An Improved Dutch Book Theorem for Conditionalization

Michael Rescorla

Abstract: Lewis proved a Dutch book theorem for Conditionalization. The theorem shows that an agent who follows any credal update rule other than Conditionalization is vulnerable to bets that inflict a sure loss. Lewis's theorem is tailored to *factive* formulations of Conditionalization, i.e. formulations on which the conditioning proposition is true. Yet many scientific and philosophical applications of Bayesian decision theory require a *non-factive* formulation, i.e. a formulation on which the conditioning proposition may be false. I prove a Dutch book theorem tailored to non-factive Conditionalization. I also discuss the theorem's significance.

§1. Diachronic Dutch books

The diachronic norm *Conditionalization* is a central plank of Bayesian decision theory. Conditionalization requires that, in certain situations, you *conditionalize* on a proposition E by replacing your old credences P with new credences P_{new} given by:

$$P_{new}(H) = P(H \mid E).$$

Here $P(H \mid E)$ is the old conditional probability of H given E, defined by the *ratio formula*:

$$P(H \mid E) = \frac{P(H \,\&\, E)}{P(E)}$$

for cases where $P(E) > 0$. A *conditioning proposition* is a proposition on which you have conditionalized. The probability calculus axioms entail that $P(E \mid E) = 1$, so conditionalizing on E requires you to become *certain* of E, i.e. to set $P_{new}(E) = 1$.

Philosophers pursue various strategies for defending Conditionalization. One prominent strategy builds upon the classic *Dutch book arguments* advanced by Ramsey (1931) and de Finetti (1937/1980). A Dutch book is a set of acceptable bets that guarantees a net loss. You are *Dutch bookable* when it is possible to rig a Dutch book against you. If you are Dutch bookable, then a devious bookie can exploit you by offering you bets that you regard as fair. Ramsey and de Finetti observed that an agent whose credences violate the probability calculus axioms is Dutch bookable. They concluded that credences should conform to the probability calculus axioms. This is a *synchronic* Dutch book argument, because it addresses credences at a moment rather than credal evolution over time. Lewis (1999) extended Dutch book argumentation into the diachronic realm by proving a Dutch book theorem for Conditionalization. He showed how to rig a *diachronic Dutch book* (containing bets offered at different times) against an agent who follows any update rule other than Conditionalization.[1] Skyrms (1987b) proved a converse result: someone who conforms to the probability calculus axioms and to Conditionalization is not Dutch bookable in certain natural learning scenarios. Lewis's Dutch book theorem and Skyrms's converse theorem show that Conditionalization is the unique credal update rule immune to guaranteed net loss in the specified learning scenarios. Lewis, Skyrms, and many other authors conclude that agents should employ Conditionalization as their credal update rule.

In this paper, I will identify and rectify a significant imperfection in the standard diachronic Dutch book theorem for Conditionalization. As I discuss in Section 2, many applications of Bayesian decision theory feature an agent who becomes certain of a *false* proposition and revises her credences accordingly. To model such an agent, we require a *non-factive* formulation of Conditionalization, i.e. a formulation that allows conditioning propositions

[1] Lewis proved the theorem in 1972 but did not publish until (Lewis, 1999). Teller (1973) first published the theorem with Lewis's permission, crediting it to Lewis.

to be false. Unfortunately, the standard diachronic Dutch book theorem is tailored to *factive* formulations of Conditionalization, i.e. formulations that treat agents as conditionalizing on true propositions (Sections 3-4). For that reason, the standard theorem is non-optimal. Luckily, there is an easy remedy. In Section 5, I prove a modified diachronic Dutch book theorem tailored to non-factive Conditionalization. The modified theorem improves upon Lewis's original, because it pertains to a more general formulation of Conditionalization. The modified theorem brings Conditionalization's unique pragmatic virtues into sharper focus (Section 6).

To understand my goals, one must distinguish between Dutch book *theorems* and Dutch book *arguments*. Dutch book theorems are uncontroversial mathematical results. Dutch book arguments cite Dutch book theorems as evidence that it is irrational to violate some credal norm. Critics raise various objections to Dutch book arguments (Hájek, 2009), especially diachronic Dutch book arguments (Christensen, 1991). Whether or not Dutch book arguments succeed, I think that Dutch book theorems have great interest. Dutch bookability is an undesirable property. It is somehow bad to leave yourself open to Dutch books, somehow good to render yourself immune to Dutch books. We have learned important information when we prove a Dutch book theorem. Even if one rejects Dutch books arguments, time spent proving and analyzing Dutch book theorems is time well spent. I offer the present paper in that spirit. By studying Dutch book theorems for Conditionalization, I do not seek to establish that violating Conditionalization is irrational. Instead, I aim to clarify notable pragmatic benefits that differentiate Conditionalization from rival credal update rules.

§2. Factive versus non-factive formulations of Conditionalization

Let us begin by comparing two prominent formulations of Conditionalization. Earman's (1992, p. 34) formulation runs as follows:

> The rule of *strict conditionalization* says that if it is learned for sure that E and if E is the strongest such proposition, then the probability functions Pr_{old} and Pr_{new}, representing respectively degree of belief prior to and after acquisition of new knowledge, are related by $Pr_{new}(.) = Pr_{old}(./E)$.

Contrast Earman's formulation with Huber's (2016):

> If evidence comes only in the form of *certainties* (that is, propositions of which you become certain), if Pr: $\mathbf{A}\rightarrow \mathbb{R}$ is your subjective probability at time *t*, and if between *t* and *t′* you become certain of $A \in \mathbf{A}$ and no logically stronger proposition in the sense that your new subjective probability for *A*, but for no logically stronger proposition, is 1 (and your subjective probabilities are not directly affected in any other way such as forgetting etc.), then your subjective probability at time *t′* should be $Pr(\cdot | A)$,

where **A** is the set of relevant propositions and $\mathbb{R}$ is the real numbers. Both formulations say that, if you begin with initial credences *P*, then under certain circumstances your new credal allocation P_{new} should satisfy the constraint $P_{new}(.) = P(. | E)$. There are various differences between Earman's formulation and Huber's. The difference that interests me concerns the relation that you bear to the conditioning proposition *E*. For Earman, you "learn for sure that *E*" and thereby acquire "new knowledge." This formulation entails that *E* is true. For Huber, you become certain of *E*. You can become certain of a false proposition. Thus, Huber's formulation does not suggest that *E* is true. Earman's formulation of Conditionalization is factive, while Huber's is non-factive.

Surveying the literature on Conditionalization, one finds both factive and non-factive formulations. Factive formulations predominate (e.g. Arntzenius, 2003, p. 367, fn. 7; Briggs, 2009, p. 61; Easwaran, 2011, p. 314; Greaves and Wallace, 2006, p. 607; Hacking, 1967, p. 314; Howson and Franklin, 1994, p. 453; Vineberg, 2011; Weisberg, 2009, p. 499; Williamson, 2000, pp. 184-223), but there are also prominent non-factive formulations (e.g. Hájek, 2009; Jeffrey, 2014; Talbott, 2015; Titelbaum, 2013; van Fraassen, 1999). Some authors do not indicate whether they intend a factive or non-factive construal. Surprisingly, authors seldom note the contrast between factive and non-factive formulations or defend one approach over the other.

In my opinion, the non-factive approach offers a crucial advantage over the factive approach: it covers a wider range of important cases. The non-factive approach accommodates scenarios involving *misplaced certainty*, i.e. scenarios where you set $P_{new}(E) = 1$ even though E is false. Misplaced certainty arises in diverse scientifically and philosophically important contexts (Rescorla, forthcoming). For example, *Bayesian statisticians* routinely conditionalize on evidence that may or may not be true. A typical statistical model might dictate how to update a probability distribution over a random variable (e.g. the age of a fossil) based on the value of another random variable (e.g. a carbon dating measurement). Determining the second variable's value is usually a fallible undertaking. Scenarios featuring misplaced certainty figure routinely in most scientific fields that employ the Bayesian framework, including economics (Fudenberg and Tirole, 1991), cognitive science (Knill and Richards, 1996), robotics (Thrun, Burgard, and Fox, 2005), medicine (Ashby, 1996), and many others. To handle these scenarios, we must formulate Conditionalization non-factively.

The literature offers various subtly different formulations of non-factive Conditionalization. I give my own preferred formulation in (Rescorla, forthcoming). For present purposes, I abstract away from subtle differences among non-factive formulations. My discussion applies equally well to any reasonably careful non-factive formulation.

One might worry that I have drawn the wrong moral from the possibility of misplaced certainty. Rather than formulate Conditionalization non-factively, perhaps we should instead abandon Conditionalization. After all, conditionalizing on E requires setting $P_{new}(E) = 1$. But how can an agent reasonably set $P_{new}(E) = 1$ when E may well be false? Stretching back at least to Shimony (1955), many philosophers have maintained that every metaphysically possible proposition should receive positive credence. This normative constraint is sometimes called *Regularity*. An agent who conditionalizes on a contingent proposition E and who conforms to the probability calculus axioms will violate Regularity by setting $P_{new}(\neg E) = 0$.

Jeffrey (1983) develops an alternative approach that lets us honor Regularity. Say that $\boldsymbol{E}$ is a *partition* iff it is a countable set of mutually exclusive, jointly exhaustive propositions such that $0 < P(E) < 1$ for each $E \in \boldsymbol{E}$. Jeffrey considers scenarios where an agent reallocates credence across a partition; on that basis, the agent must reallocate credence across all other propositions. Jeffrey formulates a credal reallocation rule (now usually called *Jeffrey Conditionalization*) tailored to such scenarios. One might propose that we replace Conditionalization with Jeffrey Conditionalization as the sole basis for credal updates. By relying solely upon Jeffrey Conditionalization, we can update credences while cleaving to Regularity.

I reply that this proposal is highly revisionary regarding current scientific practice, in at least two ways:

- Scientific disciplines that employ the Bayesian framework (including statistics, economics, cognitive science, and robotics) typically use non-factive Conditionalization rather than Jeffrey Conditionalization. To replace non-factive

Conditionalization with Jeffrey Conditionalization would be to reject the main credal reallocation strategy employed in all these disciplines.

- Orthodox probability theory demands that numerous metaphysically possible propositions receive probability 0. Let *X* be a random variable with continuum many possible values (e.g. an object's velocity) and *x* a possible value of *X* (e.g. a specific velocity). It is metaphysically possible that $X = x$. Yet orthodox probability theory requires that $P(X = x) = 0$, for all but countably many *x* (Billingsley, 1995, p. 188).[2] Orthodox probability theory plays a foundational role in most scientific applications of the Bayesian framework. Regularity mandates massive revisions to all those applications.

