

TWO APPROACHES TO BELIEF REVISION

TED SHEAR, JONATHAN WEISBERG, AND BRANDEN FITELSON

ABSTRACT. In this paper, we compare and contrast two methods for revising qualitative (*viz.*, “full”) beliefs. The first method is a naïve Bayesian one, which operates *via* conditionalization (and, more generally, *via* mechanical/minimum distance updating) and the minimization of expected inaccuracy. The second method is the AGM approach to belief revision (which can also be understood in terms of mechanical/minimum distance updating). Our aim here is to provide the most straightforward explanation of the ways in which these two methods agree and disagree with each other, when it comes to imposing diachronic constraints on agents with deductively cogent beliefs. Some novel (and surprising) convergences and divergences between the two approaches are uncovered.

1. SETUP

We will be talking about a very simple and highly idealized epistemic agent who has both numerical degrees of confidence (*viz.*, *credences*) and qualitative beliefs.¹ These attitudes will be held toward (each of) the propositions in an *agenda* \mathcal{A} of (classical, possible worlds) propositions, generated by a simple, finite classical propositional language.² The credence that the agent assigns to a proposition p will be represented by $b(p)$, while the credence she assigns to p on the (indicative) supposition that E is true will be represented by $b(p | E)$. We will assume that $b(\cdot)$ is a (classical) *probability* function, and that our agent updates her credences *via*

Date: 02/10/16. *Draft:* Do not cite or quote without permission.

We would like to thank audiences at Tilburg, Boulder, Maryland, Helsinki, Kent, Taipei, and Toronto for stimulating discussions. Individually, we must single out Luc Bovens, Kenny Easwaran, Haim Gaifman, Konstantinos Georgatos, Jeremy Goodman, Kevin Kelly, Hannes Leitgeb, Isaac Levi, Hanti Lin, Eric Pacuit, Rohit Parikh, Richard Pettigrew, Jonah Schupbach, Teddy Seidenfeld, and Julia Staffel for providing useful feedback.

¹Presently, we remain neutral on questions regarding *reduction* (or *fundamentality*) of these two kinds of doxastic attitudes. We prefer to think of the constraints we’ll describe here as joint (coherence) constraints on credence and belief. In this respect, we follow the pluralistic approach of Leitgeb [23].

²We will use W to denote the set of possible worlds underlying \mathcal{A} . And, the worlds in W will correspond to the *state-descriptions* of the propositional language which generates \mathcal{A} .

conditionalization.³ Because our present interest is in doxastic *dynamics*, we will make use of two credence functions: a *prior* $b(\cdot)$ and a *posterior* $b'(\cdot) = b(\cdot | E)$, where E is (exactly) what the agent learns between the prior and posterior times.

When it comes to the qualitative attitudes of our idealized agent, we will attend only to the *beliefs* of the agent (*i.e.*, we will not discuss disbelief or suspension of judgment here, since these attitudes are not generally discussed within the standard frameworks for qualitative belief revision). We will adopt the notation $B(p)$ to indicate that our agent believes that p , and the notation \mathbf{B} to denote the *set* of propositions (in \mathcal{A}) that our agent believes. Again, because we are interested in the dynamics of these attitudes, we will make use of two belief sets: a *prior* \mathbf{B} and a *posterior* $\mathbf{B}' = \mathbf{B} \star E$, where E is (exactly) what the agent learns between the prior and posterior times and \star is a (qualitative) *belief revision operator*.

Ultimately, we will be comparing and contrasting two belief revision operators: one provided by the traditional AGM theory (\ast) and a novel one based on *epistemic utility theory* (\otimes). AGM’s revision operator is defined on the basis of logical considerations alone and is not sensitive to changes in the agent’s credences.⁴ In contrast, our novel operator will be defined explicitly in terms of the agent’s credence function. While traditional approaches to credence and belief have been developed largely in isolation of each other, more recently various authors have been investigating joint constraints on quantitative and qualitative epistemic attitudes.⁵

One such joint constraint is that beliefs should obey what Foley [10] dubbed a (normative) *Lockean thesis*. The Lockean thesis requires that an agent believe p iff her credence in p is at least t , for some *Lockean threshold* t . Intuitively, this constraint requires that an agent believe all and only the propositions that she takes to be sufficiently probable. Accordingly, it is typically required that $t > 1/2$. This normative Lockean thesis provides an intuitively plausible rational constraint since: (1) it would seem irrational for an agent to believe a proposition that she deems improbable, (2) it would seem irrational to permit concurrent belief in p and $\neg p$ (as would be permitted if $t = 1/2$ and $b(p) = 1/2$), and (3) it would seem rational (*i.e.*, rationally permissible) for an agent to believe all and only propositions that she takes to be (sufficiently) likely.

³One can use ideas from epistemic utility theory to motivate both probabilism and conditionalization. See [27] for a survey of such arguments. Here, we simply *assume* probabilism and conditionalization as background constraints on rational credence. Most of our present results generalize to a broader family of credal updating procedures (including conditionalization and Jeffrey conditionalization as special cases), which can be characterized in mechanical/minimum-distance terms [5].

⁴As we will see later, AGM theory does appropriately respect an agent’s credences when those credences are *maximal*. But, this convergence is only a result of the fact that (probabilistic) *certainties* must behave logically. This is not due to any general sensitivity of AGM to an agent’s credences.

⁵See, most notably, Leitgeb [22], Lin and Kelly [24], Easwaran [7], and Dorst [6].

As we will see shortly, the expected utility theory (EUT) approach to belief revision requires that both the prior and the posterior belief sets and credences satisfy joint Lockean constraints. It is well known that Lockean approaches sometimes permit belief sets that are neither closed under logical consequence nor deductively consistent. On the other hand, approaches to belief revision that rely solely on logical considerations (like AGM) tend to require that agents (always) have deductively closed and consistent belief sets. In comparing AGM with EUT revision, we will not belabor this well-known source of disagreement between the two paradigms. Rather, we will focus most of our attention on some less widely known divergences (and convergences) between the two approaches — understood as constraints on agents with deductively cogent belief sets (at all times).

Accordingly, we will compare the recommendations of each approach for agents whose prior and posterior belief sets are deductively cogent. We will demonstrate that sometimes EUT will recommend that the agent give up a belief that AGM does not. Formally, this means that EUT can violate AGM’s characteristic axiom Vacuity. We will also find that these cases exhaust the divergences between the two since EUT will never recommend belief in a proposition that AGM does not. This is because (1) $\mathbf{B} * E$ will never form a proper superset of $\mathbf{B} * E$, and (2) $\mathbf{B} * E$ and $\mathbf{B} * E$ will always be comparable. We will show (1) by demonstrating that EUT satisfies AGM’s Inclusion axiom and (2) by proving that EUT violates Vacuity just in case $\mathbf{B} * E \subset \mathbf{B} * E$. In this sense, EUT is more epistemically risk averse (or more epistemically conservative) than AGM, since the most interesting cases of disagreement between EUT and AGM will involve cases in which EUT revision requires an agent to make (strictly) fewer epistemic commitments (i.e., end-up with strictly fewer beliefs) than AGM revision does.

In the next section, we sketch EUT theory and provide its veritistic motivation along with an illuminating application: a derivation of the normative Lockean thesis. In section three, we will contrast the synchronic requirements of EUT with those of AGM and briefly introduce AGM’s axioms. In the subsequent two sections, we will give a precise characterization of the similarities and differences between the two approaches to belief revision. Finally, we will close with some remarks about open questions and future work.

The ultimate aim of this paper is to characterize how approaches to belief revision based on the normative Lockean this vs. approaches based on “logical updating” relate to each other. As such, we will not extensively discuss (or attempt to justify) the background motivations for either approach.

2. EPISTEMIC UTILITY THEORY AND ITS REVISION OPERATOR

The first approach to qualitative belief revision that we will consider may be generated based on considerations from *epistemic utility theory*. The basic idea behind EUT is that an agent’s belief set should (at any given time) *maximize expected epistemic utility*, from the point of view of her (current) credence function b . Thus, EUT must presuppose some way of measuring the *epistemic utility* of a belief set \mathbf{B} .

