Propositional Reasoning that Tracks Probabilistic Reasoning

Hanti Lin (hantil@andrew.cmu.edu)
Carnegie Mellon University

Kevin T. Kelly (kk3n@andrew.cmu.edu)*
Carnegie Mellon University

Abstract. This paper concerns the extent to which propositional reasoning can track probabilistic reasoning, which addresses kinematic problems that extend the familiar Lottery paradox. An acceptance rule (Leitgeb 2010) assigns to each Bayesian credal state p a propositional belief revision method B_p, which specifies an initial belief state $\mathsf{B}_p(\top)$, that is revised into the new propositional belief state $\mathsf{B}(E)$ upon receipt of information E. The acceptance rule *tracks* Bayesian conditioning when $\mathsf{B}_p(E) = \mathsf{B}_{p|_E}(\top)$, for every E such that $p(E) > 0$; namely, when acceptance by propositional belief revision equals Bayesian conditioning followed by acceptance. Standard proposals for acceptance and belief revision do not track Bayesian conditioning. The "Lockean" rule that accepts propositions above a probability threshold is subject to the familiar lottery paradox (Kyburg 1961), and we show that it is also subject to new and more stubborn paradoxes when the tracking property is taken into account. Moreover, we show that the familiar AGM approach to belief revision (Harper 1975 and Alchourrn, Grdenfors, and Makinson 1985) cannot be realized in a sensible way by an acceptance rule that tracks Bayesian conditioning. Finally, we present a plausible, alternative approach that tracks Bayesian conditioning and avoids all of the paradoxes. It combines an odds-based acceptance rule proposed originally by Levi (1996) with a non-AGM belief revision method proposed originally by Shoham (1987). As an application, the Lottery paradox turns out to receive a new solution motivated by dynamic concerns.

Keywords: Uncertain acceptance, Lottery paradox, Belief revision, Bayesian conditioning, Gettier problem

1. An Old Riddle of Uncertain Acceptance

There are two widespread practices for modeling the doxastic state of a subject—as a probability measure over propositions or as a single proposition corresponding to the conjunction of all propositions the subject believes. One straightforward way to relate propositional belief to probabilistic belief is to accept only propositions of probability one. However, that skeptical approach severely restricts the scope and practical relevance of propositional reasoning, so it is natural to seek a more

* This work was supported generously by the National Science Foundation under grant 0750681 and the John Templeton Foundation under grant 24145. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or the John Templeton Foundation.

liberal standard for acceptance. One natural idea, called the *Lockean rule* in honor of John Locke, who proposed something like it, is to accept all and only the logical consequences of the set of all *sufficiently* probable propositions, whose probabilities are no less than some fixed threshold t strictly less than one.

Alas, however the threshold for acceptance is set, the Lockean rule leads to acceptance of inconsistency, a difficulty known as the *lottery paradox* (Kyburg 1961). Suppose that the threshold is 2/3. Now consider a fair lottery with 3 tickets. Then the degree of belief that a given ticket loses is 2/3, so it is accepted that each ticket loses. That entails that no ticket wins. Furthermore, with probability one *some* ticket wins, so that proposition is also accepted. Then, the conjunction of the accepted propositions is contradictory. In general, if t is the threshold, a lottery with more than $1/(1 - t)$ tickets suffices for acceptance of inconsistency.

2. Two New Riddles of Uncertain Acceptance

The lottery paradox concerns static consistency. But there is also the kinematic question of how to revise one's propositional belief state in light of new evidence or suppositions. Probabilistic reasoning has its own, familiar revision method, namely, Bayesian conditioning. Mismatches between propositional belief revision and Bayesian conditioning are another potential source of conundrums for uncertain acceptance. Unlike the lottery paradox, these riddles cannot be avoided by the expedient of raising the probabilistic standard for acceptance to a sufficiently high level short of full belief.

For the first riddle, suppose that there are three tickets and consider the Lockean acceptance rule with threshold 3/4, at which the lottery paradox is easily avoided. Suppose further that the lottery is not fair: ticket 1 wins with probability 1/2 and tickets 2 and 3 win with probability 1/4. Then it is just above the threshold that ticket 2 loses and that ticket 3 loses, which entails that ticket 1 wins. Now entertain the new information that ticket 3 has been removed from the lottery, so it cannot win. Since ruling out a competing ticket seems only to provide further evidence that ticket 1 will win, it is strange to then retract one's belief that ticket 1 wins. But the Lockean rule does just that. By Bayesian conditioning, the probability that ticket 3 wins is reset to 0 and the odds between tickets 1 and 2 remain 2:1, so the probability that ticket 1 wins is 2/3. Therefore, it is no longer accepted that ticket 1 wins, since that proposition is neither sufficiently probable by itself nor entailed by a set of sufficiently probable propositions, where sufficient probability means probability no less than 3/4.

