi \in \{c, g \cap s, b\}$ —where ' $C_{A,\alpha}$ ' is a rational credence function to adopt upon hearing α , given that one started with the credence function C_A , and likewise for ' $C_{B,\beta}$ ' and ' $C_{\Gamma,\gamma}$ '. From there, the argument goes through as before. Since A 's learning scenario is the partition $\{g, b\}$, **DB** entails that they must conditionalize on g if they learn it. Since B 's prior credence that $g \cap s$ and g is the same as A 's, **PE** entails that B must update their credence that $g \cap s$ by conditionalizing on g if they learn g . Since Γ 's learning scenario is the partition $\{c, s\}$, **DB** entails that they must conditionalize on s if they learn it. Since B 's prior credence that $g \cap s$ and s is the same as Γ 's, **PE** entails that B must update their credence that $g \cap s$ by conditionalizing on s if they learn s . But if B updates their credence that $g \cap s$ by conditionalizing on whatever they learn, they will be diachronically Dutch book-able, because their learning scenario is the overlapping $\{g, s\}$. Contradiction.

We thus face three options: firstly, we could either emend our understanding of CONDI, insisting that it only applies in learning scenarios which form a partition. This position is explicitly endorsed by HILD (1998b,a), BRONFMAN (2014), and SCHOENFIELD (forthcoming). Secondly, we could reject the claim that rational updating is not diachronically Dutch book-able. This position is explicitly endorsed by CHRISTENSEN (1991). Thirdly, we could deny that triads of learning scenarios like A 's, B 's, and Γ 's are possible. As the first two options are consistent with my conclusion, I will spend the rest of this section considering the final option. I will try to persuade you to reject this final option, and to accept that it is not a failure of ideal rationality to think, *e.g.*, that you might learn something, yet you might also learn nothing. However, even if you are unpersuaded, I think you should accept that, *if* triads of learning scenarios like these are not possible, then it is possible to acquire evidence outside of any learning scenario. As noted in §2.1, if it is possible to acquire evidence outside of any learning

scenario, then no DDBA is capable of vindicating CONDI. So, while I will try to persuade you that accepting CONDI is inconsistent with accepting the premises of a DDBA, even if you remain unpersuaded, you should nevertheless accept that many learning experiences fall outside of the purview of the DDBA.

If we are to deny that triads of learning scenarios like those shown in figure 6 are possible, we could do so on at least one of two grounds: firstly, we could say that the nature of evidence guarantees that learning scenarios form a partition. Secondly, we could say that it is in the nature of a learning scenario that it must form a partition. A good example of the first kind of response comes from LEWIS (1999), who appears to implicitly defend the assumption that a learning scenario will always form a partition when he describes the propositions that you might learn in any learning experience as being “mutually exclusive and jointly exhaustive propositions that specify, in full detail, all the alternative courses of experience you might undergo”.²⁰ If the evidence you acquire consists of a proposition specifying in full detail the course of experience you have undergone (along with something like a *that’s all* clause—more on that in a bit), then the potential evidence will form a partition. Nevertheless, the assumption that your evidence consists of a proposition specifying in full detail the course of experience you have undergone is rather implausible. You look up at the night sky. Perhaps you learn that Orion is visible; perhaps you learn that the big dipper is not. But, *prima facie*, you do not learn exactly how many stars are in your visual field; nor do you learn the ratio of the distance between Orion’s hands and the distance between his feet.²¹

Even Lewis’s incredibly rich conception of evidence fails to guarantee that the potential evidence forms a partition without the additional assumption that the evidence includes something like a *that’s all* proviso to the effect that experience does not consist of anything other than the features specified by the evidence. For instance, a hearing impaired person could have a visual, tactile, gustatory, *etc.* experience which is visually, tactilely, gustatorily, *etc.* identical to one accompanied with an auditory experience. Without a *that’s all* clause, the propositions describing these two experiences will overlap.

Some may want to contend—though I suspect, due to his *Elusive Knowledge*, that Lewis would not want to contend—that the model of learning that LEWIS (1999) presupposes, in which our evidence consists of a proposition specifying in full detail the course of experience we have undergone, is merely a simplifying idealization. True, agents like *us* don’t gain the evidence that there are 1,157 stars in our visual field; but a suitably idealized agent *would* gain this evidence. I certainly believe that simplifying idealizations can prove epistemically interesting and theoretically fruitful, but once we start making idealizations about the kind of evidence that agents acquire, it’s no longer clear that this idealization has any interesting connection to the epistemic rationality of agents like us. Why should we use the doxastic state of an agent with more evidence to evaluate the epistemic rationality of the doxastic state of an agent with less evidence? Simply because your degrees of belief are not those of an agent who has more evidence than you do, this doesn’t mean that you have fallen short of ideal epistemic rational-

²⁰ LEWIS (1999, p. 405).

²¹ Cf. the ‘speckled hen’ problem of AYER (1940).

ity, any more than the fact that your degrees of belief aren't those of an omniscient agent means that you have fallen short of ideal epistemic rationality. *A fortiori*, having the same degrees of belief as an agent with more evidence than you is frequently epistemically *irrational*; if you were to become absolutely certain that the number of stars in your visual field is 1,157 when you haven't counted them all up, you would be epistemically irrational, even if you happened to hit upon the correct number.

Now, we might think that ideal epistemic rationality requires acquiring all the evidence that is, in some sense, *available* to us. And, when we look at the night sky, the evidence of how many stars are in our visual field is—in some sense—available to us. Even so, this doesn't mean that, if we fail to acquire all the evidence that is available to us, we should act as though we have. Nor does it mean that, if we are absolutely certain that we will not acquire all the evidence that is in some sense available to us, we should plan to revise our degrees of belief as though we were going to acquire all of that evidence. Suppose that a man has an envelope in front of him which he knows either contains evidence that his wife has been faithful or evidence that she has cheated on him. Suppose that the man, because he has faith in his wife, resolves to not open the envelope.²² Perhaps failing to look in the envelope constitutes a failure of epistemic rationality. However, what is certainly *not* a failure of rationality is failing to either become more confident that his wife has been faithful or to become more confident that his wife has cheated. Even if there's some epistemic irrationality going on here, the irrationality has to do with the failure to collect available evidence; it has nothing to do with the way that the man's credences evolve over time, given the evidence he has collected. Similarly, even if there's some irrationality involved in failing to acquire the evidence that there are 1,157 stars in your visual field, the irrationality has nothing to do with the way that your credences evolve over time. And a full theory of diachronic rationality—even *ideal* diachronic rationality—ought to tell us something about what to do when we will fail to acquire all the evidence that we have (in some sense) at our disposal, and we know ahead of time that we will fail to acquire all of this evidence. Even if the man falls short of ideal epistemic rationality by failing to look in the envelope, we can still ask what the epistemically ideal response to his (not ideally rationally collected) evidence is. Similarly, we can ask about the epistemically ideal response to a body of evidence which lacks the information of how many stars are in your visual field, even if any such body of evidence must have been not ideally rationally collected.