Following Easwaran [7] and Dorst [6], we will assume that our agent has a very simple epistemic utility function, which cares only about whether the agent’s beliefs are *accurate*. Specifically, we will assume that our agent has the following two-parameter epistemic utility function.

$$u(\mathbf{B}(p), w) := \begin{cases} \mathbf{r} & \text{if } p \text{ is true at } w \\ -\mathbf{w} & \text{if } p \text{ is false at } w \end{cases}$$

That is, if p is true at w , then the utility associated with believing that p at w is some non-negative number \mathbf{r} ; and, if p is false at w , then the utility associated with believing that p at w is some non-positive number $-\mathbf{w}$. This naïve EUT approach can be motivated on (broadly) Jamesian grounds [20], as it may be plausibly suggested that belief aims at truth. Thus, an agent enjoys some “epistemic credit” for believing truths and incurs some “epistemic debit” for believing falsehoods.

When it comes to assumptions about the value ranges of the two parameters \mathbf{r} and \mathbf{w} , we will impose the following single constraint.⁶

$$(\dagger) \quad 1 \geq \mathbf{w} > \left(\frac{1 + \sqrt{5}}{2} \right) \cdot \mathbf{r} > 0.$$

We will discuss this key constraint — including the somewhat surprising appearance in (\dagger) of the *Golden Ratio*, which we will abbreviate hereafter as ϕ — in detail below. As we will see in section five, the inclusion of ϕ in (\dagger) will follow naturally from our desire to focus on precisely those cases in which the Bayesian/Lockean approach to belief revision can diverge from the AGM approach, as a constraint on the beliefs of *deductively cogent* EUT agents.

With this naïve, accuracy-centered utility function in hand, it is straightforward to define the *expected epistemic utility* (EEU) of an agent’s belief state.

⁶Most authors who apply EUT to Lockeanism (e.g., Dorst [6]) minimally assume that $\mathbf{w} > \mathbf{r}$. This assumption can be justified by appealing to the widely held view that all Lockean thresholds should be *strictly* greater than $1/2$. It is natural to require that Lockean thresholds be *strictly* greater than $1/2$, since allowing thresholds *equal* to $1/2$ would make it permissible for a Lockean agent to believe *both* p and $\neg p$, so long as $b(p) = 1/2$. Pruss [28] argues that \mathbf{w} should be at least $\frac{1}{\log(4)-1} \approx 2.588$ times as great as \mathbf{r} , based on some assumptions about the nature of epistemic accuracy. This is stronger than our assumption (\dagger) , which only requires that \mathbf{w} is at least $\phi \approx 1.618$ times as great as \mathbf{r} . Our (\dagger) also implies that the values of \mathbf{r} and \mathbf{w} are on the half-open unit interval $(0, 1]$. But, this choice of scale is purely conventional, as any linear transformation of (\mathbf{r}, \mathbf{w}) would yield the same requirements.

Definition The *expected epistemic utility* (EEU) of an agent's *belief* that p , $\mathbf{B}(p)$, from the point of view of her credence function b , is defined as:

$$EEU(\mathbf{B}(p), b) := \sum_{w \in W} b(w) \cdot u(\mathbf{B}(p), w).$$

Then, the overall EEU of an agent's total belief set \mathbf{B} (over an entire agenda \mathcal{A}) is provided by the sum of all of the EEUs of the agent's beliefs that comprise \mathbf{B} .

Definition The *expected epistemic utility* (EEU) of an agent's *belief set* \mathbf{B} , on agenda \mathcal{A} , from the point of view of her credence function b is defined as:

$$EEU(\mathbf{B}, b) := \sum_{p \in \mathbf{B}} EEU(\mathbf{B}(p), b).$$

Dorst [6] demonstrates that a belief set *maximizes* EEU relative to a credence function b just in case it satisfies the following precise (normative) Lockean thesis.⁷

Theorem (Dorst [6]) An agent's belief set \mathbf{B} (over agenda \mathcal{A}) *maximizes EEU from the point of view of her credence function* b if and only if, for every $p \in \mathbf{B}$

$$b(p) > \frac{w}{r+w}.$$

In other words, \mathbf{B} maximizes EEU just in case our agent believes exactly those propositions in \mathcal{A} to which they assign a sufficiently high credence, where the Lockean threshold for belief is itself implied by EUT to be the ratio $\frac{w}{r+w}$.

Dorst's theorem can be used to ground both synchronic and diachronic constraints on qualitative belief sets. Synchronically, an agent's belief set must obey the precise Lockean coherence requirement above. And, this synchronic coherence requirement immediately lifts to a diachronic coherence requirement, given the assumption that our agent updates her credences *via* conditionalization. That is to say, naive EUT gives rise to the following belief revision operator.

Definition (EUT Revision) Suppose an agent has an initial belief set \mathbf{B} , and she learns (exactly) E . Her new belief set \mathbf{B}' should maximize EEU, relative to her *conditional* credences $b(\cdot | E)$. That is, her new belief set $\mathbf{B}' = \mathbf{B} * E$ should be

$$\mathbf{B} * E := \left\{ p \mid b(p | E) > \frac{w}{r+w} \right\}^8$$

⁷It is worth noting that a similar result is also proved (independently) in Easwaran [7], although Easwaran's applications of his result are much different than Dorst's. Historically, these kinds of EUT-derived Lockean-style constraints trace back to the work of Hempel [19].

⁸This definition can be generalized significantly, by characterizing Bayesian updating in mechanical (viz., minimum-distance) terms [5]. More generally, our results regarding EUT vs. AGM revision will all continue to hold for any Bayesian update b' (on learned proposition E) which satisfies the following three constraints: (i) $b'(E) > b(E)$, (ii) $b'(E) > t$ (where t is the agent's EUT Lockean threshold), and (iii) $b(E \supset X) \geq b'(X)$. So, more generally, we can think of Bayesian revision b' on E as taking the agent from her prior b to the *closest probability function* b' , *satisfying the three constraints (i)–(iii)*.

As we will see, this EUT revision operator $*$ will have several novel features that distinguish it from the more traditional AGM revision operator (to be discussed below). We will compare and contrast these two operators in detail shortly.

Meanwhile, as a preview of how that dialectic will unfold, it is useful to note that $*$ does not (generally) satisfy the following principle, which has origins in the earliest work on belief revision⁹ and was more recently employed by Leitgeb [23] to motivate a diachronic cousin of the preface paradox.

(P2) If an agent initially believes X (i.e., if $X \in \mathbf{B}$), then updating \mathbf{B} on X *should not change* \mathbf{B} . [More formally, $X \in \mathbf{B} \Rightarrow \mathbf{B}' = \mathbf{B} * X = \mathbf{B}$.]

This principle is always satisfied by the AGM revision operator [23].¹⁰ But, it need not be satisfied by the EUT revision operator $*$. However, the following proposition offers some further insight into such failures by placing a bound on the “degree” to which (P2) can fail from the perspective of EUT.

Proposition Suppose $b(p) > \frac{w}{r+w}$ and $b(q) > \frac{w}{r+w}$ (i.e., that our deductively cogent EUT agent believes both p and q). And, following our constraint (\dagger), suppose that $\phi - 1 < \frac{w}{r+w} \leq 1$.¹¹ Then, $b(p | q) > \frac{w-r}{w}$, and $b(p | q) - b(p) < \frac{r^2}{rw+w^2}$.

If an agent learns (with certainty) something (q) she *already* believes, then the lowest possible credence that she can coherently assign to any of her previous beliefs (p) is provided by the ratio. And, the more confident our agent is in p and q , the smaller this decrement can be. Moreover, in the limit, (P2) *will* be satisfied by EUT revision *if* our agent assigns *maximal* credence to her beliefs (more on this extremal case in section 4).

Figure 1 furnishes a visual explanation of what is going on with (P2), from an EUT point of view. We think this goes some way toward explaining why (P2) may *seem* like a plausible diachronic constraint on full belief, since it is “approximately” true if full beliefs have sufficiently high credence (and it is *exactly* true in the extremal case). We will return to (P2) in the final section of the paper, where we provide a counterexample to the underlying postulate of AGM belief revision theory (Vacuity) which undergirds (P2).

⁹This principle was first identified as (C9) in Gärdenfors' seminal work [13], but has more extensively been discussed as Weak Preservation, e.g. see [2].