It is important to recognize that the first riddle is *geometrical* rather than logical (figure 1). Let H_1 be the proposition that ticket 1 wins, and

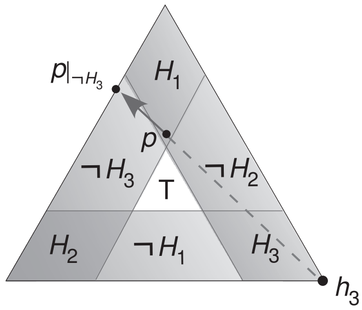

Figure 1. the first riddle

similarly for H_2, H_3. The space of all probability distributions over the three tickets consists of a triangle in the Euclidean plane whose corners have coordinates $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 1)$, which are the extremal distributions that concentrate all probability on a single ticket. The assumed distribution p over tickets then corresponds to the point $p = (1/2, 1/4, 1/4)$ in the triangle. The conditional distribution $p|_{\neg H_3} = p(\cdot|\neg H_3)$ is the point $(2/3, 1/3, 0)$, which lies on a ray through p that originates from corner 3, holding the odds $H_1 : H_2$ constant. Each zone in the triangle is labeled with the strongest proposition accepted at the probability measures inside. The acceptance zone for H_1 is a parallel-sided diamond that results from the intersection of the above-threshold zones for $\neg H_2$ and $\neg H_3$, since it is assumed that the accepted propositions are closed under conjunction. The rule leaves the inner triangle as the acceptance zone for the tautology $\top$. The riddle can now be seen to result from the simple, geometrical fact that p lies near the point of the diamond, which is so skinny that conditioning carries p outside of the diamond. If the bottom of the diamond were more blunt, to match the slope of the conditioning ray, then the paradox would not arise.

The riddle can be summarized by saying that the Lockean rule fails to satisfy the following, diachronic principle for acceptance: accepted beliefs are not to be retracted when their logical consequences are learned. Assuming that accepted propositions are closed under entailment, let B_p denote the strongest proposition accepted in probabilistic credal state p. So H is accepted at p if and only if $\mathsf{B}_p \models H$. Then the principle may be stated succinctly as follows, where $p|_E$ denotes the conditional distribution $p(\cdot|E)$:

$$\mathsf{B}_p \models H \text{ and } H \models E \implies \mathsf{B}_{p|_E} \models H. \tag{1}$$

Philosophers of science speak of *hypothetico-deductivism* as the view that observing a logical consequence of a theory provides evidence in favor of the theory. Since it would be strange to retract a theory in light of new, positive evidence, we refer to the proposed principle as *Hypothetico-deductive Monotonicity.*

One Lockean response to the preceding riddle is to adopt a higher threshold for disjunctions than for conjunctions (figure 2) so that the acceptance zone for H_1 is closed under conditioning on $\neg H_3$. But now

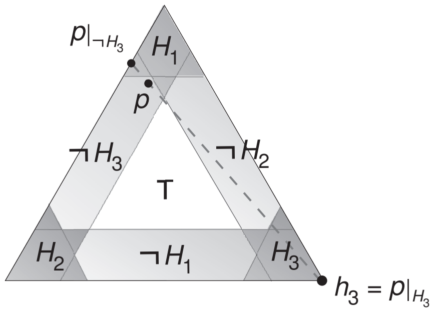

Figure 2. second riddle

a different and, in a sense, complementary riddle emerges. For suppose that the credal state is p, just inside the zone for accepting that either ticket 1 or 2 will win and close to, but outside of the zone for accepting that ticket 1 will win. Now, the rule accepts that ticket 2 loses no matter whether one learns that ticket 3 wins (i.e. p moves to $p|_{H_3}$) or leans its negation (i.e. p moves to $p|_{\neg H_3}$), but fails to accept that ticket 2 loses until one actually learns what happens with ticket 3. That violates the following principle:[1]

$$\mathsf{B}_{p|_E} \models H \text{ and } \mathsf{B}_{p|_{\neg E}} \models H \implies \mathsf{B}_p \models H, \tag{2}$$

[1] The principle is analogous in spirit to the *reflection principle* (van Fraassen 1984), which, in this context, might be expressed by saying that if you know that you will accept a proposition regardless what you learn, you should accept it already. Also, a non-conglomerable probability measure has the feature that some B is less probable than it is conditional on each H_i. Schervish, Seidenfeld, and Kadane (1984) show that every finitely additive measure is non-conglomerable in some partition. In that case, any sensible acceptance rule would fail to satisfy reasoning by cases. Some experts advocate finitely additive probabilities and others view non-conglomerability as a paradoxical feature. For us, acceptance is relative to a partition (question), a topic we discuss in detail in Lin and Kelly (forthcoming), so non-conglomerability does not necessarily arise in the given partition.

which we call *Case Reasoning*.