One final point about appeals to the nature of evidence: Suppose we grant that an agent could never be in a situation in which it is metaphysically possible for them to acquire total evidence e_1 , it is metaphysically possible for them to acquire total evidence e_2 , and e_1 and e_2 are consistent. Even so, it does not follow that the agent's learning scenario must form a partition. For a learning scenario is defined in terms of *epistemic*, and not *metaphysical*, modality. If it is possible for an agent to rationally believe that they might learn something, yet might learn nothing, then it is possible for a rational agent to be in a learning scenario which doesn't form a partition—even if, for whatever

²² The example comes from BUCHAK (2012), who argues that failure to look in the envelope does not constitute a failure of ideal epistemic rationality.

reason, it is metaphysically impossible to learn nothing. So any argument attempting to establish that every learning scenario forms a partition by appealing to some property of evidence will require an additional premise stating that ideally rational agents must be certain that evidence has that property.

So I don't think that, in general, claims about the nature of evidence are going to do much to help reach the conclusion that learning scenarios always form a partition. Even if there's some reason to think that the potential *available* evidence forms a partition, an agent could know full well that they will fail to acquire all of the available evidence. Additionally, even if the set of all (metaphysically) learnable evidence propositions forms a partition, this on its own is insufficient to establish that every learning scenario forms a partition, since an agent could fail to know this fact about evidence, and so believe (falsely) that the set of propositions they might learn doesn't form a partition. So I'll move on to arguments alleging that the very nature of a learning scenario guarantees that a learning scenario will form a partition.

Along these lines, here is a promising argument for the conclusion that every learning scenario forms a partition: in order to be in a learning scenario $\mathcal{S} = \{e_1, e_2, \dots, e_N\}$, it must be epistemically possible that your total evidence will be e_i , for each $e_i \in \mathcal{S}$, and it must be epistemically necessary that exactly one $e_i \in \mathcal{S}$ will be your total evidence. This is just the claim that $\{\mathbb{T}e_1, \mathbb{T}e_2, \dots, \mathbb{T}e_N\}$ forms a partition (where, recall, $\lceil \mathbb{T}\phi \rceil$ says that ϕ is your total evidence). Assume that evidence is factive, so that $\mathbb{T}e_i \models e_i$. Assume further that, if you acquire the total evidence e_i , then you also acquire the evidence that e_i is your total evidence, $\mathbb{T}e_i \models \mathbb{E}\mathbb{T}e_i$ (where $\lceil \mathbb{E}\phi \rceil$ says that ϕ is evidence for you). It then follows that, for each $e_i \in \mathcal{S}$, $e_i \models \mathbb{T}e_i$.²³ Thus, for every $e_i \in \mathcal{S}$, $e_i \models \mathbb{T}e_i$, and the claim that $\{\mathbb{T}e_1, \mathbb{T}e_2, \dots, \mathbb{T}e_N\}$ forms a partition is just the claim that $\{e_1, e_2, \dots, e_N\}$ forms a partition. So, every learning scenario must form a partition.²⁴

Applying this suggestion to the case of A, B , and Γ above: suppose that, prior to hearing either α, β , or γ , you hold credal opinions about the propositions $\mathbb{E}g, \mathbb{E}b, \mathbb{E}c$, and $\mathbb{E}s$. Upon hearing α , A eliminates the possibilities $\mathbb{E}c$ and $\mathbb{E}s$. Upon hearing β , B eliminates the possibilities $\mathbb{E}c$ and $\mathbb{E}b$. And upon hearing γ , Γ eliminates the possibilities $\mathbb{E}g$ and $\mathbb{E}b$. They thus end up with the credal states in the middle of figure 7. Suppose also that, if B learns that g , then they also learn that $\mathbb{E}g$; and if they learn that s , then they also learn that $\mathbb{E}s$. Thus, B 's total evidence will either be $\mathbb{E}g$ or $\mathbb{E}s$. And $\{\mathbb{E}g, \mathbb{E}s\}$ partitions B 's prior credal state, so **DB** mandates that they update by conditionalizing on this total evidence, and the tension between **PE** and **DB** is alleviated.

Because this argument appeals to the internalist thesis that you always have the evidence of what your total evidence is—that is, $\mathbb{T}e \models \mathbb{E}\mathbb{T}e$ —I will call it the *argument from internalism*. I have two worries about the argument from internalism.

²³ Suppose that $\mathbb{T}e_i \models e_i$, yet $e_i \not\models \mathbb{T}e_i$. Then, $\mathbb{T}e_i$ would be stronger than e_i . But, since $\mathbb{T}e_i \models \mathbb{E}\mathbb{T}e_i$, $\mathbb{T}e_i$ is learned if e_i is your total evidence. Since $\mathbb{T}e_i$ is stronger than e_i , this contradicts our assumption that e_i is the *strongest* proposition learned. So $e_i \models \mathbb{T}e_i$.

²⁴ HILD (1998b) and WEISBERG (2007) both consider defenses of the assumption that learning scenarios form a partition along these general lines.

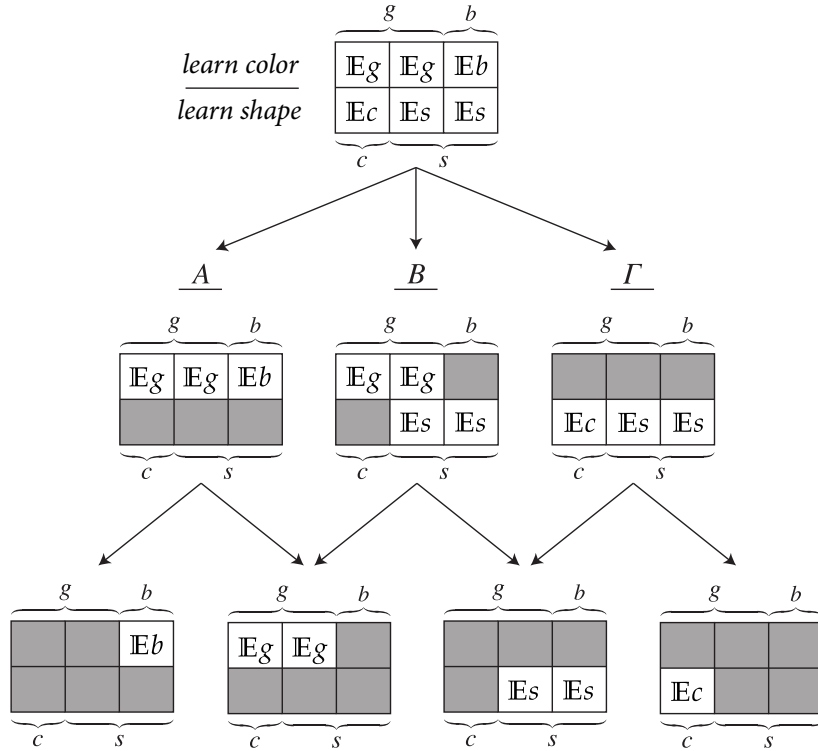


FIGURE 7: If you begin with the propositions that you acquire the evidence that g , b , c , or s — $\mathbb{E}g$, $\mathbb{E}b$, $\mathbb{E}c$, and $\mathbb{E}s$ —included in your algebra \mathcal{A} , and if $\mathbb{T}\phi \models \mathbb{E}\mathbb{T}\phi$, then, upon hearing α , β , or γ , \mathcal{A} 's B 's and Γ 's learning scenarios will each form a partition.