¹⁰(P2) holds in AGM theory *only if* \mathbf{B} is assumed to be *deductively cogent*. Although this is a standard assumption to make and is directly in line with AGM's motivations, it is absent from a number of notable presentations (e.g. [16]). We will give a proof of (P2), given this assumption, once we've introduced the postulates of AGM revision theory. Then, we will give an EUT-theoretic counterexample to the key AGM postulate (Vacuity) that undergirds (P2).

¹¹Our constraint (\dagger) and the Lockean threshold $\frac{w}{w+r}$ (which follows from the maximization of expected epistemic utility [6]), jointly entail this range of Lockean thresholds.




FIGURE 1. The “degree” to which EUT belief updating can violate Leitgeb’s (P2), as a function of the EUT Lockean threshold $t = \frac{w}{r+w}$.

In the next section, we introduce AGM theory and its revision operator. In subsequent sections, we explain the relationship(s) between EUT and AGM revision.

3. AGM THEORY AND ITS REVISION OPERATOR

3.1. The Synchronic Presuppositions of AGM vs EUT. While the AGM theory [1] is a diachronic theory of belief revision, it presupposes a *synchronic* coherence requirement for full belief that is *strictly stronger* (i.e., *strictly more demanding*) than the synchronic Lockean coherence requirement implied (via Dorst’s theorem) by EUT. Specifically, AGM theory presupposes the following.

Consistency. An agent’s belief set \mathbf{B} should (always) be *deductively consistent*, i.e., there should be some possible world w such that every member of \mathbf{B} is true at w .

It is important to recognize that **Consistency** is implied by an *extreme form* of Lockean representability, where the threshold is *maximal*. Thus, from the point of view of EUT, **Consistency** (as a universal, global requirement) is tantamount to

requiring that an agent’s utilities be such that $w = 1$ and $r = 0$.¹² As such, from the point of view of EUT, **Consistency** is *too demanding*. As Foley [10, p. 186] explains,

... if the avoidance of recognizable inconsistency were an absolute prerequisite of rational belief, we could not rationally believe each member of a set of propositions and also rationally believe of this set that at least one of its members is false. But this in turn pressures us to be unduly cautious. It pressures us to believe only those propositions that are certain or at least close to certain for us, since otherwise we are likely to have reasons to believe that at least one of these propositions is false. At first glance, the requirement that we avoid recognizable inconsistency seems little enough to ask in the name of rationality. It asks only that we avoid certain error. It turns out, however, that this is far too much to ask.

We think EUT provides a nice formal explication of the notion of epistemic rationality that Foley has in mind. While belief aims at truth, this does *not* mean that epistemic *rationality* requires **Consistency**. Rather, epistemic rationality requires something weaker — that an agent’s beliefs *minimize expected inaccuracy*, from the point of view of her credences [8].

In fact, AGM theory actually presupposes a synchronic constraint that is *even stronger* than **Consistency**. That is, AGM theory *actually* presupposes:

Cogency. An agent’s belief set \mathbf{B} should (at any given time) be *deductively cogent*, i.e., \mathbf{B} should be *both deductively consistent and closed under logic*.

Closure under logic (i.e., closure under tautological entailment) is sometimes thought of as a diachronic constraint on belief sets. Indeed, as we will see shortly, AGM theorists typically (also) state closure (as well as consistency) as a constraint on their belief revision operator. However, these purportedly diachronic constraints do not seem to be (essentially) diachronic. Instead they seem (ultimately) to require that revisions take synchronically coherent belief sets as both inputs and outputs. Thus, we prefer to think of these as (standing) synchronic constraints.¹³

¹²Strictly speaking, this extreme case does not fall out of Dorst’s approach, since this would involve a condition on belief $[b(p) > 1]$ that can never be met by a probabilistically coherent agent. Nonetheless, we can extend Dorst’s approach to allow for this extreme case, by adding a clause saying that $b(p) = 1$ is the requirement on $\mathbf{B}(p)$ when $w = 1$ and $r = 0$. This is the assumption we will make.

¹³Steinberger [31] has argued convincingly that, if failures of consistency are permitted, then closure loses much (if not all) of its normative force. It is for this reason that we prefer to (a) take consistency to be a more fundamental epistemic requirement than closure, and (b) to take both consistency and closure as standing, synchronic presuppositions of the AGM framework (see *fn*s. 17 and 20). In the

In the next subsection, we will rehearse the main principles of AGM theory (*i.e.*, the AGM axioms for revision operators). Then, in the next section, we'll begin to explore the relationship(s) between AGM revision and EUT revision.

3.2. The AGM Postulates. AGM theory can be presented in various, equivalent ways.¹⁴ But, for our purposes, it will be most perspicuous to present AGM in terms of the postulates, which (when taken as axioms) provide a representation of AGM revision operator $*$.¹⁵ Here are the six basic AGM principles (plus two supplemental postulates). Note: the function $\text{Cn}(\cdot)$ takes a set of propositions \mathbf{P} and returns the *logical closure* of \mathbf{P} , *i.e.*, the set of *tautological entailments* of \mathbf{P} .¹⁶

(*1) $\mathbf{B} * E = \text{Cn}(\mathbf{B} * E)$ Closure

- In words, (*1) says that if an agent learns (exactly) E , then their posterior belief set $\mathbf{B}' = \mathbf{B} * E$ should be *closed under logic*. As we explained above, we think Closure derives from a pre-theoretic commitment to closure as a *synchronic* requirement.

(*2) $E \in \mathbf{B} * E$ Success

- In words, (*2) says that if an agent learns (exactly) E , then their posterior belief set $\mathbf{B}' = \mathbf{B} * E$ should *contain* E .

(*3) $\mathbf{B} * E \subseteq \text{Cn}(\mathbf{B} \cup \{E\})$ Inclusion

- In words, (*3) says that if an agent learns (exactly) E , then their posterior belief set $\mathbf{B}' = \mathbf{B} * E$ should be a *subset* of $\text{Cn}(\mathbf{B} \cup \{E\})$.

(*4) If E is consistent with \mathbf{B} , then $\mathbf{B} * E \supseteq \text{Cn}(\mathbf{B} \cup \{E\})$ Vacuity

- In words, (*4) says that if an agent learns (exactly) some E that is consistent with their prior belief set \mathbf{B} , then their posterior belief set $\mathbf{B}' = \mathbf{B} * E$ should be a *superset* of $\text{Cn}(\mathbf{B} \cup \{E\})$. Note: (*3) and (*4) jointly imply that if an agent learns (exactly) some E that is consistent

next subsection, we will see that this way of thinking about **Closure** and **Consistency** is supported by AGM's formulation of these constraints, which only constrain the posterior belief set $(\mathbf{B} * E)$.

¹⁴One intuitive way to understand AGM updating is in terms of *minimal change updating*. Any procedure that takes a (coherent) prior belief set \mathbf{B} to the *closest (coherent) belief set to \mathbf{B} which accommodates E* will (inevitably) satisfy all the AGM postulates (below) for revising \mathbf{B} by E . At least, this will be true for a very broad class of measures of "distance between prior and posterior belief sets." See [29] for a nice survey of these sorts of mechanical/minimum-distance qualitative belief revision methods.

¹⁵As it turns out, the so-called "Basic Gärdenfors Postulates" (*1)–(*6) provide an axiomatization of *partial meet revision* operations, which can be thought of as emerging from the minimally mutilating revision of some prior belief set \mathbf{B} in accord with an entrenchment ordering on propositions. The addition of the supplementary postulates, (*7) and (*8), yields a characterization of a special class of partial meet contraction operations: the *transitively relational* ones. See [12] for an overview of the various ways of characterizing AGM belief revision operators.

¹⁶One can interpret these axioms more generally, in terms of a generalized entailment relation (which may be non-classical). For simplicity, we will assume a classical entailment relation here. What we say below can be generalized to various non-classical (*e.g.*, substructural) entailment relations.

with their prior belief set \mathbf{B} , then their posterior belief set $\mathbf{B}' = \mathbf{B} * E$ should be *identical to* $\text{Cn}(\mathbf{B} \cup \{E\})$.

(*5) If E is not self-contradictory, then $\mathbf{B} * E$ is consistent. Consistency

- In words, (*5) says that if an agent learns (exactly) some E that is not itself a contradictory proposition, then their posterior belief set $\mathbf{B}' = \mathbf{B} * E$ should be deductively consistent. As we explained above, we think Consistency derives from a pre-theoretic commitment to **Consistency** as a *synchronic* requirement.¹⁷

(*6) If $X \equiv Y \in \text{Cn}(\emptyset)$, then $\mathbf{B} * X = \mathbf{B} * Y$ Extensionality

- In words, (*6) says that if X and Y are *tautologically equivalent*, then updating on X should have *exactly the same effect* as updating on Y .