It is far from clear that we could implement such sweeping changes while preserving the explanatory and pragmatic benefits offered by current scientific practice.

Scientific practice is not sacrosanct. Some readers may simply insist that here is one case where scientists should revise how they normally operate. However, I contend that this revisionary viewpoint does not obviate our need to consider non-factive Conditionalization.

To see why, suppose that a scientist sets $P_{new}(E) = 1$, where *E* may be either true or false. The revisionary viewpoint insists that the scientist should not have done so. That diagnosis may seem especially plausible when *E* is false (unbeknownst to the scientist). Let us grant for the moment that the diagnosis is correct: the scientist should not have set $P_{new}(E) = 1$. Even so, there are better and worse ways for her to revise her other credences in light of her newfound certainty in *E*. A single misstep does not entail that now anything goes. For example, it seems better for the scientist to set

[2] Skyrms (1980) recommends that we avoid this consequence by revising orthodox probability theory, allowing infinitesimals to serve as probabilities. For critical discussion, see (Easwaran, 2014) and (Weisberg, 2009).

$$P_{new}(H) = P(H \mid E)$$

than for her to set

$$P_{new}(H) = P(\neg H \mid E).$$

Which credal revisions are better and which are worse? Only by answering this question can we lay down useful prescriptions for scientific practice *as it currently stands*. Even if the scientist should not have become certain of *E*, we can still assess how well she reallocates her other credences in light of her faulty certainty. Even if we hope to revise scientific practice by banning certainties in contingent propositions, there is considerable interest in studying the norms that govern scientific practice *until the revisionary ban takes effect*. I submit that non-factive Conditionalization occupies a central place among those norms.[3]

Still, it may seem that the extra generality offered by non-factive Conditionalization over factive Conditionalization is not so helpful. Suppose an agent conditionalizes on *E*, where *E* is false. Assuming that the agent conforms to the probability calculus axioms, she sets

$$P_{new}(E \mid F) = 1,$$

for all propositions *F* such that $P_{new}(F) > 0$. She cannot dislodge her newfound certainty in *E* by conditionalizing on any such *F*. So it may appear that she is saddled with permanent certainty in *E*, no matter how much evidence against *E* she subsequently acquires. Permanent certainty in an evident falsehood is highly undesirable. Non-factive Conditionalization may therefore seem a fairly useless addition to the Bayesian repertoire.

---

[3] An anonymous referee suggests that, when a scientist conditionalizes on *E*, she engages in a kind of idealization: she takes *E* to be true for certain purposes, or in certain theoretical contexts, despite knowing that *E* may be false or even that *E is* false. However, classifying some aspect of scientific practice as an idealization does not exempt us from limning the norms that govern it. Scientific idealizations also fall under norms. Even if we grant that a scientist's newfound certainty in *E* is an idealizing assumption, there are still better and worse ways for her to adjust her other credences in light of her idealizing assumption. Codifying which ways are better and which ways are worse requires us to articulate a norm governing idealizing assumptions of this general kind. Articulating such a norm carries us back to non-factive Conditionalization or some norm in that vicinity (e.g. an otherwise similar norm that relativizes credal assignments to certain theoretical contexts).

I reply that it is quite possible for a rational agent to dislodge her certainty in a proposition. The literature offers several broadly Bayesian frameworks for modeling eradication of certainties. The framework I myself favor uses *regular conditional distributions* (rcds) to generalize the ratio formula, thereby delineating conditional probabilities $P(H \mid E)$ for numerous cases where $P(E) = 0$.[4] Rcds play a foundational role in orthodox probability theory (Billingsley, 1995). They also figure prominently in scientific practice, including within Bayesian statistics (Ghosal and van der Vaart, 2017). If an agent conditionalizes using rcds, then she can dislodge credences from 1 down to 0 or anywhere in between. Indeed, certainty eradication through conditionalization using rcds is a routine occurrence in Bayesian statistics (Ghosal and van der Vaart, 2017). Conditionalization using rcds raises many mathematical and philosophical complexities that lie beyond the scope of this paper. What matters for our purposes is just that agents can eradicate certainties by using rcds to update credences. In particular, the following progression is possible: the agent uses the ratio formula to conditionalize on E, as a result setting $P_{new}(E) = 1$; then the agent receives new evidence and conditionalizes using an rcd, as a result setting $P_{new}(E) < 1$.[5] Thus, conditionalizing on E does not forever condemn a Bayesian agent to certainty in E. Agents who obey non-factive Conditionalization need not fear permanent certainty in falsehoods.[6]

[4] (Rescorla, 2015) contains a philosophically oriented introduction to rcds. There are several alternative frameworks for analyzing conditional probability when $P(E) = 0$. Easwaran (2019) gives a balanced survey, from a perspective sympathetic to rcds. Note that the word “regular” in “regular conditional distribution” has nothing to do with the doctrine (Regularity) that metaphysically possible propositions should receive positive credence. This is an unfortunate case where the literature associates the same word, “regular,” with two completely different meanings.

[5] Titelbaum’s (2013) Certainty-Loss Framework can also model situations where the agent becomes certain of E but subsequently gains evidence that eradicates her certainty in E. As Titelbaum admits (2013, pp. 296-298), though, the Certainty-Loss Framework does not fully analyze key aspects of such situations.

[6] (Rescorla, 2018) proves a Dutch book theorem and converse Dutch book theorem for conditionalization using rcds. However, the proof assumes a factive setting. In future work, I will discuss rcds from a non-factive perspective. I will also analyze in more detail how conditionalization using rcds can eradicate certainties, including certainties gained by previously conditionalizing on a proposition E.

In what follows, I restrict attention to cases where $P(E) > 0$. These cases already raise enough issues for an entire paper. In Section 3, I review Lewis's diachronic Dutch book theorem. In Section 4, I explain why Lewis's theorem is tailored to factive rather than non-factive Conditionalization. In Section 5, I prove a modified Dutch book theorem tailored to non-factive Conditionalization.

§3. Lewis's diachronic Dutch book theorem

Lewis's (1999) diachronic Dutch book theorem concerns an idealized agent who interacts with a bookie. At time t_1, the agent has initial credences over some set of propositions. Her credences are given by a function P that conforms to the probability calculus axioms. Let $\boldsymbol{E}$ be a partition. Call the members of $\boldsymbol{E}$ *partition propositions*. At time t_2, the agent learns which partition proposition E is true. More precisely, she sets $P_{new}(E) = 1$ for some $E \in \boldsymbol{E}$, and it is guaranteed that

(1) $P_{new}(E) = 1$ iff E is true.

She has determinate dispositions to reallocate credences upon learning that E is true. We model these dispositions through a credal update rule C that maps each partition proposition E to a corresponding credal allocation C_E. $C_E(H)$ is the credence assigned to H upon learning that E is true. Hence,

$$P_{new}(H) = C_E(H)$$

if partition proposition E is revealed. Assume that C_E conforms to the probability calculus axioms. Call the foregoing scenario a *Lewis learning scenario*.[7]

[7] Lewis's exposition has several idiosyncratic features that it would be distracting for us to consider in depth. For example, he assumes that partition propositions "specify, in full detail, all the alternative courses of experience [the agent] might undergo between time 0 and time 1" (1999, p. 405). I track Skyrms's (1987b) exposition rather than Lewis's. But the core ideas are due to Lewis.

The bookie proposes bets at t_1 and t_2 concerning propositions in the domain of the agent's credence function. The bookie may choose to offer the null bet, which yields net payoff 0 in all outcomes. Multiple bets offered at a given time can be combined into a single bet, so it is legitimate to assume that the bookie offers a single bet at t_1 and a single bet at t_2. The bookie learns at time t_2 which proposition E is true. He has a strategy for deciding which bet to offer at t_2, depending on which proposition E is revealed. To model this setup, say that a *bookie strategy* is a pair (B, Γ), where B is a bet and Γ is a function that carries each partition proposition E into a bet $\Gamma(E)$. Intuitively, B is the bet offered at t_1, while $\Gamma(E)$ is the bet offered at t_2 if E is revealed. Since the bookie learns the one true partition proposition:

(2) The bookie offers $\Gamma(E)$ if E is true.

Our agent accepts or rejects a bet depending on whether it is *acceptable*, where a bet is acceptable iff it has non-negative expected net payoff. At t_1, she calculates expected net payoffs relative to her original credences P. At t_2, she calculates expected net payoffs relative to her new credences P_{new}. A bet is *acceptable relative to P* (*relative to C_E*) iff it has non-negative expected net payoff when expectations are calculated relative to P (relative to C_E). A bookie strategy (B, Γ) is a *Dutch book for update rule C* iff:

- Bet B is acceptable relative to P.
- For each partition proposition E, bet $\Gamma(E)$ is acceptable relative to C_E.
- For each partition proposition E, net payoff from B and $\Gamma(E)$ is negative in all outcomes. (Here net payoffs are computed under the assumption that bet $\Gamma(E)$ is enacted if E is true.)

A Dutch book for C inflicts a guaranteed net loss upon an agent who obeys C in a Lewis learning scenario. These definitions can be formalized more rigorously, but the present level of rigor suffices for this paper. A strengthened version of Lewis's theorem runs as follows:

Diachronic Dutch Book Theorem: *If there is a partition proposition E such that* $C_E(\,.\,) \neq P(\,.\mid E)$, *then there exists a Dutch book for C.*

Skyrms (1987b) proves a converse theorem:

Converse Diachronic Dutch Book Theorem: *If* $C_E(\,.\,) = P(\,.\mid E)$ *for all partition propositions E, then there does not exist a Dutch book for C.*

These two theorems establish Conditionalization as the unique update rule that immunizes you from a sure loss in Lewis learning scenarios.

Proof of the Diachronic Dutch Book Theorem: Suppose that

$$C_E(H) \neq P(H \mid E)$$

for some propositions E and H. Consider first the case where

$$C_E(H) < P(H \mid E).$$

At t_1, the bookie offers the following complex conditional bet: if E is false, then no money changes hands; if E is true, then the agent pays $P(H \mid E)$, and she receives a payoff from the bookie according to whether H is also true:

If H is true, the bookie pays 1.

If $\neg H$ is true, the bookie pays 0.

Table 1 summarizes net payoffs for this complex conditional bet. It is easy to show that the bet has expected net payoff zero relative to P and hence is acceptable relative to P. At t_2, the agent and the bookie both learn which partition proposition is true. The bookie's strategy is to offer the null bet at t_2 unless E is true. If E is true, then the bookie will ask the agent to sell for price $C_E(H)$ a new bet that pays off as follows:

If H is true, the agent pays 1.

If $\neg H$ is true, the agent pays 0.