Firstly, I have familiar worries about internalism itself. The internalist assumption that $\mathbb{T}e \models \mathbb{E}\mathbb{T}e$ is equivalent to the S5 ‘negative access’ principle $\neg\mathbb{E}e \models \mathbb{E}\neg\mathbb{E}e$, which is equivalent to the conjunction of the S4 ‘positive access’ principle $\mathbb{E}e \models \mathbb{E}\mathbb{E}e$ and the Brouwer principle $\neg e \models \mathbb{E}\neg\mathbb{E}e$.²⁵

WILLIAMSON (2014) argues that any case in which our perceptual knowledge is inexact provides a counterexample to the S4 ‘positive access’ principle $\mathbb{E}\phi \models \mathbb{E}\mathbb{E}\phi$. Suppose that, off in the distance, you catch a glimpse of an unmarked clock (see figure 8(a)). Your vision is good enough that you get the evidence that the hand is on the right-hand side of the clock. And though you likely learn something stronger still, it seems undeniable that your evidence does not tell you the precise location of the clock hand. At most, it tells you that the clock hand is located in some interval (see figure 8(b)). It additionally seems reasonable to suppose that your evidence about the quality of your vision in these kinds of cases tells you this: if the clock hand is located at some position l , then you won’t get the evidence that the clock hand is located within some interval that has l as an endpoint (see figure 8(c)). That’s not to say that your evidence about the position of the clock hand must be symmetric around the

²⁵ These equivalences rely upon a standard Kripke semantics for the operators \mathbb{T} and \mathbb{E} , where $\lceil \mathbb{E}\phi \rceil$ is true at a world w iff $\lceil \phi \rceil$ is true at all worlds to which w bears the accessibility relation R , and $\lceil \mathbb{T}\phi \rceil$ is true at w iff $\lceil \phi \rceil$ is true at *all and only* worlds to which w bears R .




FIGURE 8: A distant and brief glimpse at the unmarked clock (8(a)) provides the evidence that the clock hand is positioned within some interval of values $[u, l]$ (8(b)); and, if the clock hand is positioned at θ degrees from 12 o'clock, then the glimpse does not provide the evidence that the clock hand is no more than θ degrees from 12 o'clock (8(c)). These assumptions contradict the 'positive access' principle for evidence, $\mathbb{E}\phi \models \mathbb{E}\mathbb{E}\phi$.

clock hand's actual position. It's just to say that you've learned enough about your visual capacities in circumstances like these to conclude that your evidence about the position of the clock hand will not be *maximally* asymmetric, ruling out all positions ever so slightly further than l , when in fact the clock hand is at l .

But these minimal assumptions are enough to conclude that the positive access principle is false. For the following three claims are inconsistent (Note: I'll use ' θ ' to stand for the position of the clock hand—see figure 8.)

- 12) The strongest proposition you learn about the position of the clock hand is that it lies in some (non-trivial) interval—call it $[u, l]$ (' u ' for *upper* and ' l ' for *lower*).
- 13) The following is entailed by your total evidence: if the clock hand is located at l , then your evidence doesn't tell you that it is located no further than l .

$$\mathbb{E}[\theta = l \rightarrow \neg\mathbb{E}(\theta \leq l)]$$

- 14) The possession of evidence is always evidence itself.

$$\mathbb{E}\phi \rightarrow \mathbb{E}\mathbb{E}\phi$$

To see that they are inconsistent, note that, by contraposition, (13) tells us that your evidence entails that, if you learn that the clock hand is no further than l , then the clock hand is not at l .

$$(15) \quad \mathbb{E}[\mathbb{E}(\theta \leq l) \rightarrow \theta \neq l]$$

If we assume that the evidence operator \mathbb{E} satisfies the K -axiom, then (15) entails that, if you learn that you've learned that the clock hand is no further than l , then you also learn that the clock hand is not *at* l .

$$(16) \quad \mathbb{E}\mathbb{E}(\theta \leq l) \rightarrow \mathbb{E}(\theta \neq l)$$

Now, (12) tells us that you *do* learn that the clock hand is no further than l .

$$\mathbb{E}(\theta \leq l)$$

And (14) tells us that, if you learn this, then you learn that you have learned it,

$$(17) \quad \mathbb{E}\mathbb{E}(\theta \leq l)$$

And from (17) and (16), it follows that you learn, not just that the clock hand is no further than l , but moreover that it is not *at* l , either.

$$\mathbb{E}(\theta \neq l)$$

This, however, contradicts (12), which assured us that the *strongest* thing you learned about the position of the clock hand was that it was within the interval $[u, l]$. Since this does not entail $\theta \neq l$, (12) tells us that you cannot have learned it.²⁶ So (12), (13), and (14) are inconsistent, and the positive access principle (14) looks to be the least plausible of the bunch. So we have reason to reject the S4 ‘positive access’ principle. And since the argument from internalism’s premise $\mathbb{T}e \models \mathbb{E}\mathbb{T}e$ entails S4, rejecting S4 means giving up on that argument as well.

Note that appeals to evidence about what you have learned will do nothing to blunt the force of WILLIAMSON’s argument. For (12) does not say that $\theta \in [u, l]$ is the strongest proposition learned *full stop*. It says only that it is the strongest proposition learned *about the position of the clock hand*. So contending that you additionally learned that you learned that $\theta \in [u, l]$ or that you guessed thus-and-so about the position of the clock hand,²⁷ does nothing to call (12) into question.

Cases of false memories provide an argument against the Brouwer principle. In normal circumstances, evidence is preserved through memory. If I acquire the evidence e at 12:00, then my memory allows me to still possess the evidence e at 12:01. Today, circumstances aren’t normal. My memory has been tampered with, and though I seem to remember that e , in fact, $\neg e$. But I have no evidence that my memory has been tampered with, so my evidence does not rule out that I am in normal circumstances. In normal circumstances, evidence is preserved in memory. So my evidence does not rule out that I now have the evidence that e . So $\neg e$ and yet $\neg\mathbb{E}\neg\mathbb{E}e$. So the Brouwer principle is false. We could of course reject the claim that, in normal circumstances, evidence is preserved through memory. If we say this, however, one begins to wonder how to make sense of claims like CONDI in the first place. When should I conditionalize on e ? Presumably *after* I’ve learned it; but if evidence isn’t preserved through memory in normal circumstances, then CONDI would advise me to be unduly confident in e .