In addition to the six basic postulates, the AGM's revision operator is often taken to satisfy two supplementary postulates (which, as we will see below, are generalizations of (*3) and (*4), respectively). While the supplementary postulates have generated more resistance in the literature¹⁸, we opt to discuss AGM Revision in its simplest and most standard form for considerations of simplicity.

(*7) $\mathbf{B} * (X \wedge Y) \subseteq \text{Cn}((\mathbf{B} * X) \cup \{Y\})$ Superexpansion

- In words, (*7) says that updating the prior belief set \mathbf{B} on the *conjunction* $X \wedge Y$ should yield a posterior belief set that is a *subset* of $\text{Cn}((\mathbf{B} * X) \cup \{Y\})$, which is the deductive closure of the union of $\{Y\}$ and the belief set $\mathbf{B} * X$ (the result of updating the prior belief set \mathbf{B} on X alone).

(*8) If Y is consistent with $\text{Cn}(\mathbf{B} * X)$, then $\mathbf{B} * (X \wedge Y) \supseteq \text{Cn}((\mathbf{B} * X) \cup \{Y\})$ Subexpansion

- In words, (*8) says that if Y is consistent with $\mathbf{B} * X$ (the result of updating the prior belief set \mathbf{B} on X alone), then updating the prior belief set \mathbf{B} on the *conjunction* $X \wedge Y$ should yield a posterior belief set that is a *superset* of $\text{Cn}((\mathbf{B} * X) \cup \{Y\})$. Note: (*7) and (*8) jointly imply

¹⁷Consider the closed, but inconsistent belief set $\mathbf{B} = \{P, \neg P, \top, \perp\}$, where P is a contingent (atomic) claim. Consistency implies that $\mathbf{B} * \top$ is consistent. Thus, according to AGM, if an agent starts out with the prior belief set \mathbf{B} and then "revises by a tautology \top ," they must (as a result of this revision) abandon either their belief in P or their belief in $\neg P$ (since, otherwise, the consistency of $\mathbf{B} * \top$ will not be ensured). But, it is counter-intuitive that "learning a tautology" should provide an agent with a conclusive reason to drop one of their contingent beliefs. This drives home the point that AGM theory really needs to presuppose consistency as a standing, synchronic constraint on all belief sets.

¹⁸Most notably, Stalnaker [30] suggests that the supplementary postulates provide untoward treatments in the context of iterated belief revision since they fail to acknowledge certain evidential considerations. We are especially sympathetic to Stalnaker's general point. However, his insight is incidental with respect to our current purposes and is worthy of independent consideration.

that if an agent learns (exactly) some Y that is consistent with $\mathbf{B} * X$, then updating the prior belief set \mathbf{B} on the *conjunction* $X \wedge Y$ should yield a posterior belief set that is *identical to* $\text{Cn}((\mathbf{B} * X) \cup \{Y\})$.

It is worth noting that Superexpansion and Subexpansion can (for present purposes) be seen as *generalizations of* Inclusion and Vacuity, respectively. More precisely, Superexpansion and Subexpansion *entail* Inclusion and Vacuity, respectively — provided that the following additional postulate is accepted:¹⁹

(*T) $\mathbf{B} * \top = \mathbf{B}$ Idempotence

Intuitively, Idempotence says that “learning a tautology” should not change your belief set (in any way). Of course, the EUT approach to belief updating will always respect Idempotence. Moreover, this constraint seems right to us (in general), and so we think that AGM-theorists should also accept Idempotence (in general).²⁰ Indeed, we will presuppose Idempotence for the remainder of the paper.

The following derivation establishes that Superexpansion, Extensionality, and Idempotence jointly entail Inclusion:

- | | |
|---|--|
| 1. $Y \in \mathbf{B} * X$ | Assumption |
| 2. $\mathbf{B} * X = \mathbf{B} * (\top \wedge X)$ | (1), Extensionality |
| 3. $Y \in \mathbf{B} * (\top \wedge X)$ | (1), (2), Logic |
| 4. $Y \in \text{Cn}((\mathbf{B} * \top) \cup \{X\})$ | (3), Logic, Superexpansion |
| 5. $\text{Cn}((\mathbf{B} * \top) \cup \{X\}) = \text{Cn}(\mathbf{B} \cup \{X\})$ | Idempotence, Logic |
| 6. $Y \in \text{Cn}(\mathbf{B} \cup \{X\})$ | (4), (5), Logic □ |

Similarly, the following derivation establishes that Subexpansion, Extensionality, and Idempotence jointly entail Vacuity:

- | | |
|--|--|
| 1. X is consistent with \mathbf{B} | Assumption |
| 2. $Y \in \text{Cn}(\mathbf{B} \cup \{X\})$ | Assumption |
| 3. $Y \in \text{Cn}((\mathbf{B} * \top) \cup \{X\})$ | (2), Idempotence, Logic |
| 4. $Y \in \mathbf{B} * (\top \wedge X)$ | (3), Subexpansion |
| 5. $Y \in \mathbf{B} * X$ | (4), Extensionality □ |

These simplifications will play a role in the next section. They will allow us to provide a single EUT-counterexample to *both* Subexpansion *and* Vacuity.

¹⁹A slightly weaker version of this postulate first appeared in the belief revision literature in [11].

²⁰We saw above (fn. 17) that Idempotence can be violated by AGM theory, if **Consistency** is *not* assumed as a standing, synchronic requirement. We think the most plausible and natural interpretation of AGM is one that accepts *both* **Consistency** (as a standing, synchronic constraints) *and* Idempotence. Thus, we think the most natural interpretation of AGM makes Inclusion and Vacuity *consequences of* Superexpansion and Subexpansion. This will simplify some of our arguments below.

In the next subsection, we’ll look at a proof of Leitgeb’s (P2) from the AGM postulates above. This will help to clarify the content of the postulates (and how they can be applied), as well as the precise role of consistency and closure in AGM.

3.3. A Proof of Leitgeb’s (P2) from the AGM Postulates. Here is a more careful statement of the version of (P2) that follows from the AGM postulates.

- (P2) If an agent initially believes X (i.e., if $X \in \mathbf{B}$), then — provided that \mathbf{B} satisfies **Cogency** — updating \mathbf{B} on X *does not change* \mathbf{B} . [More formally, $X \in \mathbf{B} \Rightarrow \mathbf{B}' = \mathbf{B} * X = \mathbf{B}$.]

Here is a proof of (P2) from the AGM postulates Closure, Inclusion and Vacuity.

- | | |
|---|---|
| 1. \mathbf{B} is consistent. | Assumption |
| 2. \mathbf{B} is closed, i.e., $\mathbf{B} = \text{Cn}(\mathbf{B})$. | Assumption |
| 3. $X \in \mathbf{B}$. | Assumption |
| 4. X is consistent with \mathbf{B} . | (1), (3), Logic |
| 5. $\mathbf{B} * X = \text{Cn}(\mathbf{B} \cup \{X\})$. | (4), Vacuity, Inclusion |
| 6. $\mathbf{B} * X = \text{Cn}(\mathbf{B})$. | (5), (3), Logic |
| 7. $\mathbf{B} * X = \text{Cn}(\mathbf{B} * X)$ | Closure |
| 8. $\text{Cn}(\mathbf{B} * X) = \text{Cn}(\mathbf{B})$ | (6), (7), Logic |
| 9. $\mathbf{B} * X = \mathbf{B}$ | (7), (8), (2), Logic □ |

Note: assumptions (1) and (2) are *crucial* here. To see why (2) is needed, note that if \mathbf{B} is *not* closed, then $\mathbf{B} * X \neq \mathbf{B}$, since Closure requires that $\mathbf{B} * X$ be closed. To see why (1) is needed, note that if \mathbf{B} is *not* consistent, then the precondition for Vacuity [*viz.*, (4)] will not be met, and the derivation of (5) will be blocked.²¹

In the next section, we’ll briefly discuss the (well-known) fact that *extremal* EUT revision (i.e., EUT revision *with extremal epistemic utilities*) is AGM revision. Then, in the subsequent sections, we’ll explore the ways in which EUT revision and AGM revision *diverge* in the general (non-extremal) case.