Table 2 summarizes net payoffs for this new bet. The bet has expected net payoff zero relative to C_E and hence is acceptable relative to C_E.

INSERT TABLES 1-2 ABOUT HERE

Table 3 gives net payoffs for the overall gambling scenario as described thus far. Net payoff is $C_E(H) - P(H \mid E) < 0$ when E is true and 0 otherwise. A Dutch book as defined above requires a loss in *all* outcomes. We can fix this by adding a sidebet at t_1. The agent buys the sidebet for price $[P(H \mid E) - C_E(H)]P(E)$, and she receives a payoff from the bookie according to whether E is true:

If E is true, then the bookie pays $P(H \mid E) - C_E(H)$.

If $\neg E$ is true, then the bookie pays 0.

Table 4 summarizes net payoffs for the sidebet. The sidebet is easily shown to have expected net payoff zero relative to P. With the sidebet in place, net payoff for the entire gambling scenario is $P(E)[C_E(H) - P(H \mid E)] < 0$ in all outcomes. Thus, the overall gambling scenario inflicts a guaranteed net loss. For the case where $C_E(H) > P(H \mid E)$, the bookie can inflict a guaranteed net loss by multiplying all payoffs from Table 1-4 by –1. □

INSERT TABLES 3-4 ABOUT HERE

A bet is *favorable relative to P* (*relative to C_E*) iff it has positive expected net payoff relative to P (relative to C_E). Skyrms (1992) suggests that a genuine Dutch book should contain bets that are favorable, not just acceptable. One can strengthen Lewis's theorem to accommodate Skyrms's viewpoint. We need merely add a "sweetener" $-P(E)[C_E(H) - P(H \mid E)]/100$ to the agent's payoff for each bet and each possible outcome. The sweetener ensures that each bet is favorable. Net payoff from the sweetened book is negative no matter the outcome.

§4. Mismatch between theorem and norm

The proof of Lewis's theorem assigns a crucial role to the factivity assumptions (1) and (2). It assumes that at t_2 the non-conditionalizer become certain of the *true* partition proposition. This leads her to adopt new credences, which she uses to compute expectations at t_2. If the true partition proposition is E, then she uses C_E to compute expectations. She therefore deems the bet from Table 2 acceptable. The bookie also learns the true partition proposition at t_2, and he offers the bet from Table 2 precisely when E is true. As a result,

(3) The bet from Table 2 is enacted iff E is true.

This reasoning underlies the computation of net payoffs in the proof of Lewis's theorem. Absent (3), net payoff differs from $P(E)[C_E(H) - P(H \mid E)]$ in some outcomes.

More specifically, consider what happens when H & E is true but both gamblers become certain at t_2 of a different partition proposition E^*. Since the bookie does not believe that E is true, and since his goal is to offer the bet from Table 2 precisely when E is true, he does not offer the bet from Table 2. Net payoff is determined by the bets from Tables 1 and 4:

$$[1 - P(H \mid E)] + [P(H \mid E) - C_E(H)][1 - P(E)].$$

The first summand $[1 - P(H \mid E)]$ is nonnegative. The second summand $[P(H \mid E) - C_E(H)][1 - P(E)]$ is the product of two positive numbers. Thus, net payoff is positive. In short: the bookie does not offer any bet at t_2, so the non-conditionalizer's winnings from the bets enacted at t_1 are not cancelled out. What this shows is that Lewis's Dutch book does not inflict a sure loss in scenarios where participants may acquire misplaced certainty in a partition proposition. Lewis's proof goes through only because the theorem restricts attention to Lewis learning scenarios, i.e. scenarios where misplaced certainty in partition propositions cannot arise. The theorem and its converse may establish that Conditionalization offers distinctive pragmatic benefits *when participants are guaranteed to become certain of the true partition proposition*, but the theorems leave open whether Conditionalization offers distinctive pragmatic benefits *when participants can acquire misplaced certainty in a partition proposition*.[8]

If we formulate Conditionalization factively, then perhaps there is a principled reason to restrict attention to Lewis learning scenarios. Factive Conditionalization only applies to situations where the conditioning proposition is true. This restricted applicability arguably supplies a rationale for only considering Lewis learning scenarios when evaluating any distinctive pragmatic advantages afforded by factive Conditionalization. Once we formulate

[8] Lewis's proof also depends on the assumption that the agent follows a deterministic update rule. The proof does not apply to an agent who violates Conditionalization by employing a *stochastic* credal reallocation strategy (e.g. an agent who randomly selects a new credal allocation at t_2). I will follow Lewis by restricting attention to deterministic rather than stochastic credal reallocation strategies.

Conditionalization non-factively, though, the putative rationale evaporates. Non-factive Conditionalization applies to situations where the agent acquires misplaced certainty in a partition proposition. Situations of that kind arise routinely in applications of Bayesian decision theory. We should consider them for a more complete assessment of the pragmatic benefits provided by non-factive Conditionalization versus rival norms. I conclude that Lewis's theorem does not optimally specify the distinctive pragmatic benefits that attach to Conditionalization. There is a fundamental mismatch between Lewis's theorem and non-factive Conditionalization.

Some readers may deny that there is any need to consider the pragmatics of non-factive Conditionalization once we have proved a Dutch book theorem and converse theorem for factive Conditionalization. How are you supposed to determine whether proposition E is true, beyond evaluating whether it deserves credence 1? If you have some reason for suspecting that E is *not* true, then shouldn't you assign it credence below 1? All evidence regarding E should already be incorporated into your new credal allocation. Surely, then, you should respond to your new credence $P_{new}(E) = 1$ by updating credences in the same way whether or not E is true. One might conclude that a Dutch book theorem tailored to non-factive Conditionalization is superfluous, since anyone who conforms to factive Conditionalization should also conform to non-factive Conditionalization.

Let us grant that anyone who conforms to factive Conditionalization should also conform to non-factive Conditionalization. We still want to elucidate the distinctive pragmatic benefits offered by non-factive Conditionalization. Doing so requires a Dutch book theorem not specifically tailored to Lewis learning scenarios. We must investigate learning scenarios where the conditioning proposition is allowed to be false.

§5. Generalizing the Dutch book theorem and its converse

A *generalized Lewis learning scenario* proceeds as follows. An agent's initial credences at t_1 are given by a function P conforming to the probability calculus axioms. $\boldsymbol{E}$ is once again a partition. At t_2, the agent assigns $P_{new}(E) = 1$ for some $E \in \boldsymbol{E}$, where E may or may not be true. For any partition proposition E, she has determinate dispositions to reallocate credence in light of setting $P_{new}(E) = 1$. A credal update rule C codifies these dispositions. C_E is the credal allocation that results when the agent becomes certain of partition proposition E. Assume that C_E conforms to the probability calculus axioms. Assume also that

$$C_E(E) = 1,$$

since C_E is the credal allocation induced by her new certainty in E. In the context of Lewis's theorem, C codifies how the agent would reallocate credence upon learning the true partition proposition with complete certainty. In the present context, C codifies how she would reallocate credence upon acquiring certainty --- possibly misplaced certainty --- in a partition proposition.

Having broadened attention to generalized Lewis learning scenarios, we must reconsider our treatment of gambling. We still assume that the bookie has a strategy Γ for deciding what bet to offer at t_2, depending on the partition proposition of which he becomes certain. However, it no longer seems appropriate to assume (2). After all, if the bookie learns the true partition proposition but the agent has no such luck, then the bookie's ability to exploit his superior knowledge for sure gain hardly suggests that the agent's credal update rule is problematic. It now seems appropriate to assume merely that agent and bookie both become certain of the same partition proposition E at t_2, where E may be false. Under this revised conception, a gambling interaction proceeds much as under Lewis's conception. On both conceptions, there is a partition proposition that agent and bookie come to invest with complete certainty. The key difference is that under our revised conception the partition proposition may be false. As with Lewis's theorem, we can model the gambling setup through a bookie strategy (B, Γ). In the context of Lewis's theorem, $\Gamma(E)$ is the bet the bookie will offer at t_2 upon learning the true partition proposition E with complete certainty. In the present context, $\Gamma(E)$ is the bet the bookie will offer at t_2 upon becoming certain of E (where E may or may not be true).

Payoffs are to be computed in terms of the true outcome, which may differ from the agent's certain opinion at t_2. To illustrate, suppose the agent and bookie enact a bet at t_1 that yields a net gain if E is true and a net loss otherwise. Suppose the agent and bookie both become certain at t_2 of E, even though E is false. The bet yields a net loss *but the agent and bookie are certain at* t_2 *that it yields a net gain*. How, then, can they compute the correct payoff from the bet? The answer is that they can compute the correct payoff only if they eradicate their certainty in E. As noted in Section 2, the literature offers several frameworks for modeling certainty eradication. We could use one of these frameworks to model how the gamblers transition away from certainty that E is true. For present purposes, we may proceed informally. We may simply stipulate that at some future time t_3 the gamblers learn the truth-value of each proposition relevant to their bets and revise their other credences accordingly. Under this stipulation, the gamblers can compute true payoffs at t_3 for all bets they have enacted.

Skyrms (1987b), in his treatment of diachronic Dutch book, postulates an oracle whom the gamblers trust completely and who reveals the truth about all relevant propositions at t_3. This is a vivid way of imagining how the gamblers might compute true payoffs. However, there are obviously many different ways besides an oracle that the agent and bookie may gain new evidence at t_3. Accordingly, I do not assume any oracle. I assume only that at t_3 the gamblers *in some way or other* become certain of all true propositions relevant to computing net payoffs.

A similar assumption is needed even if we restrict attention to Lewis learning scenarios *simpliciter*, rather than generalized Lewis learning scenarios. An agent in a Lewis learning scenario *simpliciter* may become certain at t_2 of a false proposition H: the agent may set $C_E(H) = 1$ even though H is false and E is true. To compute the correct payoff for a bet on H, the agent must eventually learn that H is false. So payoff computation in Lewis learning scenarios also sometimes requires certainty eradication. Hence, retrenching from generalized Lewis learning scenarios to Lewis learning scenarios *simpliciter* would not eliminate the need to postulate certainty eradication. In either setup, cases arise where correct payoff computation requires eradicating certainties gained at t_2.

Our new setting requires a slightly different formal notion of Dutch book. Say that a bookie strategy (B, Γ) is a *non-factive Dutch book for update rule C* iff

- Bet B is acceptable relative to P.
- For each partition proposition E, bet $\Gamma(E)$ is acceptable relative to C_E.
- For each partition proposition E, net payoff from B and $\Gamma(E)$ is negative in all outcomes. (Here net payoffs are *not* computed under the assumption that bet $\Gamma(E)$ is enacted if E is true.)