²⁶ The reader may be wondering whether this contradiction may be avoided by exchanging (12)’s closed interval $[u, l]$ for an open one (u, l) . This claim would be inconsistent with (14) and the following principle, for any choice of $\epsilon > 0$, no matter how small—the reasoning is exactly the same as in the body, *mutatis mutandis*:

$$\mathbb{E}[\theta = l - \epsilon \rightarrow \neg\mathbb{E}(\theta < l)]$$

²⁷ See STALNAKER (2009) and SALOW (forthcoming).

So I think the conditionalizer has good reason to steer clear of that skeptical road; and therefore, has good reason to reject the Brouwer principle. Since the argument from internalism's premise $\mathbb{T}e \models \mathbb{E}\mathbb{T}e$ entails the Brouwer principle, rejecting the Brouwer principle means giving up on that argument as well.

Note again that appeals to evidence about what you have learned will do nothing to blunt the force of this argument. For the argument does not say anything at all about what your *total* evidence is. It only appeals to the premise that, in normal circumstances, evidence (whatever that evidence may be) is preserved through memory, and the premise that, in abnormal circumstances in which your memory has been tampered with, you don't have the evidence that you're not in normal circumstances.

Secondly: even if the argument from internalism is able to establish that every learning scenario partitions the agent's credal state, it is only able to do so by making it rather plausible that not all evidence is acquired in a learning scenario. The argument from internalism concludes that, in order to be in a learning scenario $\mathcal{S} = \{e_1, e_2, \dots, e_N\}$, each $e_i \in \mathcal{S}$ must be such that $\mathbb{T}e_i \models e_i$ —i.e., you recognize no possibilities in which your total evidence is true without it being the case that you acquire that total evidence. If we think that it's possible to acquire evidence without this condition being met—if, for instance, we think that it is possible to learn something even when you thought that you might have learned nothing—then we will think that it's possible to acquire evidence outside of any learning scenario, even if we're persuaded that all learning scenarios form a partition of the agent's credal state. And, if it's possible to acquire evidence outside of any learning scenario, then the DDBA for CONDI fails to establish that you ought to update by conditionalization in those cases.

For an illustration of this second problem, look back to the internalist argument that B 's learning scenario must form a partition, illustrated in figure 7. This argument was able to establish that B 's learning scenario formed a partition only by conceding that the respective certainties rationalized by hearing α , β , and γ overlap. Look to the set of total evidence propositions which were supposed to be responsible for getting A , B , and Γ into their partitioned learning scenarios in the first place, viz., $\{\mathbb{E}g \cup \mathbb{E}b, \mathbb{E}g \cup \mathbb{E}s, \mathbb{E}s \cup \mathbb{E}c\}$. The internalist response to the case of A , B , and Γ relied upon it being possible for you to learn any of the members of this set. Either this set (or some superset of it) could constitute a learning scenario or it could not. If it could constitute a learning scenario, then it is possible to be in a learning scenarios which doesn't form a partition, and accepting the premises of the DDBA is inconsistent with accepting CONDI. If it couldn't, then the internalist response commits us to saying that it's possible to acquire evidence outside of any learning scenario, and the DDBA says nothing about how we ought to revise our degrees of belief in such situations.

3.2 THE PROBLEM FOR THE DDBA FOR JCONDI, AND YET ANOTHER PROBLEM FOR THE DDBA FOR CONDI

Suppose that you're in the same experiment described at the start of §3.1, with two minor variations. If the room is pitch black, then a guide will take your hand and brush it quickly against the object. This quick tactile sensation will either rationalize raising your credence that the object is cubical or it will rationalize raising your credence

that the object is spherical—but it will not rationalize raising either of those credences all the way to 1. If the room is illuminated, on the other hand, then the lighting will be very dim, and you will only be allowed a quick glimpse of the object. This glimpse will either rationalize raising your credence that the object is green or it will rationalize raising your credence that the object is blue—but it won't rationalize raising either of those credences all the way to 1. Additionally, at the conclusion of the experiment, once everyone has been in the room, you will be told the object's shape and color. (This experiment is, therefore, analogous to Skyrms's two-stage learning scenario.)

Once again, your friend is ahead of you in line; and after they leave the room, they say one of the following sentences to you:

- α^*) You'll either end up more confident that it's green or more confident that it's blue.
- β^*) You'll either end up more confident that it's green or more confident that it's spherical.
- γ^*) You'll either end up more confident that it's cubical or more confident that it's spherical.

Call the version of you that hears α^* ' A^* '; call the version of you that hears β^* ' B^* '; and call the version of you that hears γ^* ' Γ^* '. A^* 's learning scenario consists of a Jeffrey shift raising their credence that g and a Jeffrey shift raising their credence that b . B^* 's learning scenario consists of a Jeffrey shift raising their credence that g and a Jeffrey shift raising their credence that s . Γ^* 's learning scenario consists of a Jeffrey shift raising their credence that c and a Jeffrey shift raising their credence that s .

As we saw in §2.3, A^* is diachronically Dutch book-able if they fail to update by Jeffrey conditionalization in response to a Jeffrey shift raising their credence that g ; and Γ^* is diachronically Dutch book-able if they fail to update by Jeffrey conditionalization in response to a Jeffrey shift raising their credence that s . However, as we saw in §2.4, B^* is diachronically Dutch bookable if they update by Jeffrey conditionalization both in response to a Jeffrey shift raising their credence that g and in response to a Jeffrey shift raising their credence that s .

This case brings out a tension between the following two theses.

Rational Updating is Determined by Priors and Evidence* (PE*) The rational way to update your credences is determined by your prior credence function and the acquired evidence alone.

Rational Updating is not Diachronically Dutch Book-able (DB) If you rationally update your prior credences, then you will not be diachronically Dutch book-able.

The tension is that, if triads of learning scenarios like A^* 's, B^* 's, and Γ^* 's are possible, then PE* and DB are inconsistent. The argument is precisely the same as in §3.1, with ' A^* ' swapped out for ' A ', ' B^* ' swapped out for ' B ', and ' Γ^* ' swapped out for ' Γ '.

The DDBAs for CONDI and JCONDI employ **DB** as a premise, and both CONDI and JCONDI, as they are usually understood, entail **PE***. This means that the possibility of triads of learning scenarios like A^* 's, B^* 's, and Γ^* 's doesn't merely show that the DDBA for JCONDI is invalid. It shows that the DDBA for CONDI is invalid as well. That is: if you accept the possibility of triads of learning scenarios like A^* 's, B^* 's, and Γ^* 's, then you must accept that the DDBA for CONDI is invalid, even if you think that all learning scenarios consisting of *propositional* evidence form a partition.