4. *Extremal* EUT REVISION is AGM REVISION

If our agent has *extremal* epistemic utilities (i.e., if $r = 0$ and $w = 1$), then our EUT agent’s Lockean threshold is *maximal* (see fn. 12). In this case, our agent will believe p iff they assign p *maximal* credence. It is easy to see that, in this extremal case, our agent’s belief set \mathbf{B} will *always* satisfy **Cogency**. As a result, extremal EUT agents must satisfy both Closure and Consistency. Furthermore, it has been known for some time that extremal EUT agents must satisfy *all of the other AGM postulates as well*. To wit, we have the following (classic) theorem.

²¹This drives home the point that **Cogency** must be *presupposed* as a standing, synchronic requirement, in order for AGM to have the consequences it is normally taken to have (see fns. 17 and 20).

Theorem (Gärdenfors [15]) Suppose $r = 0$, $w = 1$, \mathbf{B} is synchronically coherent in the EUT sense, and that for all propositions X and Y that our agent might learn, $b(X | Y) > 0$.²² Then \ast satisfies all eight of the AGM postulates above.

The situation is much more interesting when our agent's epistemic utilities (and, hence, Lockean thresholds) are *non-extremal*. In the general case, the relationship between EUT revision and AGM revision is more nuanced.

5. Non-EXTREMAL EUT REVISION IS *Not* AGM REVISION

In the general, non-extremal case, EUT revision and AGM revision *diverge significantly*. In this section, we explore this divergence. Counterexamples are available for *four* of the eight AGM postulates. As we have suggested, EUT revision's violation of Consistency and Closure has been widely discussed and so we will not bother to rehearse their counterexamples. Moreover, since in the presence of Idempotence Vacuity is implied by Subexpansion, the counterexample to the former will suffice to establish its failure to satisfy the latter. But, before we get to those four negative results, we begin with four *positive* results.

5.1. Convergences between EUT Revision and AGM Revision. The following four propositions exhaust the AGM postulates that are satisfied by EUT.

Proposition 1 EUT Revision (Generally) *Satisfies Success*.

Proof. Provided that $b(E) > 0$ (which we're assuming about *all* potential pieces of evidence — see *fn.* 22), it is a theorem of probability calculus that $\Pr(E | E) = 1$. Therefore, $b(E | E) = 1 \geq t$, for *any* Lockean threshold t . So, $E \in \mathbf{B} \ast E$. \square

Proposition 2 EUT Revision (Generally) *Satisfies Inclusion*.

Proof. Suppose $X \in \mathbf{B} \ast E$. Then, $b(X | E) \geq t$. And, it is a theorem of probability calculus that $\Pr(E \supset X) \geq \Pr(X | E)$. Therefore, $b(E \supset X) \geq t$. So, $E \supset X \in \mathbf{B}$. Hence, by *modus ponens* (for material implication), $X \in \mathbf{Cn}(\mathbf{B} \cup \{E\})$. \square

Proposition 3 EUT Revision (Generally) *Satisfies Extensionality*.

Proof. Suppose X and Y are tautologically equivalent. Then, X and Y are *probabilistically indistinguishable* (under *every* probability function). Therefore, EUT revisions on X are indistinguishable from EUT revisions on Y . \square

Proposition 4 EUT Revision (Generally) *Satisfies Superexpansion*.

²²In Kolmogorvian probability theory, we cannot conditionalize on propositions that have zero probability. For this reason, we must assume that our agents only learn things to which they assign non-zero credence. There are generalizations of Gärdenfors's Theorem which make use of Popper functions [17, 18]. Such generalizations allow for EUT-style modeling of agents who learn propositions with zero credence. This allows for an EUT-style approach to *belief contravening revision* (and not merely *expansion*). Here, we focus exclusively on Kolmogorvian EUT agents (*i.e.*, expansion).

Proof. Suppose $Z \in \mathbf{B} \ast (X \wedge Y)$. Then, $b(Z | X \wedge Y) \geq t$. And, it is a theorem of probability calculus that $\Pr(Y \supset Z | X) \geq \Pr(Z | X \wedge Y)$. So, $b(Y \supset Z | X) \geq t$. Therefore, $Y \supset Z \in \mathbf{B} \ast X$. And, by *modus ponens* (for material implication), $Z \in \mathbf{Cn}((\mathbf{B} \ast X) \cup \{Y\})$. \square

These four positive results provide the ways in which EUT and AGM converge.

EUT's satisfaction of Inclusion and Superexpansion is of particular interest, since at first sight it may not have been so obvious that this would obtain. Since EUT Revision is driven by the Bayesian apparatus, we might have been inclined to think that there may be cases in which an agent may acquire a new belief in some proposition X , which is probabilistically (but not logically) dependent on the learned proposition E . However, we have seen that EUT's synchronic requirements ensure that any time such a new belief is acquired, it might have equally well been acquired through *modus ponens*. Towards the end of the next section, we will examine whether a convergence between EUT and AGM's *new* beliefs holds more generally. Ultimately, we will show that *sometimes* AGM will require the agent to form more new beliefs than EUT.

5.2. Divergences between EUT and AGM. Now that we have identified the similarities between the two approaches, we turn our attention to the ways in which the two differ.

The first divergence that we note has been extensively discussed in the literature and, as such, we omit the proof.

Proposition 5 Non-Extremal EUT Revision *Violates Consistency and Closure*.

It has been known since the early 1960's [21] that non-extremal Lockean representability is compatible with failures of Consistency (*e.g.*, the lottery paradox). And, of course, if Consistency fails, then Closure must also fail (on pain of epistemic triviality — see *fn.* 13). So, the well-known paradoxes of consistency will (inevitably) yield examples of *non-extremal* EUT updating which violate both Consistency and Closure. For present purposes, we are not so interested in this well-known divergence between EUT and AGM (see [3] for a survey). Rather, we are more interested in cases where belief sets *satisfy* deductive cogency, but EUT revision and AGM revision *still* (substantively) disagree.

Our next (central) counterexample highlights this deeper kind of divergence between these two approaches to belief revision.

Proposition 6 Non-Extremal EUT Revision *Violates Vacuity* — *even if it is restricted to deductively cogent agents*.

Proof. Suppose a Lockean threshold of $t = 0.85$ (*i.e.* $w = 0.17$ and $r = 0.03$). And, suppose that our agent's prior $b(\cdot)$ and posterior $b(\cdot | E)$ are given by the

probability assignments depicted in Table 1 over the Boolean algebra \mathcal{A} generated by a propositional language containing two atomic sentences: E and X .

p	$b(p)$	$b(p E)$	$p \in \mathbf{B}$?	$p \in \text{Cn}(\mathbf{B} \cup \{E\})$?	$p \in \mathbf{B} * E$?	
$E \wedge X$	2/10	2/3	No	Yes	No	✓
$E \wedge \neg X$	1/10	1/3	No	No	No	
$\neg E \wedge X$	4/10	0	No	No	No	
$\neg E \wedge \neg X$	3/10	0	No	No	No	
E	3/10	1	No	Yes	Yes	
X	6/10	2/3	No	Yes	No	✓
$E \equiv X$	5/10	2/3	No	Yes	No	✓
$E \neq X$	5/10	1/3	No	No	No	
$\neg E$	7/10	0	No	No	No	
$\neg X$	4/10	1/3	No	No	No	
$E \vee X$	7/10	1	No	Yes	Yes	
$E \vee \neg X$	6/10	1	No	Yes	Yes	
$\neg E \vee X$	9/10	2/3	Yes	Yes	No	✓
$\neg E \vee \neg X$	8/10	1/3	No	No	No	

TABLE 1. Counterexample to Vacuity for EUT Revision

It will be instructive to present an intuitive urn case that represents the agent's epistemic situation. Suppose we have an urn containing four types of objects: black squares, red squares, black circles, and red circles. We are going to sample an object from the urn at random. And, we assume the following interpretations of the two atomic sentences E and X :

E := 'The object sampled from the urn is red', and

X := 'The object sampled from the urn is a circle'.