The difference between non-factive Dutch books as just defined and Dutch books as defined in Section 3 and lies in the third clause: we drop the assumption that bet $\Gamma(E)$ is enacted if E is true. The reason why we drop this assumption is that we have dropped the factivity assumptions (1) and (2). If agent and bookie acquire misplaced certainty in partition proposition E^*, then the bookie will offer bet $\Gamma(E^*)$ even if some different partition proposition E is true. See Figure 1. A non-factive Dutch book for C inflicts a guaranteed net loss upon someone who follows C in a generalized Lewis learning scenario: she loses money no matter the partition proposition of which she becomes certain and no matter the true outcome. She loses money no matter how the gambling scenario unfolds.

INSERT FIGURE 1 ABOUT HERE

Lewis's book from Section 3 is a Dutch book *simpliciter*, but it is not a non-factive Dutch book. In terms of Figure 1: Lewis's book yields a net loss on all diagonal cells, but it does not yield a net loss on all cells. Once we recognize the possibility of misplaced certainty, we see that Lewis's theorem does not adequately elucidate Conditionalization's distinctive pragmatic benefits. We require a new Dutch book theorem that targets non-factive Dutch books.

Non-factive Diachronic Dutch Book Theorem: *If there is a partition proposition E such that* $C_E(\,.\,) \neq P(\,.\mid E)$, *then there exists a non-factive Dutch book for C.*

Proof**:** Suppose that

$$C_E(H) \neq P(H \mid E),$$

for some propositions E and H. We again consider the case where

$$C_E(H) < P(H \mid E).$$

The case where $C_E(H) > P(H \mid E)$ can be handled by multiplying all payoffs described below by –1. At t_1, the bookie offers the bet given by Table 1. This bet is acceptable relative to P, as noted in the proof of the diachronic Dutch book theorem. At t_2, the agent newly assigns credence 1 to some partition proposition E^*, so her new credences are given by C_{E^*}. The bookie's strategy is to ignore E^* and propose a new complex conditional bet: if E is false, then no money changes hands; if E is true, then the bookie pays $C_E(H)$, and the agent pays according to whether H is true:

If H is true, the agent pays 1.

If $\neg H$ is true, the agent pays 0.

Table 5 summarizes net payoffs for this bet. Let us now confirm that the bet is acceptable relative to C_{E^*}, for all $E^* \in \boldsymbol{E}$.

INSERT TABLE 5 ABOUT HERE

Either $E = E^*$ or $E \neq E^*$. Suppose first that $E = E^*$. Since $C_E(E) = 1$, and since C_E satisfies the probability calculus axioms,

$C_E(\neg E) = 0$

$C_E(H \,\&\, \neg E) = 0$

$C_E(H \,\&\, E) = C_E(H)$.

The bet's expected net payoff is therefore

$$C_{E^*}(H \,\&\, E)\mathit{Payoff}\,(H \,\&\, E) + C_{E^*}(\neg H \,\&\, E)\mathit{Payoff}\,(\neg H \,\&\, E) + C_{E^*}(\neg E)\mathit{Payoff}\,(\neg E)$$

$$= C_E(H \,\&\, E)[C_E(H) - 1] + C_E(\neg H \,\&\, E)C_E(H) + 0$$

$$= C_E(H)[C_E(H \,\&\, E) + C_E(\neg H \,\&\, E)] - C_E(H \,\&\, E)$$

$$= C_E(H)C_E(E) - C_E(H)$$

$$= C_E(H) - C_E(H) = 0\,.$$

On the other hand, suppose that $E \neq E^*$. Since $C_{E^*}(E^*) = 1$, and since C_{E^*} satisfies the probability calculus axioms,

$$C_{E^*}(E) = 0$$

$$C_{E^*}(H \,\&\, E) = 0$$

$$C_{E^*}(\neg H \,\&\, E) = 0$$

so that the bet's expected net payoff is

$$C_{E^*}(H \,\&\, E) Payoff\,(H \,\&\, E) + C_{E^*}(\neg H \,\&\, E) Payoff\,(\neg H \,\&\, E) + C_{E^*}(\neg E) Payoff\,(\neg E)$$

$$= 0 + 0 + 0.$$

Thus, the bet is acceptable relative to C_{E^*} whether or not $E \neq E^*$. Table 6 gives net payoffs for the entire gambling scenario. Net payoff is $C_E(H) – P(H \mid E) < 0$ when E is true and 0 otherwise. To inflict a guaranteed net loss in all outcomes, the bookie can add the sidebet at t_1 given by Table 4. For the supplemented gambling scenario, net payoff is $P(E)[C_E(H) – P(H \mid E)] < 0$ no matter the outcome. □

INSERT TABLE 6 ABOUT HERE

The key difference from Lewis's proof lies in the bookie strategy pursued at t_2. The new proof employs a different bookie strategy, one that requires no factivity assumptions to secure a net loss. As with Lewis's original theorem, we can strengthen the theorem by sweetening all bets, so that each individual bet becomes favorable and not just acceptable.

Converse Non-factive Diachronic Dutch Book Theorem: *If C_E (.) = P(. | E) for all partition propositions E, then there does not exist a non-factive Dutch book for C.*

Proof**:** Suppose that C_E (.) = P(. | E) for all partition propositions E. Suppose for reductio that there exists a non-factive Dutch book (B, Γ) for C. Thus:

- Bet B is acceptable relative to P.
- For each partition proposition E, the bet Γ(E) is acceptable relative to C_E.
- For each partition proposition E, net payoff from B and Γ(E) is negative in all outcomes.

For the third clause, net payoffs are not computed under the assumption that bet Γ(E) is enacted if E is true. If net payoffs are negative even without that assumption, then they are certainly negative under the assumption. (In terms of Figure 1: if (B, Γ) yields negative net payoff in all cells, then it certainly yields negative net payoff along the diagonal.) It follows that (B, Γ) is a Dutch book for C as defined in Section 3. By the converse diachronic Dutch book theorem, there exists no such Dutch book. By reductio, there does not exist a non-factive Dutch book for C. □

§6. Significance of the two theorems

From a technical perspective, Section 5's theorems are fairly trivial variants on theorems proved by Lewis and Skyrms. The new theorems are significant for philosophical rather than mathematical reasons. They show that non-factive Conditionalization offers distinctive benefits over rival update rules one might follow in generalized Lewis learning scenarios.

Consider an agent who accepts all acceptable bets. The non-factive diachronic Dutch book theorem shows that, if she employs any update rule other than Conditionalization in a generalized Lewis learning scenario, then a devious bookie with no superior knowledge can inflict a sure loss upon her. The converse theorem shows that, if the agent instead employs Conditionalization as her update rule, then she is immune to such exploitation. Just as factive Conditionalization is the unique update rule that immunizes the agent from a sure loss in Lewis learning scenarios, non-factive Conditionalization is the unique update rule that immunizes her from a sure loss in generalized Lewis learning scenarios.

My stipulation that the agent accepts all acceptable bets may seem questionable. If we construe payoffs as monetary amounts, then rationality does not demand that an agent accept every bet whose expected net payoff is non-negative. For example, an agent may rationally reject certain such bets if money has diminishing marginal utility for her (as it does for most people).

Following Armendt (1993) and Skyrms (1987a), I respond that we may construe payoffs as *utilities* rather than monetary amounts. An agent who maximizes expected utility should be willing to accept the bets comprising the non-factive diachronic Dutch book, with payoffs construed as utilities. If we consider sweetened bets that are favorable rather than just acceptable, then accepting the sweetened bets is the *unique* choice that maximizes expected utility. Thus, Section 5 establishes the following contrast: an expected utility maximizer is vulnerable to a sure loss of utility in generalized Lewis learning scenarios if she follows any update rule besides Conditionalization; she is immune to this prospect if she instead employs Conditionalization as her update rule. The contrast strikes me as a significant pragmatic advantage for non-factive Conditionalization. Losing utility is bad. Therefore, it is bad to find yourself in a situation where you are *guaranteed* to lose utility. Non-factive Conditionalization is the only update rule through which an expected utility maximizer can insulate herself from that predicament.[9]

Some readers may object that expected utility maximization is a non-mandatory decision-making rule. The literature offers numerous arguments that people do not (Kahneman and Tversky, 1979) or need not (Buchak, 2013) maximize expected utility. If an agent declines to maximize expected utility, then why should she care about any advantages that Conditionalization offers to expected utility maximizers? This objection applies whether we consider Lewis learning scenarios (for factive Conditionalization) or generalized Lewis learning scenarios (for non-factive Conditionalization).

Luckily, my argument that Conditionalization offers pragmatic advantages does not assume that expected utility maximization is the unique rational decision-making rule. My argument assumes only that expected utility maximization is *one reasonable decision-making rule* that agents might follow. Given that expected utility maximization figures crucially across a range of disciplines, I feel comfortable leaving this assumption undefended. If credal update rule C_1 when combined with reasonable decision-making rule D offers pragmatic benefits over credal update rule C_2 when combined D, then there is at least one respect in which C_1 is pragmatically superior to C_2. Thus, Section 5 shows that there is at least one respect in which non-factive Conditionalization is pragmatically superior to rival credal update rules. Perhaps there are *other* respects in which Conditionalization is *inferior* to certain rival update rules. This paper does not aim for a comprehensive assessment of Conditionalization's merits and demerits. I aim only to establish one advantage uniquely offered by Conditionalization.

[9] Some authors claim that a Dutch bookable agent should see the losses coming and opt out from all bets, thereby avoiding any losses. Levi (1987) and Maher (1992) develop this viewpoint as applied to diachronic Dutch books. Skyrms (1993) responds convincingly to Levi and Maher.

Overall, then, non-factive Conditionalization offers pragmatic benefits in generalized Lewis scenarios comparable to the pragmatic benefits that factive Conditionalization offers in non-generalized Lewis learning scenarios. That non-factive Conditionalization offers distinctive pragmatic benefits was no doubt antecedently plausible, but it is reassuring when rigorous proof confirms the antecedently plausible.

The pragmatic benefits I have identified may seem rather marginal. It is hardly as if Dutch bookies lurk on every street corner. Implementing a diachronic Dutch book (whether Lewis's original Dutch book or my modified Dutch book) requires knowing the agent's credal update rule in advance. The bookie must then engineer a complex series of exploitive bets. Few bookies are so knowledgeable and resourceful. Why should an agent feel at all troubled by her vulnerability to such a remote threat? Why should a conditionalizer feel particularly relieved that she is immune to that same threat?