Moreover, here we cannot appeal to propositions about our own experience, or propositions like 'I underwent the Jeffrey shift E ' as we did in the case of CONDI. For if there are propositions like these in \mathcal{A} for the agent to conditionalize upon, then there would be no need for JCONDI in the first place. If they learn a proposition saying that they've acquired the Jeffrey shift E , then the agent may simply conditionalize upon this proposition, and there is no need for JCONDI. JCONDI was motivated by precisely the thought that, in scenarios like A^* 's, B^* 's, and Γ^* 's, *there is no* proposition for them to conditionalize upon. To deny this is to deny JCONDI's *raison d'être*.

3.3 TWO THEORIES OF EVIDENTIAL IMPORT

Stepping back from the argument that **PE** (or **PE***) and **DB** are inconsistent, I believe we can diagnose what's going on at a more general level by noting that, if non-partitioning learning scenarios are possible, then a DDBA and an update rule like CONDI entail incompatible claims about evidential import—incompatible claims about what determines the rational response to a piece of total evidence. Because it requires a specification of the agent's learning scenario, and because its verdicts are sensitive to whether this learning scenario forms a partition, if non-partitioning learning scenarios are possible, the premises of a DDBA entail a *contrastivist* thesis about evidential import.

Contrastive Evidential Import (CEI) The rational response to evidence depends upon what other evidence might have been received instead.

CONDI and JCONDI, on the other hand, because they do not require a specification of which other propositions the agent might have acquired instead, entail a *noncontrastivist* thesis about evidential import.

Noncontrastive Evidential Import (NEI) The rational response to evidence does not depend upon what other evidence might have been received instead.

According to **CEI**, there's at least some prior doxastic state and some collection of propositions $e, e_1^*, e_2^*, \dots, e_N^*, e'_1, e'_2, \dots, e'_M$ such that learning that e , *rather than* e_1^*, e_2^*, \dots , or e_N^* , could warrant a different reaction than learning that e , *rather than* e'_1, e'_2, \dots , or e'_M . According to **NEI**, there is no such prior doxastic state and collection of propositions.

If there could be two different learning scenarios, \mathcal{S}_1 and \mathcal{S}_2 , both containing the very same evidence e , such that a response to e which is not diachronically Dutch book-able in \mathcal{S}_1 is diachronically Dutch book-able in \mathcal{S}_2 , then accepting a DDBA means accepting **CEI**, which means rejecting rules like CONDI and JCONDI, which are inconsistent with **CEI**.

4 CONCLUSION

The problem from §3.2 leaves very little wiggle room. As I see things, accepting the *DDBA* for *JCONDI* means rejecting *NEI*, and therefore, rejecting both *CONDI* and *JCONDI*.

With respect to the problem for the *DDBA* for *CONDI* (§3.1), there is more wiggle room. How devastating the problem is depends in part upon whether we are persuaded that non-partitioning learning scenarios are possible. If not—if we end up thinking that every learning scenario must form a partition—then we might want to hold on to the *DDBA* for *CONDI* but concede that, while the *DDBA* shows that you should conditionalize in learning scenarios, it doesn't show that you ought to conditionalize when you acquire evidence outside of a learning scenario.

If we don't think that non-partitioning learning scenarios are possible, then we would be free to appeal to other arguments that we ought to conform to *CONDI* (and *JCONDI*) when we acquire evidence outside of a learning scenario. Such arguments are available. There is an argument from symmetry,²⁸ from calibration,²⁹ from suppositional consistency,³⁰ and a handful of arguments that attempt to justify *CONDI* and *JCONDI* on the grounds that conforming to those rules will maximize the expected accuracy of your posterior credences,³¹ or that all other belief-revision norms are accuracy-dominated.³² With the exception of GREAVES & WALLACE (2006)'s and EASWARAN (2013)'s expected accuracy maximization arguments, and BRIGGS & PETTIGREW (ms)'s 'accuracy-dominance' argument for *CONDI*, none of these justificatory strategies require any specification of the agent's learning scenario. Thus, if we think that all learning scenarios form a partition, then, while we're not forced to abandon the *DDBA* for *CONDI*, we will end up thinking that they are less general than many other arguments for that update rule.

Suppose, on the other hand, that we think that non-partitioning learning scenarios are possible. Then, we face a choice. We could either endorse *CEI* or *NEI*. If we opt for *CEI*, then we would have to revise *CONDI*. That belief-revision rule, recall, says:

CONDI

Upon undergoing an experience which rationalizes absolute certainty that e , nothing stronger, and nothing else, you should be disposed to conditionalize on e .

$$C_e(\cdot) \stackrel{!}{=} C(\cdot | e)$$

There is no mention of a contrast class of other potential evidence propositions. If we opt for *CEI*, then we would have to exchange *CONDI* for a more qualified belief revision norm like *CONDI**.

²⁸ See VAN FRAASSEN (1989, ch. 13)

²⁹ See LANGE (1999)

³⁰ See TITELBAUM (2013, ch. 7)

³¹ See WILLIAMS (1980), DIACONIS & ZABELL (1982), ODDIE (1997), GREAVES & WALLACE (2006), LEITGEB & PETTIGREW (2010), LEVINSTEIN (2012), and EASWARAN (2013)

³² See BRIGGS & PETTIGREW (ms)

CONDI*

Upon undergoing an experience which rationalizes absolute certainty that e , and might have rationalized absolute certainty in any e_i in a partition $\{e_1, \dots, e_N\}$, you should be disposed to conditionalize on e .

$$C_e(\cdot) \stackrel{!}{=} C(\cdot | e)$$

And, presumably, we would want to generalize this rule to cover cases where the contrast class of evidence propositions fail to be disjoint from e .³³ Additionally, if we think that non-partitioning learning scenarios are possible and we accept **CEI**, then we must reject all the justifications of **CONDI** and **JCONDI** which entail **NEI** mentioned in the previous paragraph.

Alternatively, we could accept that non-partitioning learning scenarios are possible and reject **CEI** in favor of **NEI**. If we take this tack, then we must reject premise (4) of the **DDBA** for **CONDI**. This doesn't mean forsaking **CONDI**. As I noted above, there are other arguments for that belief-revision rule which don't involve specifying which other evidence an agent might have acquired instead, and are thus consistent with **NEI**. However, it does mean forsaking the **DDBA** for **CONDI**, as well as **DDBAs** for any other putative principle of rationality—like, for instance, the **DDBA** for **REFLECTION** and **HITCHCOCK** (2004)'s **DDBA** for the thirder's solution to the Sleeping Beauty puzzle.

³³ See [GALLOW \(ms\)](#) for suggestions.