The urn contains ten (10) objects, distributed in the following way: four (4) black circles, three (3) black squares, one (1) red square and two (2) red circles (see Figure 2 for a graphical representation of the urn). We will assume our agent has credences in propositions about the shapes and colors of the objects in the urn which are calibrated to this distribution. In this case, our (EUT) agent's prior belief set will be the following singleton.²³

$$\mathbf{B} = \{\neg E \vee X\}$$

Upon learning that the sampled object was red (*i.e.* upon learning E), the agent

²³We omit reference to the contradictory proposition \perp and the tautological proposition \top , since all coherent EUT agents will always have the same attitudes toward those two propositions.




FIGURE 2. Visualization of counterexample to Vacuity for EUT Revision

loses her prior belief that the next ball drawn will either be non-red or a circle since her credence in $\neg E \vee X$ has now dropped below the threshold from 0.9 to $2/3$. Having learned E and lost belief in $\neg E \vee X$, the only other new beliefs that she acquires are the logical consequences of the learned proposition (because they now are assigned maximal credence). That is, after learning E , the agent's posterior belief set is:

$$\mathbf{B}' = \mathbf{B} * E = \{E, E \vee X, E \vee \neg X\}.$$

Note, we have the following four crucial facts in this example (which can all be verified by inspection of Table 1).

- Both the prior belief set \mathbf{B} and the posterior belief set $\mathbf{B} * E$ are *deductively cogent*. That is, the agent in question is *deductively cogent at all times*.
- E is consistent with \mathbf{B} .
- Since $E \supset X \in \mathbf{B}$, it follows (by *modus ponens* for \supset) that $X \in \text{Cn}(\mathbf{B} \cup \{E\})$.
- *But*, $X \notin \mathbf{B} * E$.

Therefore, this is a counterexample to Vacuity for EUT revision — even for some deductively cogent agents. \square

It is important to note that this counterexample *satisfies Cogency*. As such, *this* disagreement between EUT revision and AGM revision is *orthogonal* to the traditional disputes between “Bayesian” vs “logical” schools of thought in formal epistemology (which tend to be obsessed with **Cogency** [3]). In this sense, our counterexample to Vacuity reveals a more fundamental disagreement about diachronic epistemic requirements.²⁴

Proposition 7 *Non-Extremal EUT Revision Violates Subexpansion — even if it is restricted to deductively cogent agents.*

Because Vacuity follows from Subexpansion (assuming Idempotence), the demonstration that EUT fails to satisfy Vacuity suffices to establish the final divergence.

5.2.1. *Constraints on EUT Lockean Thresholds for Vacuity Violators.* In our counterexample to Vacuity, the EUT Lockean threshold is rather high (0.85). This raises the question: *how low* can the EUT Lockean threshold of such an agent be? That is, how low can the Lockean threshold be for an agent who violates Vacuity (*via* learning some E that they do not already believe²⁵). Interestingly, the answer to this question involves the *Golden Ratio* (ϕ). Specifically, the range of Lockean thresholds for which such violations of Vacuity are possible is the half-open interval $[\phi - 1, 1) \approx [0.618, 1)$. That is, we have the following result.

Proposition 8 *If an EUT/Lockean agent is deductively cogent (at all times), then they can only violate Vacuity (*via* learning some E that they do not already believe) if their Lockean threshold is on the half-open interval $[\phi - 1, 1)$.*

We have placed the proof of this proposition in the APPENDIX. Note that this result explains why we imposed constraint (+) on the epistemic utility parameters \mathfrak{w} and \mathfrak{r} . Because rational (*i.e.*, expected epistemic utility maximizing) EUT agents will believe propositions only if their credence is above the threshold $\frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}$, the

²⁴We have seen very few *arguments in favor of* Vacuity (or, for that matter, (P2) or Subexpansion). One such argument is due to Gärdenfors [14]. Ironically, that argument makes use of an analogy between belief revision and Bayesian conditionalization — an analogy that is undermined by our EUT counterexample above. Moreover, in a recent paper, Lin & Kelly [24] give arguments *against* Vacuity which are also based on its deviation from “Bayesian updating” (broadly construed). However, the alternative to AGM that is offered by Lin & Kelly (which is based on odds-ratio thresholds as opposed to conditional probability thresholds) is not in the spirit of EUT. Specifically, their revision procedure (a) is partition-sensitive, and (b) allows propositions with arbitrarily low credence to be believed by the agent. Neither of these properties is compatible with EUT. Ultimately, the reason their approach deviates from EUT in these ways is that they insist on **Cogency** as a (universal) epistemic rational requirement on belief sets. But, setting **Cogency** aside, we sympathize with Lin & Kelly’s objections to AGM. However, we think our counterexamples are more direct and probative. Moreover, our discussion below explains the precise conditions under which Lockean agents can violate Vacuity. These conditions (which, somewhat mysteriously, involve the *Golden Ratio*) were not previously known.

²⁵We already saw that EUT Lockean agents can violate Vacuity by learning something *they already believe*. These are violations of Leitgeb’s (P2). And, as our discussion of (P2) above suggested, *those* sorts of Vacuity violations can happen for *any* Lockean threshold $t \geq 1/2$.

constraint above implies that $\phi - 1 \leq \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}} < 1$, which is equivalent to (+). And, that is the constraint on \mathfrak{w} and \mathfrak{r} which is compatible with (all possible kinds of) violations of AGM principles by deductively cogent, Lockean agents.

5.3. **Non-Extremal EUT Revision is More Conservative than AGM Revision (when the two approaches interestingly diverge).** Interesting divergences between EUT and AGM revision have another crucial feature, which is also revealed by our counterexample to Vacuity. Because EUT satisfies Inclusion, we know that it never requires the agent to acquire *more* new beliefs than AGM permits. More precisely, because EUT and AGM both satisfy Inclusion, they both *rule out* posterior belief sets $\mathbf{B}' = \mathbf{B} \star E$ that are *proper supersets* of $\text{Cn}(\mathbf{B} \cup \{E\})$. In other words, neither EUT nor AGM will ever require an agent to be committed to new beliefs that *go beyond* the logical consequences of $\mathbf{B} \cup \{E\}$.

On the other hand, it can be shown that in all (interesting) cases of disagreement between EUT and AGM, EUT will require the agent to have *strictly fewer* new beliefs than are mandated by AGM. For example, consider our EUT-counterexample to Vacuity, above (see Table 1). In this case, we know that EUT revision requires the agent to give up some beliefs that AGM revision does not. Specifically, in our EUT-counterexample to Vacuity, we know that AGM’s posterior must include $X \vee \neg E$. So, when E is learned, AGM revision will (by Closure) result in the agent believing X , $X \wedge E$ and $X \equiv E$, while EUT revision will *preclude* these new beliefs. And, there is nothing special about this particular example. *Whenever* EUT and AGM disagree *in an interesting way* (*i.e.*, because EUT violates Vacuity), EUT will require *strictly fewer* new beliefs than AGM. And, the converse holds as well. That is, we have the following (final) theorem, which asserts that EUT violates Vacuity *if and only if* EUT requires *strictly fewer* new beliefs than AGM (provided E is consistent with \mathbf{B}).

Theorem *EUT violates Vacuity (wrt \mathbf{B}, E) $\Leftrightarrow E$ is consistent with \mathbf{B} and $\mathbf{B} \star E \subset \mathbf{B} \star E$.*

Proof. (\Rightarrow) Suppose EUT violates Vacuity (wrt \mathbf{B} and E). Then, (a) E is consistent with \mathbf{B} ; and, (b) $\mathbf{B} \star E \not\subseteq \text{Cn}(\mathbf{B} \cup \{E\})$. By (b), there exists an X such that $X \in \text{Cn}(\mathbf{B} \cup \{E\})$ but $X \notin \mathbf{B} \star E$. It follows from (a), Vacuity and Inclusion that $\text{Cn}(\mathbf{B} \cup \{E\}) = \mathbf{B} \star E$. Therefore, $X \in \mathbf{B} \star E$ and $X \notin \mathbf{B} \star E$. And, by Inclusion, $\mathbf{B} \star E \subseteq \text{Cn}(\mathbf{B} \cup \{E\}) = \mathbf{B} \star E$. \square

(\Leftarrow) Suppose E is consistent with \mathbf{B} and $\mathbf{B} \star E \subset \mathbf{B} \star E$. Then, there exists an X such that $X \in \mathbf{B} \star E$ but $X \notin \mathbf{B} \star E$. Because E is consistent with \mathbf{B} , Vacuity and Inclusion imply that $\mathbf{B} \star E = \text{Cn}(\mathbf{B} \cup \{E\})$. Therefore, $X \in \text{Cn}(\mathbf{B} \cup \{E\})$; but, $X \notin \mathbf{B} \star E$. \square

In other words, when EUT and AGM (interestingly) diverge, AGM will be *more demanding* on an agent’s beliefs (insofar as they are maintained *via* revision). Since AGM will require agents to maintain beliefs in the face of counter-evidence (such as in our counter-example to Vacuity), it may be seen as an epistemically risk-seeking

policy for belief revision.²⁶ On the other hand, EUT will recommend that agents suspend belief in many cases and so it may be seen as epistemically risk-averse.²⁷

Moreover, the cases in which AGM is more demanding than EUT in this sense are *precisely* those cases in which EUT violates Vacuity. Thus, Vacuity is truly at the heart of the (interesting) divergences between EUT and AGM revision.