I agree that one is unlikely to encounter a Dutch bookie. However, prudence counsels that one guard against many unlikely dangers. If you own a home, then buying homeowner's insurance leaves you better off in one important respect than you were before, even though you are unlikely to file a claim on your insurance. There is a benefit to guarding against an unlikely outcome when the potential downside is sufficiently bad. And the potential downside from Dutch bookability is devastating: the Dutch bookie from Section 5 can enforce *arbitrarily* high losses, simply multiplying all payoffs by an arbitrary constant. Hence, there is an important respect in which someone immunized from sure loss in generalized Lewis scenarios is better off than someone vulnerable to sure loss in generalized Lewis learning scenarios, even though the chance of encountering a Dutch bookie is small. Just as there is a benefit to having homeowner's insurance, so is there a benefit to insulating oneself from a diachronic Dutch book.

The literature offers various additional arguments that Dutch bookability is not a pragmatic defect (Vineberg, 2011). I have not addressed all these arguments, so some readers may remain unconvinced that non-factive Conditionalization offers pragmatic advantages over alternative update rules. These readers will not accept everything I say in the paper, but they can still accept my main thesis: the non-factive diachronic Dutch book improves upon Lewis's original diachronic Dutch book. The non-factive Dutch book is an improvement because it works for scenarios that feature misplaced certainty. By allowing for misplaced certainty, the non-factive diachronic Dutch book theorem provides a fuller perspective on the worst case result that attaches to every possible update rule one might follow. Every update rule besides Conditionalization is vulnerable to a devastating worst case result in generalized Lewis learning scenarios: arbitrarily large loss of utility. Whether or not one agrees that this is a "pragmatic advantage" offered by Conditionalization, it is a striking difference in worst case results, and it is a difference far more sweeping than any isolated by the original Lewis-Skyrms results. Comparison of worst case results should inform any contrastive evaluation of update rules. Thus, Section 5's theorems should inform any comprehensive comparison of Conditionalization with rival update rules.

§6.1 Comparison with Skyrms

Skyrms proves some theorems in the same vicinity as Section 5's non-factive Dutch book theorem. Let us do a comparison so as to clarify what the present paper contributes.

Following Jeffrey (1983), Skyrms analyzes learning scenarios where probability mass shifts across a partition $\boldsymbol{E}$. The agent begins with credences P and subsequently adopts new credences over the partition propositions, which leads to new credences P_{new} for all other propositions. As discussed in Section 2, Jeffrey Conditionalization is one possible update rule for learning scenarios of this kind. Building on Armendt's (1980) discussion, Skyrms proves a Dutch book theorem and converse Dutch book theorem for Jeffrey Conditionalization. Generalized Lewis learning scenarios are a special case of Jeffrey-style learning scenarios: they are the special case where all probability mass shifts to a single partition proposition E^*, so that $P_{new}(E^*) = 1$ and $P_{new}(E) = 0$ for $E \neq E^*$.

By inspecting Skyrms's discussion (1987b, pp. 9-10), we can extract a Dutch book theorem for non-factive Conditionalization. However, the theorem employs a somewhat different notion of diachronic Dutch book than figures in Section 5. First, Skyrms posits a third time stage t_3 where an oracle reveals the true partition proposition E. Skyrms assumes that the agent will invest complete certainty in the oracle's pronouncements. In Skyrms's setup, a bookie strategy is a plan for offering bets at three times stages t_1, t_2, and t_3 (whereas a bookie strategy as defined in Section 5 only involves bets at t_1 and t_2). Second, and more importantly, the diachronic book showcased by Skyrms's theorem does not guarantee a net loss. Instead, it guarantees that the agent will never achieve a net gain and will *sometimes* suffer a net loss. Books with this property are called *semi-Dutch books*.[10]

Skyrms proves that (given reasonable assumptions) one can rig a diachronic semi-Dutch book against an agent who follows any update rule besides non-factive Conditionalization. Specifically, suppose there are propositions E^* and H such that the agent's credence in H at t_2 upon becoming certain of E^* differs from $P(H \mid E^*)$. Skyrms isolates a bookie strategy that never yields a net gain and that inflicts a net loss in scenarios where the agent become certain of E^* at t_2. If the agent's credences at t_2 are given by P_{new}, then the bookie strategy yields a net loss when

$$P_{new}(E^*) = 1$$

[10] Skyrms's discussion of Jeffrey Conditionalization assumes that the agent does not become certain of a partition proposition at t_2. Nevertheless, his proof readily generalizes to yield the results summarized in the next paragraph.

and net payoff 0 otherwise. See Figure 2.[11]

INSERT FIGURE 2 ABOUT HERE

This semi-Dutch book theorem is non-optimal. The reason is that semi-Dutch books are not so troubling.

Here it becomes important to distinguish between *semi-Dutch books* and what I will call *weak Dutch books*. A *semi-Dutch book* offers a possibility of net loss and no possibility of net gain. A *weak Dutch book* offers a positive probability of net loss and no positive probability of net gain. Semi-Dutch books are not necessarily weak Dutch books, because the probability of net loss from a semi-Dutch book may be 0. A weak Dutch book offers a *positive probability* of net loss, not just a *possibility* of net loss.

As Hájek (2009, pp. 188-189) urges, a possibility of net loss need not seem troubling if you are certain that the possibility will not occur. Focusing first on a synchronic semi-Dutch book, suppose you attach probability 0 to some remote metaphysical possibility H (e.g. that aliens have replaced your cat with an indistinguishable duplicate). Taking expected net payoffs as a guide, you should be willing to pay price 1 for a wager that returns payoff 1 when H is false and payoff 0 when H is true. *From your viewpoint,* the possibility that you will lose money on this wager is so remote that you assign the wager expected net payoff 0. The wager is acceptable *from your viewpoint*, even though it offers a possibility of net loss and no compensating possibility of net gain. The wager is a semi-Dutch book but not a weak Dutch book. *From your*

[11] I assume that the agent's credences at t_3 are given by $P(. \mid E)$, where E is the true partition proposition revealed by the oracle at t_3. (If the assumption is false, then the bookie can impose a Dutch book *simpliciter* using Lewis's strategy for times t_1 and t_3, bypassing t_2.) Under my assumption, the bookie can mount a semi-Dutch book by exploiting the disparity between the agent's credences at t_2 and t_3 in the case where $P_{new}(E^*) = 1$.

viewpoint, the wager imposes no serious risk of net loss. Perhaps you are worried by the remote possibility of net loss, perhaps not. If you are not worried, it does not follow that you violate any norm of rationality. Thus, semi-Dutch bookability does not in itself suggest that any serious pragmatic defect afflicts your credal assignments. Whether weak Dutch bookability reveals a pragmatic defect in your credal assignments is a trickier question that we may luckily set aside for purposes of this paper.[12]

These points generalize to diachronic semi-Dutch books, including the diachronic semi-Dutch book secured by Skyrms's theorem. That an agent will suffer a net loss if $P_{new}(E^*) = 1$ and otherwise achieve net payoff 0 is not *in itself* worrisome. The mere possibility of net loss need not worry the agent if she attaches credence 0 to that possibility. In particular, the prospect of net loss when $P_{new}(E^*) = 1$ need not disturb her *if she attaches credence 0 to the possibility that* $P_{new}(E^*) = 1$. We have assumed that she attaches non-zero initial credence to E^*, but it does not follow that she attaches non-zero initial credence to the proposition that $P_{new}(E^*) = 1$. For example, she may feel certain that an evil demon will intervene at time t_2 to ensure that she does not set $P_{new}(E^*) = 1$, even if E^* is in fact true. *From her viewpoint*, a semi-Dutch book that only inflicts net loss when $P_{new}(E^*) = 1$ need not be disturbing. Hence, the semi-Dutch Book theorem does not indicate the presence of any pragmatic defect.

Skyrms is attuned to these worries. He addresses them by making additional assumptions about second-order credences (1987b, pp. 11-14). Basically, he assumes that the non-conditionalizer initially sets $P(P_{new}(E^*) = 1) > 0$. Given Skyrms's second-order assumption, the semi-Dutch book secured by his diachronic semi-Dutch book theorem is a weak Dutch book. We can therefore add a sidebet at t_1 that yields net gain when $P_{new}(E^*) = 1$ and net loss otherwise.

---

[12] As this paragraph illustrates, one can rig a semi-Dutch book against any agent who violates Regularity. Shimony (1955) defends Regularity on that basis. I agree with Hájek (2009) that semi-Dutch bookability in itself is not worrisome and hence that Shimony's argument is not compelling.

(No such sidebet would be acceptable if the non-conditionalizer assigned credence 0 to the proposition that $P_{new}(E^*) = 1$.) When supplemented with a suitable sidebet, the weak Dutch book becomes a non-factive Dutch book in the spirit of Section 5 (except that bookie strategies now extend over the three time stages t_1, t_2, and t_3). In effect, then, Skyrms proves a second-order theorem along the following lines: an agent is vulnerable to a guaranteed net loss if she initially sets $P(P_{new}(E^*) = 1) > 0$, subsequently becomes certain of E^*, and responds to this newfound certainty by applying any update rule other than Conditionalization.[13]

This second-order theorem is potentially less helpful than Section 5's non-factive diachronic Dutch book theorem. Theorems should assume as little as possible. The second-order theorem requires an additional second-order assumption, so in that respect it is weaker than the non-factive diachronic Dutch book theorem. Specifically, suppose that an agent has second-order credence $P(P_{new}(E^*) = 1) = 0$ and employs an update rule other than Conditionalization. The second-order theorem leaves open whether she is Dutch bookable. So the second-order theorem leaves open whether any pragmatic defect afflicts her credal reallocation strategy. In contrast, the non-factive diachronic Dutch book theorem shows that her credal reallocation strategy shares a crucial pragmatic defect with the credal reallocation strategy of an otherwise similar agent who sets $P(P_{new}(E^*) = 1) > 0$.

Second-order credences obscure the true nature of the pragmatic advantages that accrue to Conditionalization. Whether formulated factively or non-factively, Conditionalization does not have a second-order character. It dictates the proper evolution of an agent's credences, where those credences may be either first-order or higher-order. Just as Lewis's Dutch book theorem for factive Conditionalization does not involve second-order credences, Section 5's non-factive

[13] Here one must allow *second-order bets*, i.e. bets over the agent's own credences at t_2. See (Skyrms, 1987b) for mathematical and philosophical details.

Dutch book theorem does not involve second-order credences. An agent who employs a credal update rule other than Conditionalization is exploitable for very general reasons quite independent of her initial second-order credences.

§6.2 Dutch book arguments for Conditionalization

I have not offered a Dutch book argument for Conditionalization. Nevertheless, many philosophers *do* offer such arguments. My discussion bears upon that enterprise.