A A GENERAL DIACHRONIC DUTCH-BOOK STRATEGY AGAINST THE CONDITIONALIZER

Consider a learning scenario in which some of the potential evidence propositions overlap—that is, consider a learning scenario $\mathcal{S} = \{e_1, e_2, \dots, e_N\}$ such that, for some $\{e_1^* \dots e_M^*\} \subseteq \mathcal{S}$, $\bigcap_i e_i^* \neq \emptyset$. Say that \mathcal{S}^* is a *maximal overlapping subset* of \mathcal{S} , iff both i) $\bigcap_{e_i^* \in \mathcal{S}^*} e_i^* \neq \emptyset$, and ii) for any $e_j \in \mathcal{S} - \mathcal{S}^*$, $e_j \cap \bigcap_i e_i^* = \emptyset$. And say that a maximal overlapping subset of \mathcal{S} has *breathing room* iff, conditional on at least one of the $e_i^* \in \mathcal{S}^*$ being true and all of the $e_j \in \mathcal{S} - \mathcal{S}^*$ being false, the agent is neither certain that all of the $e_i^* \in \mathcal{S}^*$ are true, nor certain that not all of the $e_i^* \in \mathcal{S}^*$ are true,

$$0 < C \left(\bigcap_{e_i^* \in \mathcal{S}^*} e_i^* \mid \bigcup_{e_i^* \in \mathcal{S}^*} e_i^* \cap \bigcap_{e_j \in \mathcal{S} - \mathcal{S}^*} \neg e_j \right) < 1.$$

In any learning scenario \mathcal{S} in which a maximal overlapping subset of \mathcal{S} has breathing room, conditionalization is diachronically Dutch book-able. To see this, consider the proposition b (for *breathing room*), defined as follows,

$$b \stackrel{\text{def}}{=} \bigcup_{e_i^* \in \mathcal{S}^*} e_i^* \cap \bigcap_{e_j \in \mathcal{S} - \mathcal{S}^*} \neg e_j$$

and consider the proposition o (for *overlap*), defined as follows,

$$o \stackrel{\text{def}}{=} \bigcap_{e_i^* \in \mathcal{S}^*} e_i^*$$

If \mathcal{S}^* has breathing room, then $0 < C(o \mid b) < 1$, for all $e_i^* \in \mathcal{S}^*$, $C(o \mid b \cap e_i^*) \geq C(o \mid b)$, and for some $e_i^* \in \mathcal{S}^*$, $C(o \mid b \cap e_i^*) > C(o \mid b)$.

Then, a diachronic Dutch book strategy against the conditionalizer may be constructed as follows. First, prior to their learning which $e_i \in \mathcal{S}$ is true, sell the conditionalizer bets 7 and 8.

Bet 7	Bet 8										
<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; text-align: center;">\$$C(o \mid b) - 1$</td> <td style="width: 50%; text-align: center;">if $o \cap b$</td> </tr> <tr> <td style="text-align: center;">\$$C(o \mid b)$</td> <td style="text-align: center;">if $\neg o \cap b$</td> </tr> <tr> <td style="text-align: center;">\$0</td> <td style="text-align: center;">if $\neg b$</td> </tr> </table>	\$ $C(o \mid b) - 1$	if $o \cap b$	\$ $C(o \mid b)$	if $\neg o \cap b$	\$0	if $\neg b$	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; text-align: center;">\$$\delta[1 - C(b)]$</td> <td style="width: 50%; text-align: center;">if b</td> </tr> <tr> <td style="text-align: center;">\$$-\delta C(b)$</td> <td style="text-align: center;">if $\neg b$</td> </tr> </table>	\$ $\delta[1 - C(b)]$	if b	\$ $-\delta C(b)$	if $\neg b$
\$ $C(o \mid b) - 1$	if $o \cap b$										
\$ $C(o \mid b)$	if $\neg o \cap b$										
\$0	if $\neg b$										
\$ $\delta[1 - C(b)]$	if b										
\$ $-\delta C(b)$	if $\neg b$										

where $\delta \stackrel{\text{def}}{=} \min_{e_k^* \in \mathcal{S}^*} \{C(o \mid b \cap e_k^*) - C(o \mid b)\}$. (If, for some $e_k^* \in \mathcal{S}^*$, $C(o \mid b \cap e_k^*) = C(o \mid b)$, then $\delta = 0$ and bet 10 is guaranteed to pay out \$0 no matter what.) If the conditionalizer learns some proposition $e_j \notin \mathcal{S}^*$, then they will break even on bet 9 and lose $\delta C(b)$ on bet 10 (or, in the event that $\delta = 0$, they will break even on both bets). If, however, they learn one of the $e_i^* \in \mathcal{S}^*$, then they may be sold bet 9.

Bet 9						
<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; text-align: center;">\$$1 - C(o \mid b \cap e_i^*)$</td> <td style="width: 50%; text-align: center;">if $o \cap b$</td> </tr> <tr> <td style="text-align: center;">\$$-C(o \mid b \cap e_i^*)$</td> <td style="text-align: center;">if $\neg o \cap b$</td> </tr> <tr> <td style="text-align: center;">\$0</td> <td style="text-align: center;">if $\neg b$</td> </tr> </table>	\$ $1 - C(o \mid b \cap e_i^*)$	if $o \cap b$	\$ $-C(o \mid b \cap e_i^*)$	if $\neg o \cap b$	\$0	if $\neg b$
\$ $1 - C(o \mid b \cap e_i^*)$	if $o \cap b$					
\$ $-C(o \mid b \cap e_i^*)$	if $\neg o \cap b$					
\$0	if $\neg b$					

This strategy will net the conditionalizer at most $-\delta C(b)$. If $C(o \mid b \cap e_i^*) = C(o \mid b)$ for some $e_i^* \in \mathcal{S}^*$, so that $\delta = 0$, then this strategy will net the conditionalizer \$0 in some contingencies, and lose them money in others. They will have opened themselves up to the risk of losing money without any possibility of winning money. If $C(o \mid b \cap e_i^*) > C(o \mid b)$ for

every $e_i^* \in \mathcal{S}^*$, then this strategy will lose the conditionalizer money come what may.

B DIACHRONIC DUTCH BOOK-ABILITY IN NON-PARTITIONING LEARNING SCENARIOS

Given a probability space $\langle \mathcal{W}, \mathcal{A}, C \rangle$, we may, utilizing the technique of DIACONIS & ZABELL (1982, theorem 2.1, p. 824), construct a new probability space $\langle \mathcal{W}^*, \mathcal{A}^*, C^* \rangle$, such that a) $\mathcal{W}^* = (\mathcal{W} - e_1 \cap e_2) \cup (e_2 \cap e_2 \times \{1, 2\})$; b) for every $p \in \mathcal{A}$, there is a $p^* \in \mathcal{A}^*$ such that $p^* \stackrel{\text{def}}{=} (p \cap \neg e_1) \cup (p \cap \neg e_2) \cup (p \cap e_2 \cap e_2)^*$, and $(p \cap e_2 \cap e_2)^* \stackrel{\text{def}}{=} \{\langle w, 1 \rangle, \langle w, 2 \rangle \mid w \in p \cap e_1 \cap e_2\}$; and c) for each $p \in \mathcal{A}$, $C^*(p^*) = C(p)$.