6. CONCLUSION AND FUTURE WORK

We have pinpointed the precise ways in which a (broadly Bayesian) EUT approach to belief revision agrees (and disagrees) with the more traditional AGM theory of belief revision. Setting aside issues surrounding deductive cogency as a rational requirement for belief sets, EUT revision and AGM revision exhibit a surprising degree of convergence. Our analysis reveals that, once we hold **Cogency** (and Idempotence) fixed, the two approaches to belief revision disagree *only* regarding the universal validity of Vacuity (and, hence, of Subexpansion and Leitgeb's (P2)).

One common complaint made against AGM is that it fails to easily accommodate iterated revisions. On the other hand, EUT has no special problem with iterated revision.²⁸ In future work, we plan to compare EUT revision to other systems of belief revision aside from AGM (especially, ones that seem better suited than AGM to handle *iterated* revision). In this connection, it will be of particular interest to compare EUT with the Darwiche and Pearl postulates for iterated revision [4].

Another interesting next step in the exploration of Bayesian qualitative revision is to investigate how EUT revision changes when the credence function is a non-classical probability function (and so would allow for conditionalizing on propositions with zero unconditional credence) or when conditionalization is replaced by other, more general credal updating procedures.²⁹ One especially interesting

²⁶Pettigrew [26] argues that **Cogency** may be aptly analyzed as an epistemically risk-seeking synchronic constraint on belief. The connection between his analysis and ours is suggestive. Since we think of AGM's requirement of **Cogency** as a synchronic constraint, Pettigrew's analysis should have led us to expect that AGM can be seen as more epistemically risk-seeking in its *synchronic* requirements; however, our result goes beyond those expectations and shows that AGM is *diachronically* (more) risk-seeking as well.

²⁷It is worth pointing out that EUT's risk-aversion should not be wholly surprising since it is driven by the expected utility calculus given the specified concave utility function. Nonetheless, it is interesting to notice that in full generality, non-extremal EUT revision is more risk-averse than AGM.

²⁸That said, insofar as we have relied on conditionalization to define \ast , there is a problem with conditionalizing on any proposition assigned a prior probability of 0 (*fn.* 22).

²⁹As we mentioned in *fn.* 8, all of the results we reported here will continue to hold for any mechanical/minimal change Bayesian credal update procedure that satisfies the following three constraints: (i) $b'(E) > b(E)$, (ii) $b'(E) > t$ (where t is the agent's EUT Lockean threshold), and (iii) $b(E \supset X) \geq b'(X)$. It would be nice to explore these (and other) non-standard Bayesian updating procedures in more depth (especially, in conjunction with Lockeanism).

application along these lines would be to the problem of explicating a Bayesian notion of *contraction*. We have some preliminary ideas about "Bayesian contraction," which we plan to explore in a sequel to this paper.³⁰

Finally, we would (ideally) like to have a *purely qualitative axiomatization* of the EUT revision operator. Some progress toward such an axiomatization has recently been made [25, 9]. However, much theoretical work remains to be done here, in order to determine precisely which axioms are needed to characterize EUT revision (over and above the axioms on which EUT revision and AGM revision agree, and setting aside deductive cogency).

APPENDIX: Proof of Proposition 8

Proposition 8 If an EUT/Lockean agent is deductively cogent (at all times), then they can only violate Vacuity (*via* learning some E that they do not already believe) if their Lockean threshold is on the half-open interval $[\phi - 1, 1)$.

Proof. It is well-known that if a Lockean agent's threshold is *extremal* (*i.e.*, equal to 1), then such an agent will (a) have deductively cogent beliefs, and (b) satisfy Vacuity. This explains why the interval $[\phi - 1, 1)$ is open on the right.

The interesting (and new) case involves the left-hand (closed) side of the interval. To demonstrate that deductively cogent, Lockean Vacuity violators (like our agent above) must have Lockean thresholds of *at least* $\phi - 1$, we proceed *via* two lemmas.

Lemma 1 Let $\langle W, \text{Pr} \rangle$ be a (finite) probability space, over a set of possible worlds W , and suppose that the three propositions $\{X, Y, Z\}$ form a *partition* of W .³¹ Then, the following four conditions:

- (1) $\text{Pr}(X \vee Y) > t$.
- (2) $\text{Pr}(X \vee Z) \leq t$.
- (3) $\text{Pr}(Y \vee Z) \leq t$.
- (4) $\text{Pr}(X \mid X \vee Z) \leq t$.

jointly entail

$$(5) \quad t \geq \frac{\sqrt{5} - 1}{2} = \phi - 1$$

³⁰The basic idea behind our approach to "contracting a Bayesian belief set \mathbf{B} on proposition p " would involve (a) defining b' as *the closest probability function to b such that $b'(p) \leq t$* , and then (b) checking which propositions x are such that $b'(x) > t$. The set $\mathbf{B} \div p := \{x \mid b'(x) > t\}$ would be our (initial) explication of what it means to "contract a Bayesian agent's belief set \mathbf{B} on proposition p ."

³¹That is, the disjunction $X \vee Y \vee Z$ is a tautology (*i.e.*, a proposition that's true in all members of W), and each pair from $\{X, Y, Z\}$ is a contradiction (*i.e.*, a proposition that's false in all members of W).

Proof. Suppose conditions (1)–(4) hold. We adopt the following abbreviations.

$$x := \Pr(X)$$

$$y := \Pr(Y)$$

$$z := \Pr(Z)$$

Conditions (1)–(4) then may be re-written as follows

$$(1) \ x + y > t.$$

$$(2) \ x + z \leq t.$$

$$(3) \ y + z \leq t.$$

$$(4) \ \frac{x}{x+z} \leq t.$$

Because $\{X, Y, Z\}$ form a partition of W , it follows from probability calculus that $x, y, z \in [0, 1]$ and $z = 1 - x - y$. Substitution for z in (1)–(4) then yields

$$(1) \ x + y > t.$$

$$(2) \ 1 - y \leq t.$$

$$(3) \ 1 - x \leq t.$$

$$(4) \ \frac{x}{1-y} \leq t.$$

It is then simply a matter of elementary algebraic reasoning to derive that

$$t^2 + t - 1 \geq 0$$

from which it follows (since $t \geq 0$) that

$$(5) \quad t \geq \frac{\sqrt{5} - 1}{2} = \phi - 1$$

which completes the proof of Lemma 1. \square

Lemma 2 Let $\langle W, \Pr \rangle$ be a (finite) probability space, over a set of possible worlds W , and let $t \geq 0$ be an EUT agent's Lockean threshold. Suppose that the agent's prior belief set \mathbf{B} is deductively cogent, and let E be some proposition consistent with \mathbf{B} , and which is not already contained in \mathbf{B} . If such an agent's EUT-updated belief set $\mathbf{B} * E$ is deductively cogent, but yields a violation of Vacuity, then there exists a partition $\{X, Y, Z\}$ of W such that the following four conditions are met:

$$(1) \ \Pr(X \vee Y) > t.$$

$$(2) \ \Pr(X \vee Z) \leq t.$$

$$(3) \ \Pr(Y \vee Z) \leq t.$$

$$(4) \ \Pr(X \mid X \vee Z) \leq t.$$

Proof. Because \mathbf{B} is cogent, there is a proposition C (the conjunction of all the members of \mathbf{B}) which is a member of \mathbf{B} and which entails every member of \mathbf{B} . Let

$$X := E \& C$$

$$Y := \neg E \& C$$

$$Z := \neg C$$

Clearly, $\{X, Y, Z\}$ constitute a partition of W . So, we just need to show that $\{X, Y, Z\}$ satisfy conditions (1)–(4) above.