In schematic form, diachronic Dutch book arguments for Conditionalization move from the following two premises:

(4) An agent who employs any credal update rule other than Conditionalization is vulnerable to a guaranteed net loss.

(5) A conditionalizer whose credences conform to the probability calculus axioms is immune to a guaranteed net loss.

to the conclusion:

(6) Rationality requires that agents employ Conditionalization as their credal update rule.

There are many ways one might challenge the argument from (4) and (5) to (6). I want to focus on premise (4). In what sense does Lewis's diachronic book *guarantee* a net loss? The book inflicts a sure loss *in all outcomes where the agent and the bookie learn the true partition proposition*. Yet there are possible outcomes where the agent and the bookie both become certain of a false partition proposition. As seen in Section 4, the book yields a net gain in some of those outcomes. Thus, Lewis's diachronic book does not in any straightforward sense inflict a sure loss. We can say that the book "guarantees" a net loss only if, when defining what it is to "guarantee" a net loss, we ignore all outcomes where the book yields a net gain.

Theorists who formulate Conditionalization factively may be justified in ignoring those outcomes. You violate factive Conditionalization only when you become newly certain of a true proposition. A restricted focus on outcomes where (1) and (2) prevail is therefore perhaps justified. However, I argued in Section 2 that we should formulate Conditionalization non-factively. Having chosen a non-factive formulation, there is no good reason to focus exclusively on scenarios where (1) and (2) prevail. The potential for misplaced certainty is latent within virtually all scientific and philosophical applications of the Bayesian framework. Situations where (1) and (2) fail can arise just as easily as situations where they prevail. We should consider such situations when assessing the relative merits of potential updates rules. Once we take "Conditionalization" in (6) to mean "non-factive Conditionalization," the Dutch book argument only looks remotely compelling if "guaranteed" in (4) means *guaranteed even when the agent and the bookie both become certain of a false partition proposition*. Lewis's diachronic book looks ill-suited to underwrite a compelling Dutch book argument for non-factive Conditionalization.

This problem is seldom if ever noted in the literature, perhaps because few authors explicitly distinguish between factive and non-factive formulations of Conditionalization. Admittedly, I have not conclusively forestalled all attempts at using Lewis's diachronic book to support non-factive Conditionalization. One might hope to mount a compelling argument that somehow surmounts the mismatch between theorem and norm. But a more straightforward remedy is to replace Lewis's theorem with a generalized theorem that allows misplaced certainty.

That is exactly what the non-factive diachronic Dutch book theorem accomplishes. The theorem shows that any update rule besides Conditionalization induces vulnerability to a sure loss *even in scenarios where misplaced certainty in partition propositions can arise*. By employing a more suitable notion of "guaranteed net loss," the theorem offers dialectical advantages over Lewis's theorem. It also offers dialectical advantages over Skyrms's second-order theorem. The second-order theorem assumes $P(P_{new}(E^*) = 1\ ) > 0$. A Dutch book argument for non-factive Conditionalization based on the second-order theorem seems destined to concede that you can rationally decline to conditionalize on E^* as long as you initially set $P(P_{new}(E^*) = 1\ ) = 0$. The non-factive diachronic Dutch book theorem avoids such worries by assuming nothing about second-order credences. I do not say that the non-factive diachronic Dutch book theorem subserves a compelling Dutch book argument. I say only that philosophers who hope to defend non-factive Conditionalization using a Dutch book argument would do well to invoke this theorem rather than Lewis's original or Skyrms's second-order variant.

In conversation, I have sometimes heard philosophers espouse a *non-strategic formulation* of the diachronic Dutch book argument. They claim that the argument does not depend in any essential way upon *strategies* for exploiting the non-conditionalizer. One could just as well consider a machine that randomly offers bets at t_2, rather than a bookie strategy for offering bets at t_2. In some situations, the machine will randomly offer the bet from Table 2, which combined with bets already offered at t_1 ensures a net loss. According to the non-strategic formulation, this possibility establishes irrationality. If the non-strategic formulation is correct, then a satisfying diachronic Dutch book argument does not require anything as strong as the non-factive diachronic Dutch book theorem or even Lewis's original theorem. The argument requires only that the non-conditionalizer have the following property:

(7) There are situations where the agent accepts bets at t_1 and bets at t_2 that collectively yield a net loss in all outcomes.

(7) is weaker than the consequent of Lewis's original theorem, let alone the non-factive diachronic Dutch book theorem.

I find the non-strategic formulation unconvincing. I doubt that (7) or anything in its vicinity suffices to establish irrationality. To see why, suppose an agent has credence $P(E) = .5$ at t_1 and then conditionalizes on $\neg E$ at t_2 so that $P_{new}(E) = 0$. We may stipulate that it was perfectly rational for her to do so (since all parties to the present dispute agree that Conditionalization in some form or other is a rational credal update norm). The agent will accept a bet at t_1 that yields a small profit if E is true and a small loss otherwise, along with a second bet at t_2 that yields a huge loss if E is true and payoff 0 otherwise. The two bets jointly guarantee a net loss whether or not E is true. So our agent satisfies (7): she has accepted bets at t_1 and t_2 that guarantee a net loss. Nevertheless, the agent proceeded rationally. Thus, (7) does not suffice to establish irrationality. One cannot establish irrationality simply by describing a scenario where the agent accepts bets at t_1 and t_2 that collectively ensure a net loss. No doubt that is why Lewis and Skryms, when advancing the diachronic Dutch book argument, took great pains to formulate it in strategic terms. Both Lewis and Skyrms emphasize that the argument hinges upon a *strategy* through which the bookie ensures a net loss no matter how the gambling scenario unfolds.

The problem highlighted by Section 4 is that existing versions of the argument do not actually specify such a strategy. Lewis's book ensures a net loss *only when the agent and the bookie learn the true partition proposition*. The non-factive diachronic Dutch book theorem fixes the problem by specifying a strategy that ensures a net loss *even when misplaced certainty in a partition proposition occurs*.

Even with this problem fixed, many objections face any philosopher who hopes to mount a compelling Dutch book argument for Conditionalization. For example:

- As mentioned earlier, the most natural way to construe “payoffs” from bets is as utilities rather than monetary amounts. But this construal raises the question of how we should understand utilities (Kaplan, 1996, p. 160). To what extent do Dutch book arguments depend upon contestable assumptions regarding utility? For discussion of the relation between Dutch books and utility theory, see (Armendt, 1993), (Kaplan, 1996, pp. 155-180), (Maher, 1993), and (Skryms, 1987a).[14]
- An increasingly popular objection to Dutch book arguments is that they misguidedly prioritize *pragmatic* over *epistemic* factors (Joyce, 1998; Kaplan, 1996, pp. 158-159). Even if we grant that violations of Conditionalization are *pragmatically defective*, why conclude that such violations are *epistemically irrational*? Why should we think that premises (4) and (5), no matter how we interpret them, support conclusion (6)?

A vast literature addresses these and many other objections to Dutch book arguments (Hájek, 2009), including some objections specifically tailored to diachronic Dutch book arguments (Christensen, 1991). I do not pretend to have done justice to this literature. I do claim that the non-factive diachronic Dutch book theorem fixes *one* significant problem with previous versions of the diachronic Dutch book argument for Conditionalization.[15]

[14] A common way to explicate utility is through a *representation theorem*. A typical representation theorem entails that, when an agent’s preferences satisfy certain constraints, we can represent her as having credences and utilities that determine her preferences via expected utility maximization. Ramsey (1931) proved the first such representation theorem and, on that basis, defended a version of Bayesianism. The ensuing literature offers additional representation theorems that improve upon Ramsey’s, often assigning the theorems a foundational role. For example, Jeffrey (1983) highlights a representation theorem proved by Bolker (1967), while Joyce (1999) proves a representation theorem for causal decision theory. Kaplan (1996, pp. 155-180) and Maher (1993, pp. 94-104) suggest that, given the apparent need to bring utilities into Dutch book arguments, we are better off abandoning Dutch book argumentation and appealing instead to a suitable representation theorem. However, representation theorems only apply to credences *at a moment of time*. Even if they yield a compelling justification for *synchronic* Bayesian norms, they do not seem well-suited to support a *diachronic* norm such as Conditionalization.

[15] van Fraassen (1984) gives a diachronic Dutch book argument for the *Principle of Reflection*. His argument, like Lewis’s original diachronic book argument for Conditionalization, assumes a factive setting. In future work, I will discuss how the diachronic Dutch book argument for Reflection fares once we reject factivity assumptions.

§6.3 Expected accuracy arguments for Conditionalization

In recent years, philosophers have grown increasingly focused upon more properly epistemic arguments for Bayesian norms. The goal is to justify Bayesian norms by arguing that they promote *accuracy*, where "accuracy" measures how much your credences differ from actual truth-values. Building on de Finetti's (1974) discussion, Joyce (1998) gives an accuracy-based argument that credences should conform to the probability calculus axioms. Easwaran (2013), Greaves and Wallace (2006), and Leitgeb and Pettigrew (2010a; 2010b) extend accuracy-based argumentation from the synchronic realm to the diachronic realm. They argue that Conditionalization has a distinctive epistemic virtue that privileges it over rival credal reallocation rules: it maximizes the *expected accuracy* of your future credences.

Schoenfield (2017) shows that expected accuracy arguments for Conditionalization rest upon a key factivity assumption. The arguments assume that you are certain you will learn the true partition proposition *E* at time t_2. That assumption looks highly suspect once we acknowledge the possibility of misplaced certainty in partition propositions. Schoenfield shows that, if we lift the assumption, expected accuracy maximization no longer favors Conditionalization. She argues that it instead favors a different norm:

(8) If you learn *E*, then you should conditionalize on the proposition that you learned *E*.

Schoenfield here construes "learn" in non-factive terms, so that you can learn *E* even if *E* is false. She remains neutral as to whether expected accuracy arguments are compelling. She takes no stand on whether Conditionalization versus (8) is the rational update rule. Whatever one thinks of (8), Schoenfield convincingly establishes that Conditionalization does not maximize expected accuracy when you recognize a non-negligible possibility of future misplaced certainty in partition propositions.

This situation marks a significant contrast with the Dutch book theorem for Conditionalization. As I have shown, the Lewis-Skyrms Dutch book results extend smoothly from Lewis learning scenarios to generalized Lewis learning scenarios. Conditionalization offers unique pragmatic benefits *whether or not we allow misplaced certainty in partition propositions*. It offers those benefits quite independently of whatever initial first-order or second-order credences you happen to have. By comparison, Conditionalization does not maximize expected accuracy *once you recognize the non-negligible possibility of misplaced certainty in partition propositions*. The Dutch book theorem for Conditionalization naturally extends from a factive to a non-factive setting, while the expected accuracy theorem for Conditionalization does not.[16]

What should we conclude from the contrast? Should we reappraise diachronic Dutch book arguments for Conditionalization as perhaps not so bad after all? Or should we insist that there is a principled reason to assume factivity when evaluating expected accuracy? Or should we say that diachronic Dutch book arguments are unconvincing and that Schoenfield's results support (8) over Conditionalization as the rational credal update rule? All these options, and many others besides, are worth considering. As the debate moves forward, the theorems proved in this paper can serve as fixed points that constrain the space of viable options.