That is to say, we may enrich the algebra of the probability space to include two new propositions $x_1, x_2 \subseteq e_2 \cap e_2$, where $x_1 \stackrel{\text{def}}{=} \bigcup_{w \in e_1 \cap e_2} \langle w, 1 \rangle$, and $x_2 \stackrel{\text{def}}{=} \bigcup_{w \in e_1 \cap e_2} \langle w, 2 \rangle$. (Note that C^* will not be unique—there will be several choices of C^* consistent with the constraints above.)

Suppose that an agent A^* with a credal state $\langle \mathcal{W}^*, \mathcal{A}^*, C^* \rangle$ will either learn that $\neg e_2 \cup x_1$ or that $\neg e_1 \cup x_2$. Then, conditionalization is the strategy $\langle C_1(p), C_2(p) \rangle$, where

$$\begin{aligned} C_1(p) &= C^*(p \mid \neg e_2 \cup x_1) & C_2(p) &= C^*(p \mid \neg e_1 \cup x_2) \\ &= \frac{C^*(p \cap \neg e_2) + C^*(p \cap x_1)}{C^*(\neg e_2) + C^*(x_1)} & &= \frac{C^*(p \cap e_2) - C^*(p \cap x_1)}{C^*(e_2) - C^*(x_1)} \end{aligned}$$

By the construction of C^* , $C^*(p \cap \neg e_2) = C(p \cap \neg e_2)$, $C^*(\neg e_2) = C(\neg e_2)$, $C^*(p \cap e_2) = C(p \cap e_2)$, and $C^*(e_2) = C(e_2)$. However, the construction of C^* does not fix the values of $C^*(p \cap x_1)$ and $C^*(x_1)$ beyond the constraints that

$$\begin{aligned} 0 &\leq C^*(p \cap x_1) \leq C^*(x_1) \leq C(e_1 \cap e_2) \\ C^*(p \cap x_1) &\leq C(p \cap e_1 \cap e_2) \\ C(x_1) - C(p \cap x_1) &\leq C(e_1 \cap e_2) - C(p \cap e_1 \cap e_2) \end{aligned}$$

Thus, for every $\langle x, y \rangle$ within the parallelogram with vertices

$$\begin{aligned} &\langle 0, 0 \rangle, \\ &\langle C(p \cap e_1 \cap e_2), C(p \cap e_1 \cap e_2) \rangle, \\ &\langle C(e_1 \cap e_2), C(p \cap e_1 \cap e_2) \rangle, \\ &\langle C(e_1 \cap e_2) - C(p \cap e_1 \cap e_2), 0 \rangle \end{aligned}$$

there is a credal state $\langle \mathcal{W}^*, \mathcal{A}^*, C^* \rangle$ such that $\langle C^*(x_1), C^*(p \cap x_1) \rangle = \langle x, y \rangle$. (See figure 9.) Call the set of points $\langle C^*(x_1), C^*(p \cap x_1) \rangle$ contained within this parallelogram ' S_p^C '. Then, the set of conditionalization strategies $\langle C_1(p), C_2(p) \rangle$ for credal states $\langle \mathcal{W}^*, \mathcal{A}^*, C^* \rangle$ constructible as specified above is found by taking the points inside S_p^C under the transformation \mathcal{T}_p^C

$$\mathcal{T}_p^C : \langle x, y \rangle \mapsto \left\langle \frac{C(p \cap \neg e_2) + y}{C(\neg e_2) + x}, \frac{C(p \cap e_2) - y}{C(e_2) - x} \right\rangle$$

This transformation maps the vertices of S_p^C as shown below.

$$\begin{aligned} \langle 0, 0 \rangle &\mapsto \langle C(p \mid \neg e_2), C(p \mid e_2) \rangle \\ \langle C(e_1 \cap e_2), C(p \cap e_1 \cap e_2) \rangle &\mapsto \langle C(p \mid e_1), C(p \mid \neg e_1) \rangle \\ \langle C(p \cap e_1 \cap e_2), C(p \cap e_1 \cap e_2) \rangle &\mapsto \left\langle \frac{C(p \mid e_1)}{C(p \cup \neg e_2 \mid e_1)}, \frac{C(p \cap \neg e_1 \mid e_2)}{C(\neg(p \cap e_1) \mid e_2)} \right\rangle \end{aligned}$$




FIGURE 9: The set S_p^C of pairs of values $\langle C^*(x_1), C^*(p \cap x_1) \rangle$ permitted by the construction of C^* .

$$\langle C(e_1 \cap e_2) - C(p \cap e_1 \cap e_2), 0 \rangle \mapsto \left\langle \frac{C(p \cap \neg e_2 | e_1)}{C(\neg(p \cap e_2) | e_1)}, \frac{C(p | e_2)}{C(p \cup \neg e_1 | e_2)} \right\rangle$$

(See figures 4 and 10.)

Since $\{\neg e_2 \cup x_1, \neg e_1 \cup x_2\}$ partitions \mathcal{W}^* , given any credal state $\langle \mathcal{W}^*, \mathcal{A}^*, C^* \rangle$, the agent, A^* , with that credal state will not be diachronically Dutch book-able if they adopt the strategy of conditionalization (by SKYRMS, 1987, theorem 5, p. 16).

A conditionalizer, A^* , with a credal state $\langle \mathcal{W}^*, \mathcal{A}^*, C^* \rangle$ constructed as defined above will take all the same bets on the propositions in the algebra \mathcal{A} as an agent, A , with the credal state $\langle \mathcal{W}, \mathcal{A}, C \rangle$. So, if there is some conditionalizer A^* whose update strategy agrees with A 's on the propositions in \mathcal{A} , then A is not diachronically Dutch book-able.

In the special case in which $p = e_1 \cap e_2$, the set of non-diachronically Dutch book-able update strategies $\langle C_1(e_1 \cap e_2), C_2(e_1 \cap e_2) \rangle$ is given by

$$C_1(e_1 \cap e_2) = \frac{C^*(x_1)}{C(\neg e_2) + C^*(x_1)} \quad C_2(e_1 \cap e_2) = \frac{C(e_1 \cap e_2) - C^*(x_1)}{C(e_2) - C^*(x_1)}$$

where $C^*(x_1)$ can take on any value between 0 and $C(e_1 \cap e_2)$. So long as $0 < C(e_1 \cap e_2) < 1$, the set of non Dutch book-able strategies $\langle C_1(e_1 \cap e_2), C_2(e_1 \cap e_2) \rangle$ will lie on a curve with endpoints $\langle 0, C(e_1 | e_2) \rangle$ and $\langle C(e_2 | e_1), 0 \rangle$, as shown in figure 3.