(1) $\Pr(X \vee Y) > t$. To see this, note that $X \vee Y \models C$, and $C \in \mathbf{B}$ by hypothesis.

(2) $\Pr(X \vee Z) \leq t$. To see this, note that $X \vee Z \models C \supset E$. Suppose for *reductio* that $\Pr(X \vee Z) > t$. Then, $C \supset E \in \mathbf{B}$, since $\Pr(C \supset E) > t$. But, since $C \in \mathbf{B}$ and \mathbf{B} is cogent, this would imply (*via modus ponens*) that $E \in \mathbf{B}$, which contradicts one of the preconditions of the lemma (that $E \notin \mathbf{B}$).

(3) $\Pr(Y \vee Z) \leq t$. To see this, note that $Y \vee Z \models C \supset \neg E$. Suppose for *reductio* that $\Pr(Y \vee Z) > t$. Then, $C \supset \neg E \in \mathbf{B}$, since $\Pr(C \supset \neg E) > t$. But, since $C \in \mathbf{B}$ and \mathbf{B} is cogent, this would imply (*via modus ponens*) that $\neg E \in \mathbf{B}$, which contradicts one of the preconditions of the lemma (that $\mathbf{B} \neq \neg E$).

(4) $\Pr(X \mid X \vee Z) \leq t$. To see this, note that

$$\begin{aligned} \Pr(X \mid X \vee Z) &= \frac{\Pr(X \& (X \vee Z))}{\Pr(X \vee Z)} \\ &= \frac{\Pr(X)}{\Pr(X \vee Z)} \\ &= \frac{\Pr(E \& C)}{\Pr(\neg C \vee E)} \\ &\leq \frac{\Pr(C \& E)}{\Pr(E)} \end{aligned}$$

Thus, $\Pr(X \mid X \vee Z) \leq \Pr(C \mid E)$. Moreover, $\Pr(C \mid E) \leq t$. To see this, suppose (for *reductio*) that $\Pr(C \mid E) > t$. Then, $C \in \mathbf{B} * E$. But, this implies that *both* C and E are members of $\mathbf{B} * E$. And, by the cogency of $\mathbf{B} * E$, this implies $\mathbf{B} * E \supseteq \text{Cn}(\mathbf{B} \cup \{E\})$, which contradicts one of the preconditions of the lemma (that this is a counterexample to Vacuity). Therefore, $\Pr(X \mid X \vee Z) \leq \Pr(C \mid E)$ and $\Pr(C \mid E) \leq t$, which implies $\Pr(X \mid X \vee Z) \leq t$.

This completes the proof of Lemma 2. \square

Proposition 8 is an immediate corollary of Lemma 1 and Lemma 2. \square

REFERENCES

- [1] Carlos Alchourron, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal*

- of *Symbolic Logic*, 50(2): 510–530, June 1985.
- [2] Horacio Arló Costa. Conditionals and Monotonic Belief Revisions: the Success Postulate. *Studia Logica*, 49: 557–566, 1990.
- [3] David Christensen. Putting Logic in its Place. Oxford University Press, 2004.
- [4] Adnan Darwiche and Judea Pearl. On the Logic of Iterated Belief Revision *Artificial Intelligence*, 89: 1–29, 1996.
- [5] Persi Diaconis and Sandy Zabell. Updating subjective probability *Journal of the American Statistical Association* 77: 822–830, 1982.
- [6] Kevin Dorst. An Epistemic Utility Argument for the Threshold View of Outright Belief. Manuscript, 2014.
- [7] Kenny Easwaran. Dr. Truthlove, or how I learned to stop worrying and love Bayesian probabilifies. *Noûs*, Forthcoming.
- [8] Kenny Easwaran and Branden Fitelson. Accuracy, Coherence, and Evidence. *Oxford Studies in Epistemology*, Vol. 5. Oxford University Press, to appear.
- [9] Jan van Eijck and Bryan Renne. Belief as Willingness to Bet. Manuscript, 2014.
- [10] Richard Foley. Working Without a Net. Oxford University Press, 1992.
- [11] Peter Gärdenfors, Sten Lindström, Michael Morreau, and Włodzimierz Rabinowicz. The negative Ramsey test: Another triviality result. In André Fuhrmann and Michael Morreau (eds.), *The Logic of Theory Change*, pages 127–134. Springer, 1991.
- [12] Peter Gärdenfors and Hans Rott. Belief revision. *Handbook of logic in artificial intelligence and logic programming*, Vol. 4, pages 35–132. Oxford University Press, 1995.
- [13] Peter Gärdenfors. Conditionals and Changes of Belief. *Acta Philosophica Fennica*, 30: 381–404, 1978.
- [14] Peter Gärdenfors. Belief Revisions and the Ramsey Test for Conditionals. *The Philosophical Review*, 95(1): 81–93, 1986.
- [15] Peter Gärdenfors. The dynamics of belief: Contractions and revisions of probability functions. *Topoi*, 5: 29–37, 1986.
- [16] Peter Gärdenfors and David Makinson. Revision of Knowledge Systems Using Epistemic Entrenchment *TARK '88 Proceedings of the Second Conference of Theoretical Aspects of Reasoning about Knowledge*, 83–95, 1988.
- [17] William Harper. Rational belief change, popper functions and counterfactuals. *Synthese*, 30(1-2): 221–262, 1975.
- [18] James Hawthorne. A primer on rational consequence relations, popper functions, and their ranked structure. *Studia Logica*, 104: 1–18, 2013.
- [19] Carl Hempel. Deductive-nomological vs. statistical explanation. *Minnesota studies in the philosophy of science* 3: 98–169, 1962.
- [20] William James. The will to believe. *The New World*, 5: 327–347, 1896.

- [21] Henry Kyburg. Probability and The Logic of Rational Belief. Wesleyan University Press, 1961.
- [22] Hannes Leitgeb. The Stability Theory of Belief. *Philosophical Review*, 123(2): 131–171, 2014.
- [23] Hannes Leitgeb. The review paradox: On the diachronic costs of not closing rational belief under conjunction. *Noûs*, 78(4): 781–793, 2013.
- [24] Hanti Lin and Kevin Kelly. Propositional Reasoning that Tracks Probabilistic Reasoning. *Journal of Philosophical Logic*, 41(6): 957–981, 2012.
- [25] David Makinson and James Hawthorne. Lossy Inference Rules and their Bounds: A Brief Review. *The Road to Universal Logic: Festschrift for 50th Birthday of Jean-Yves Béziau*, Vol. 1, pages 385–407, Springer, 2015.
- [26] Richard Pettigrew. Accuracy and Risk. Manuscript, 2015.
- [27] Richard Pettigrew. Epistemic Utility Arguments for Probabilism. *The Stanford Encyclopedia of Philosophy* (Winter 2011 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2011/entries/epistemic-utility/>.
- [28] Alexander Pruss. The badness of being certain of a falsehood is at least $\frac{1}{\log(4)-1}$ times greater than the value of being certain of a truth. *Logos & Episteme*, 3(2): 229–238, 2012.
- [29] Odinaldo Rodrigues, Dov Gabbay and Alessandra Russo. Belief Revision. *Handbook of Philosophical Logic v.16*, D. Gabbay and F. Guenther (eds.), pages 1–114, 2011
- [30] Robert Stalnaker. Iterated belief revision. *Erkenntnis*, 70(2):189–209, 2009.
- [31] Florian Steinberger. Explosion and the normativity of logic. *Mind*, to appear.

PHILOSOPHY DEPARTMENT, UNIVERSITY OF CALIFORNIA, DAVIS

E-mail address: ebshear@ucdavis.edu

URL: <http://www.tedshear.com/>

PHILOSOPHY DEPARTMENT, UNIVERSITY OF TORONTO

E-mail address: jonathan.weisberg@utoronto.ca

URL: <http://www.utm.utoronto.ca/~weisber3/>

PHILOSOPHY DEPARTMENT, NORTHEASTERN UNIVERSITY, BOSTON

E-mail address: branden@fitelson.org

URL: <http://fitelson.org/>