§7. Going non-factive

All too often, philosophical discussion of Conditionalization proceeds in a factive vein. Many philosophers suggest that Conditionalization dictates how one should react upon becoming certain of a true proposition. This picture fits naturally with the widespread emphasis upon Lewis's Dutch book theorem, which analyzes learning scenarios where one gains certainty in a

[16] Briggs and Pettigrew (forthcoming) give an *accuracy-dominance* argument for Conditionalization. They prove that an agent who violates Conditionalization could have guaranteed an improved accuracy score by instead obeying Conditionalization. Their proof implicitly assumes a factive setting.

true proposition. I favor a more ecumenical picture. Most fundamentally, Conditionalization governs the evolution of credences in response to new certainties, where those new certainties may be misplaced. Conditionalization dictates how old credences should relate to new credences when the latter result from an exogenous change that instills newfound certainty in some proposition. I have promoted my favored approach by situating Conditionalization in the broader context of generalized Lewis learning scenarios and by proving an appropriate Dutch book theorem. Future inquiry into Conditionalization should abandon unhelpful factivity restrictions, adopting a more comprehensive approach that countenances misplaced certainty.

Acknowledgments

I am grateful to Brian Skyrms, Simon Huttegger, and Alan Hájek for helpful discussion of these issues. I also thank two anonymous referees for this journal for comments that improved the paper.

Works Cited

Armendt, B. 1980. "Is There a Dutch Book Argument for Probability Kinematics?". *Philosophy of Science* 47: pp. 583-588.
---. 1993. "Dutch Books, Additivity, and Utility Theory." *Philosophical Topics* 21: pp. 1-20.
Arntzenius, F. 2003. "Some Problems for Conditionalization and Reflection." *The Journal of Philosophy* 100: pp. 356-370.
Ashby, D. 2006. "Bayesian Statistics in Medicine: A 25 Year Review." *Statistics in Medicine* 25: pp. 3589-3631.
Billingsley, P. 1995. *Probability and Measure*. 3rd ed. New York: Wiley.
Bolker, E. 1967. "A Simultaneous Axiomatization of Utility and Subjective Probability." *Philosophy of Science* 34: pp. 333-340.
Briggs, R. 2009. "Distorted Reflection." *Philosophical Review* 118: 59-85.
Briggs, R., and Pettigrew, R. (Forthcoming). "An Accuracy-dominance Argument for Conditionalization." *Nous*.
Buchak, L. 2013. *Risk and Rationality*. Oxford: Oxford University Press.
de Finetti, B. 1937/1980. "Foresight. Its Logical Laws, Its Subjective Sources." Rpt. in *Studies in Subjective Probability*, eds. H. E. Kyburg, Jr. and H. E. Smokler. Huntington: Robert E.

Krieger.
---. 1974. *Theory of Probability,* vol. 1. New York: Wiley.
Christensen, David. 1991. "Clever Bookies and Coherent Beliefs." *Philosophical Review* 100: pp. 229–47.
Earman, J. 1992. *Bayes or Bust?* Cambridge: MIT Press.
Easwaran, K. 2011. "Bayesianism I: Introduction and Arguments in Favor." *Philosophy Compass* 6: pp. 312-320.
---. 2013. "Expected Accuracy Supports Conditionalization --- and Conglomerability and Reflection." *Philosophy of Science* 80: pp. 119-142.
---. 2014. "Regularity and Hyperreal Credences." *Philosophical Review* 123: pp. 1-41.
---. 2019. "Conditional Probabilities." *The Open Handbook of Formal Epistemology*, eds. R. Pettigrew and J. Weisberg. PhilPapers.
Fudenberg, D., and Tirole, J. 1991. *Game Theory*. Cambridge: MIT Press.
Ghosal, S., and van der Vaart, A. 2017. *Fundamentals of Nonparametric Bayesian Inference*. Cambridge: Cambridge University Press.
Greaves, H., and Wallace, D. 2006. "Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility." *Mind* 115: pp. 607-632.
Hacking, I. 1967. "Slightly More Realistic Personal Probability." *Philosophy of Science* 34: pp. 311-325.
Hájek, A. 2009. "Dutch Book Arguments." In *The Handbook of Rationality and Social Choice*, eds. P. Anand, P. Pattanaik, and C. Puppe. Oxford: Oxford University Press.
Howson, C., and Franklin, A. 1994. "Bayesian Conditionalization and Probability Kinematics." *The British Journal for the Philosophy of Science* 45: pp. 451-466.
Huber, F. 2016. "Formal Representations of Belief." *The Stanford Encyclopedia of Philosophy* (Spring 2016), ed. E. Zalta. http://plato.stanford.edu/archives/spr2016/entries/formal-belief/.
Jeffrey, R. 1983. *The Logic of Decision*. 2nd edition. Chicago: University of Chicago Press.
---. 2004. *Subjective Probability: The Real Thing*. Cambridge: Cambridge University Press.
Joyce, J. 1998. "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65: pp. 575-603.
---. 1999. *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
Kaplan, M. 1996. *Decision Theory as Philosophy*. Cambridge: Cambridge University Press.
Kahneman, D., and Tversky, A. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47: pp. 263-291.
Knill, D. and Richards, W., eds. 1996. *Perception as Bayesian Inference*. Cambridge: Cambridge University Press.
Leitgeb, H., and Pettigrew, R. 2010a. "An Objective Justification of Bayesianism I: Measuring Inaccuracy." *Philosophy of Science* 77: pp. 201-235.
---. 2010b. "An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy." *Philosophy of Science* 77: pp. 236-72.
Maher, P. 1992. "Diachronic Rationality." *Philosophy of Science* 59: pp. 120-141.
---. 1993. *Betting on Theories*. Cambridge: Cambridge University Press.
Levi, I. 1987. "The Demons of Decision." Monist 70: 193–211.
Lewis, D. 1999. "Why Conditionalize?". In *Papers in Metaphysics and Epistemology* Cambridge: Cambridge University Press.
Ramsey, F. P. 1931. "Truth and Probability." In *The Foundations of Mathematics and Other*

Logical Essays, ed. R. B. Braithwaite. London: Routledge and Kegan.
Rescorla, M. 2018. "A Dutch Book Theorem and Converse Dutch Book Theorem for Kolmogorov Conditionalization." *The Review of Symbolic Logic* 11: pp. 705-735.
---. Forthcoming. "On the Proper Formulation of Conditionalization." *Synthese*. https://doi.org/10.1007/s11229-019-02179-9.
Schoenfield, M. 2017. "Conditionalization Does Not (in general) Maximize Expected Accuracy." *Mind* 504: pp. 1155-1187.
Shimony, A. 1955. "Coherence and the Axioms of Confirmation." *Journal of Symbolic Logic* 20: pp. 1-28.
Skyrms, B. 1980. *Causal Necessity*. New Haven: Yale University Press.
---. 1987a. "Coherence." In *Scientific Inquiry in Philosophical Perspective*, ed. N. Rescher. Pittsburgh: University of Pittsburgh Press.
---. 1987b. "Dynamic Coherence and Probability Kinematics." *Philosophy of Science* 54: pp. 1–20.
---. 1992. "Coherence, Probability and Induction." *Philosophical Issues* 2: pp. 215-226.
---. 1993. "A Mistake in Dynamic Coherence Arguments?". *Philosophy of Science* 60: pp. 320-328.
Talbott, W. 2015. "Bayesian Epistemology." *The Stanford Encyclopedia of Philosophy* (Summer 2015), ed. E. Zalta. http://plato.stanford.edu/archives/sum2015/entries/epistemology-bayesian/.
Teller, P. 1973. "Conditionalization and Observation." *Synthese* 26: pp. 218–258.
Thrun, S., Burgard, W., and Fox, D. 2005. *Probabilistic Robotics*. Cambridge: MIT Press.
Titelbaum, M. 2013. *Quitting Certainties*. Oxford: Oxford University Press.
van Fraassen, B. 1984. "Belief and the Will." *Journal of Philosophy* 81: pp. 235–56.
---. 1999. "Conditionalization, A New Argument For." *Topoi* 18: pp. 93-96.
Vineberg, S. 2011. "Dutch Book Arguments." *The Stanford Encyclopedia of Philosophy* (Summer 2011), ed. E. Zalta. http://plato.stanford.edu/archives/sum2011/entries/dutch-book/.
Weisberg, J. 2009. "Varieties of Bayesianism." In *Handbook of the History of Logic*, vol. 10, eds. D. Gabbay, S. Hartman, and J. Woods. New York: Elsevier.
Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

Outcome	*Net payoff*
H & E	$1 - P(H \mid E)$
$\neg H$ & E	$- P(H \mid E)$
$\neg E$	0

Table 1. Net payoff for the bet offered at t_1.

Outcome	*Net payoff*
H	$C_E(H) - 1$
$\neg H$	$C_E(H)$

Table 2. Net payoff for the bet offered at t_2 when partition proposition E is revealed.

Outcome	*Net payoff*
H & E	$C_E(H) - P(H \mid E)$
$\neg H$ & E	$C_E(H) - P(H \mid E)$
$\neg E$	0

Table 3. Net payoff for the bets from Tables 1 and 2, computed under the assumption that the second bet is enacted iff E is true.

Outcome	*Net payoff*
E	$[P(H \mid E) - C_E(H)][1 - P(E)]$
$\neg E$	$- [P(H \mid E) - C_E(H)]P(E)$

Table 4. Sidebet on E offered at t_1.

Outcome	*Net payoff*
H & E	$C_E(H) - 1$
$\neg H$ & E	$C_E(H)$
$\neg E$	0

Table 5. Net payoff for the bet offered at t_2 in the proof of the non-factive Dutch book theorem.

Outcome	***Net payoff***
H & E	$C_E(H) - P(H \mid E)$
$\neg H$ & E	$C_E(H) - P(H \mid E)$
$\neg E$	0

Table 6. Net payoff for the bets from Tables 1 and 5. Although the entries on this table are the same as for Table 3, the underlying computations are different. For Table 6, we assume that the bet from Table 5 is offered no matter the partition proposition E^* of which the non-conditionalizer becomes certain.