REFERENCES

- ARMENDT, BRAD. 1980. "Is There a Dutch Book Argument for Probability Kinematics?" *Philosophy of Science*, vol. 47 (4): 583–588. [2], [14]
- AYER, A. J. 1940. *The Foundations of Empirical Knowledge*. Macmillan, New York. [18]
- BRIGGS, RACHAEL & RICHARD PETTIGREW. ms. "An Accuracy-Dominance Argument for Conditionalization." Available at <https://dl.dropboxusercontent.com/u/9797023/Papers/Conditionalization-RBRP.pdf>. [27]
- BRONFMAN, AARON. 2014. "Conditionalization and not Knowing that One Knows." *Erkenntnis*, vol. 79 (4): 871–892. [12], [17]
- BUCHAK, LARA. 2012. "Can it be Rational to Have Faith?" In *Probability in the Philosophy of Religion*, JAKE CHANDLER & VICTORIA S. HARRISON, editors, chap. 12. Oxford University Press, Oxford. [19]




FIGURE 10: The set S_p^C under the transformation \mathcal{T}_p^C is shown in the grey region. The update strategies $\langle C_{e_1}(p), C_{e_2}(p) \rangle$ within $\mathcal{T}_p^C(S_p^C)$ are immune from diachronic Dutch book-ability in a learning scenario $\{e_1, e_2\}$. (This particular region was obtained from a prior distribution such that $C(\neg e_2) = C(e_1 \cap e_2) = C(\neg e_1) = 1/3$, $C(p \cap \neg e_2) = 1/12$, $C(p \cap e_1 \cap e_2) = 1/6$, and $C(p \cap \neg e_1) = 1/4$.)

CHRISTENSEN, DAVID. 1991. “Clever Bookies and Coherent Beliefs.” *The Philosophical Review*, vol. 100 (2): 229–247. [2], [12], [17]

DIACONIS, PERSI & SANDY ZABELL. 1982. “Updating Subjective Probability.” *Journal of the American Statistical Association*, vol. 77 (380): 822–830. [27], [30]

EASWARAN, KENNY. 2013. “Expected Accuracy Supports Conditionalization—and Conglomerability and Reflection.” *Philosophy of Science*, vol. 80: 119–142. [27]

GALLOW, J. DMITRI. 2014. “How to Learn from Theory-Dependent Evidence; or Commutativity and Holism: A Solution for Conditionalizers.” *The British Journal for the Philosophy of Science*, vol. 65 (3): 493–519. [12]

—. ms. “Updating for Externalists.” [28]

GARDNER, MARTIN. 1961. *The Second Scientific American Book of Mathematical Puzzles and Diversions*. Simon and Schuster, New York. [12]

GREAVES, HILARY & DAVID WALLACE. 2006. “Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility.” *Mind*, vol. 115 (495): 607–632. [27]

HÁJEK, ALAN. 2008. “Dutch Book Arguments.” In *The Oxford Handbook of Rational and Social Choice*, PAUL ANAND, PRASANTA PATTANAİK & CLEMENS PUPPE, editors, 173–195. Oxford University Press, Oxford. [2]

HILD, MATTHIAS. 1998a. “Auto-Epistemology and Updating.” *Philosophical Studies*, vol. 92: 321–361. [17]

- . 1998b. “The Coherence Argument Against Conditionalization.” *Synthese*, vol. 115: 229–258. [10], [11], [17], [20]
- HITCHCOCK, CHRISTOPHER. 2004. “Beauty and the Bets.” *Synthese*, vol. 139: 405–420. [28]
- JEFFREY, RICHARD. 1965. *The Logic of Decision*. McGraw-Hill, New York. [2], [4]
- LANGE, MARC. 1999. “Calibration and the Epistemological Role of Bayesian Conditionalization.” *Journal of Philosophy*, vol. 96 (6): 294–324. [27]
- LEITGEB, HANNES & RICHARD PETTIGREW. 2010. “An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy.” *Philosophy of Science*, vol. 77 (2): 236–272. [27]
- LEVINSTEIN, BENJAMIN ANDERS. 2012. “Leitgeb and Pettigrew on Accuracy and Updating.” *Philosophy of Science*, vol. 79 (3): 413–424. [27]
- LEWIS, DAVID K. 1981. “Causal Decision Theory.” *Australasian Journal of Philosophy*, vol. 59 (1): 5–30. [7]
- . 1996. “Elusive Knowledge.” *Australasian Journal of Philosophy*, vol. 74 (4): 549–567. [4], [7]
- . 1999. “Why Conditionalize?” In *Papers in Metaphysics and Epistemology*, vol. 2, chap. 23, 403–407. Cambridge University Press, Cambridge. [4], [5], [7], [8], [18]
- ODDIE, GRAHAM. 1997. “Conditionalization, Cogency, and Cognitive Value.” *British Journal for the Philosophy of Science*, vol. 48: 533–41. [27]
- SALOW, BERNHARD. forthcoming. “The Externalist’s Guide to Fishing for Compliments.” *Mind*. [23]
- SCHOENFIELD, MIRIAM. forthcoming. “Conditionalization does not (in general) Maximize Expected Accuracy.” *Mind*. [11], [17]
- SKYRMS, BRIAN. 1987. “Dynamic Coherence and Probability Kinematics.” *Philosophy of Science*, vol. 54 (1): 1–20. [2], [6], [12], [13], [14], [31]
- STALNAKER, ROBERT C. 2009. “On Hawthorne and Magidor on Assertion, Context, and Epistemic Accessibility.” *Mind*, vol. 118 (470): 399–409. [23]
- TELLER, PAUL. 1973. “Conditionalization and Observation.” *Synthese*, vol. 26 (2): 218–258. [6]
- . 1976. “Conditionalization, Observation, and Change of Preference.” In *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, W. L. HARPER & C. A. HOOKER, editors, vol. I, 205–253. D. Reidel Publishing Company, Dordrecht. [2], [6], [7], [8]
- TITELBAUM, MICHAEL G. 2013. *Quitting Certainties: A Bayesian Framework Modeling Degrees of Belief*. Oxford University Press, Oxford. [27]
- VAN FRAASSEN, BAS C. 1984. “Belief and the Will.” *The Journal of Philosophy*, vol. 81 (5): 235–256. [10]
- . 1989. *Laws and Symmetry*. Oxford University Press, Oxford. [27]

- . 1995. “Belief and the Problem of Ulysses and the Sirens.” *Philosophical Studies*, vol. 77: 7–37. [10]
- WEISBERG, JONATHAN. 2007. “Conditionalization, Reflection, and Self-Knowledge.” *Philosophical Studies*, vol. 135 (2): 179–97. [20]
- WILLIAMS, P. M. 1980. “Bayesian Conditionalization and the Principle of Minimum Information.” *The British Journal for the Philosophy of Science*, vol. 31 (2): 131–144. [27]
- WILLIAMSON, TIMOTHY. 2000. *Knowledge and its Limits*. Oxford University Press, Oxford. [12]
- . 2014. “Very Improbable Knowing.” *Erkenntnis*, vol. 79 (5): 971–999. [21], [23]