The two new riddles add up to one big riddle: there is, in fact, *no* ad hoc manipulation of distinct thresholds for distinct propositions that avoids both riddles.[2] The first riddle picks up where the second riddle leaves off and there are even thresholds that generate both riddles at once. Unlike the lottery paradox, which must increase the number of tickets as the Lockean threshold is raised, one of the two new riddles obtains for every possible combination of thresholds, as long as there are at least three tickets and the thresholds have values less than one. So although it may be tempting to address the lottery paradox by raising the thresholds in response to the number of tickets, even that possibility is ruled out by the new riddles. All of the Lockean rules have the wrong *shape*.

3. The Propositional Space of Reasons

Part of what is jarring about the riddles is that they undermine one of the most plausible motives for considering acceptance at all: reasoning directly with propositions, without having to constantly consult the underlying probabilities. In the first riddle, observed logical consequences H result in rejection of H. In the second riddle, propositional reasoning by cases fails so that, for example, one could not rely on logic to justify policy (e.g., the policy achieves the desired objective in any case). Although one accepts propositions, the riddles witness that one has not really entered into a purely propositional "space of reasons" (Sellars 1956). The accepted propositions are mere, epiphenomenal shadows cast by the underlying probabilities, which evolve according to their own, more fundamental rules. Full entry into a propositional space of reasons demands a tighter relationship between acceptance and probabilistic conditioning.

The riddles would be resolved by an improved acceptance rule that allows one to enter the propositional system, kick away the underlying probabilities, and still end up exactly where a Bayesian conditionalizer would end up—i.e., by an acceptance rule that realizes a pre-established harmony between propositional and probabilistic reasoning. The realization of such a perfect harmony, without peeking at the underlying probabilities, is far more challenging than merely to avoid acceptance of mutually inconsistent propositions. This ideal will be shown to be impossible to achieve if one insists on employing the popular AGM approach to propositional belief revision. Then, with a distinct approach

[2] The claim is a special case of theorem 3 in Lin and Kelly (forthcoming).

to belief revision, we exhibit a broad collection of rules that do achieve perfect harmony with Bayesian conditioning.

4. Questions, Answers, and Credal States

Let $\mathcal{Q} = \{H_i : i \in I\}$ be a countable collection of mutually exclusive and exhaustive propositions, representing a *question* to which $H_1, \ldots, H_i, \ldots$ are the *(complete) answers*. Let $\mathcal{A}$ be the set of all complete and incomplete answers to $\mathcal{Q}$, i.e. the σ-algebra containing $\mathcal{Q}$. Let $\mathcal{P}$ denote the set of all countably additive probability measures on $\mathcal{A}$, which will be referred to as *credal states*. In the three-ticket lottery, for example, $\mathcal{Q} = \{H_1, H_2, H_3\}$, H_i says that ticket i wins, and $\mathcal{P}$ is the triangle (simplex) of probability distributions over the three answers.

5. Belief Revision

A *belief state* is just a deductively closed set of propositions; but for the sake of convenience, in this paper we always identity a belief state with the strongest proposition in it. A *belief revision method* is a mapping $\mathsf{B} : \mathcal{A} \to \mathcal{A}$, understood as specifying the initial belief state $\mathsf{B}(\top)$, which would evolve into new belief state $\mathsf{B}(E)$ upon revision on information E.[3] *Hypothetico-deductive Monotonicity*, for example, can now be stated in terms of belief revision, rather than in terms of Bayesian conditioning:[4]

$$\mathsf{B}(\top) \models H \text{ and } H \models E \implies \mathsf{B}(E) \models H. \tag{3}$$

Case Reasoning has a similar statement:[5]

$$\mathsf{B}(E) \models H \text{ and } \mathsf{B}(\neg E) \models H \implies \mathsf{B}(\top) \models H. \tag{4}$$

[3] Readers more familiar with the belief revision operator notation $*$ (Alchourrn, Grdenfors, and Makinson 1985) may employ the translation rule: $\mathsf{B}(\top) * E = \mathsf{B}(E)$. Note that $\mathsf{B}(\top)$ is understood as the initial belief state rather than revision on the tautology.