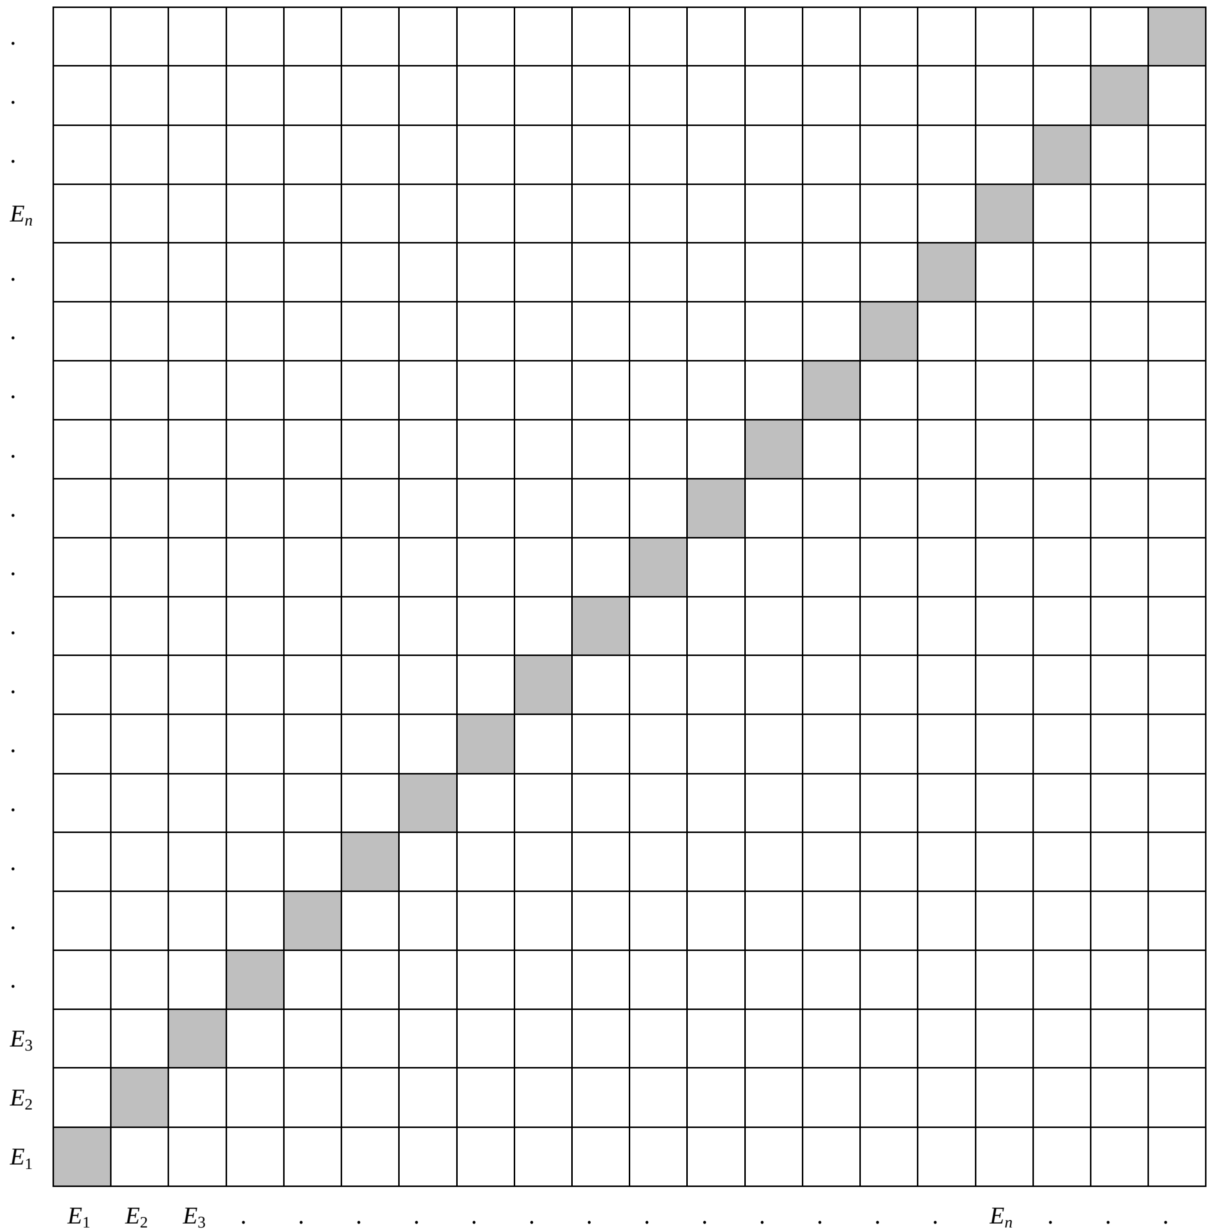

Figure 1. Assume some enumeration E_1, E_2, …, E_n, … of $\boldsymbol{\mathsf{E}}$. The horizontal axis depicts the partition proposition of which the agent becomes certain at t_2. Thus, the horizontal axis determines which bet is selected at t_2. The vertical axis depicts the true partition proposition. In a Lewis learning scenario, only cells along the diagonal are possible: agent and bookie both become certain at t_2 of the true partition proposition. In a generalized Lewis learning scenario, all cells are possible. A non-factive Dutch book ensures a net loss at every cell. A Dutch book *simpliciter* only ensures a net loss along the diagonal. (The partition $\boldsymbol{\mathsf{E}}$ may be countably infinite, in which case Figure 1 only depicts a finite portion of the infinitely many possible outcomes.)

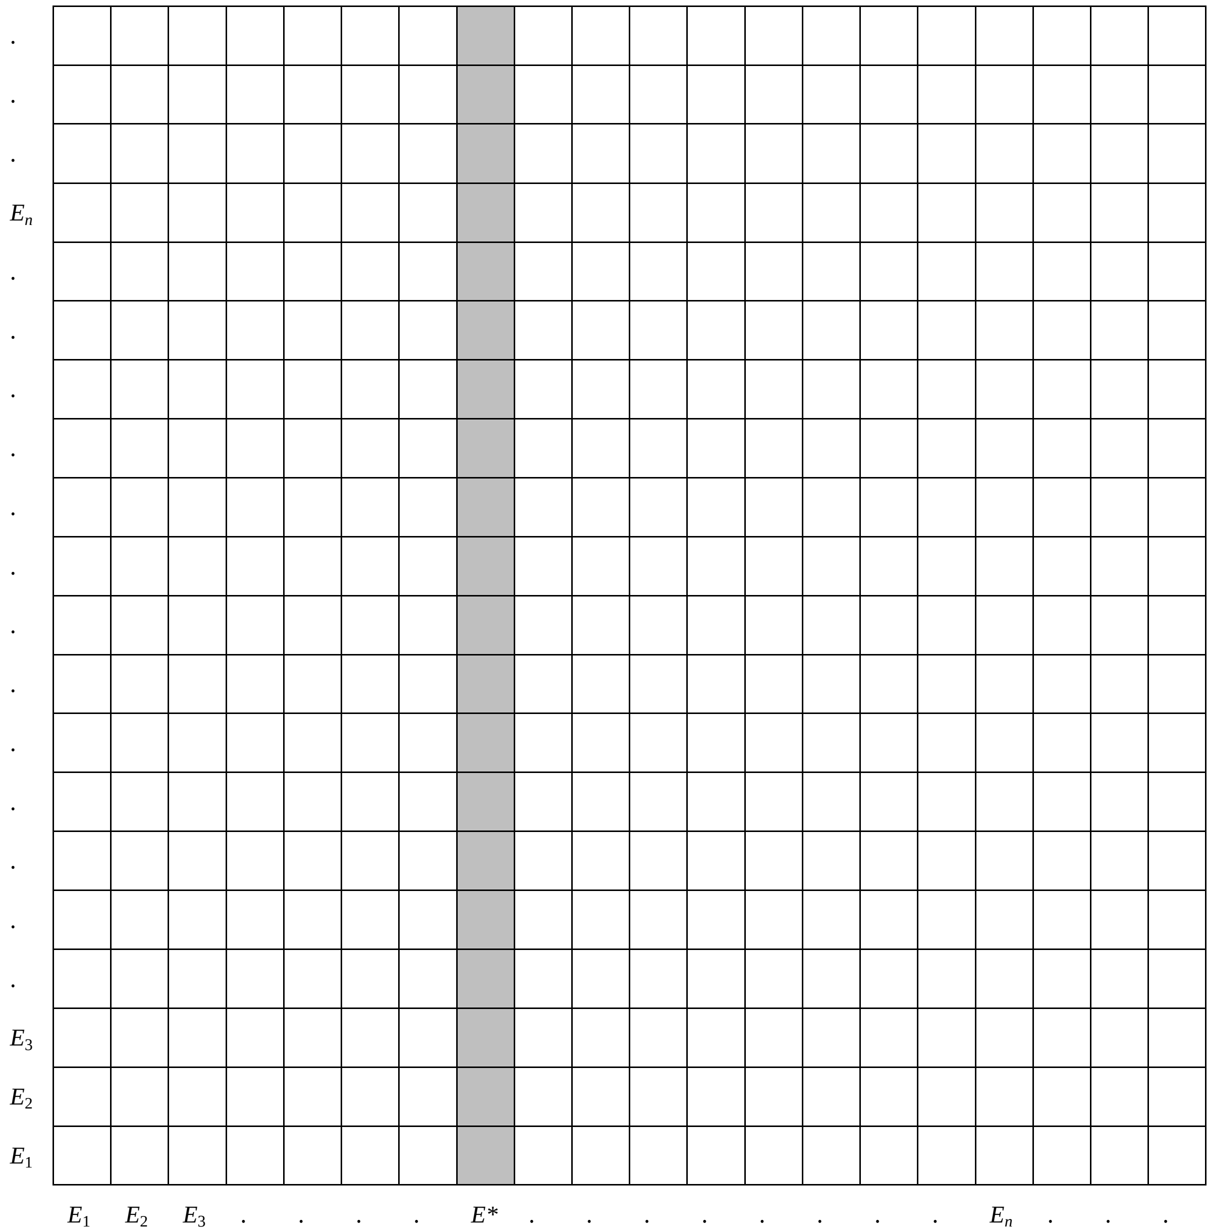

Figure 2. The horizontal axis depicts the partition proposition of which the agent becomes certain at t_2. The vertical axis depicts the true partition proposition, revealed by the oracle at t_3. Skyrms constructs a diachronic semi-Dutch book that inflicts negative net payoff in the grey column where $P_{new}(E^*) = 1$ and that otherwise yields net payoff 0.