[4] Hypothetico-deductive Monotonicity is strictly weaker than the principle called *Cautious Monotonicity* in the nonmonotonic logic literature: $\mathsf{B}(X) \models Y$ and $\mathsf{B}(X) \models Z \Longrightarrow \mathsf{B}(X \wedge Z) \models Y$.

[5] Case Reasoning is an instance of the principle called *Or* in the nonmonotonic logic literature: $\mathsf{B}(X) \models Z$ and $\mathsf{B}(Y) \models Z \Longrightarrow \mathsf{B}(X \vee Y) \models Z$.

6. When Belief Revision Tracks Bayesian Conditioning

A credal state represents not only one's degrees of belief but also how they should be updated according to the Bayesian ideal. So the qualitative counterpart of a credal state should be an initial belief state *plus* a qualitative strategy for revising it. Accordingly, define an *acceptance rule* to be a mapping that assigns to each credal state p a belief revision method B_p, which should be written as $(\mathsf{B}_p : p \in \mathcal{P})$ but will be abbreviated as B by abusing notation. (Think of B as a mapping $\mathsf{B}_{(\,.\,)}$ that sends p to B_p.) Note that $\mathsf{B}_p(\top)$ is the initial belief state that the subject accepts at credal state p; accordingly, say that proposition H is *accepted* by rule B at credal state p if and only if $\mathsf{B}_p(\top) \models X$.[6]

Each revision allows for a choice between two possible courses of action, starting at credal state p. According to the first course of action, the subject accepts propositional belief state $\mathsf{B}_p(\top)$ and then *propositionally* revises it to obtain the new propositional belief state $\mathsf{B}_p(E)$ (i.e., the left-lower path in figure 3). According to the second

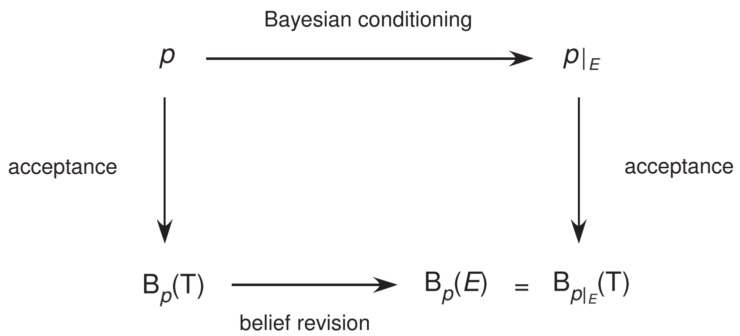

Figure 3. Belief revision tracks Bayesian conditioning

course of action, she first conditions p to obtain the posterior credal state $p|_E$ and then accepts $\mathsf{B}_{p|_E}(\top)$ (i.e., the upper-right path in figure 3). According to the pre-established harmony, the two processes should

[6] The following, conditional acceptance *Ramsey tests* translates our framework into notation familiar in the logic of epistemic conditionals:

$$p \Vdash E \Rightarrow H \iff \mathsf{B}_p(E) \models H; \qquad (5)$$

$$E \mathrel{|\!\!\sim}_p H \iff \mathsf{B}_p(E) \models H. \qquad (6)$$

We are indebted to Hannes Leitgeb for the idea of framing our discussion in terms of acceptance of belief revision methods, at the Opening Celebration of the Center for Formal Epistemology at Carnegie Mellon University. Our own approach (Lin and Kelly (forthcoming)), prior to seeing his work, was to formulate the issues in terms of conditional epistemic logic, via a probabilistic Ramsey test, which involves more cumbersome notation and an irrelevant commitment to an epistemic interpretation of conditionals.

always agree (i.e., the diagram should always commute). Accordingly, say that acceptance rule B *tracks conditioning* if and only if:

$$\mathsf{B}_p(E) \ = \ \mathsf{B}_{p|_E}(\top), \tag{7}$$

for each credal state p and proposition E such that $p(E) > 0$. In short, acceptance followed by belief revision equals Bayesian conditioning followed by acceptance.

7. Accretive Belief Revision

It remains to specify what would count as a propositional approach to belief revision that does not essentially peek at probabilities to decide what to do. An obvious and popular idea is simply to conjoin new information with one's old beliefs to obtain new beliefs, as long as no contradiction results. This idea is usually separated into two parts: say that belief revision method B satisfies *Inclusion* if and only if:[7]

$$\mathsf{B}(\top) \wedge E \models \mathsf{B}(E); \tag{8}$$

say that method B satisfies *Preservation* if and only if:

$$\mathsf{B}(\top) \text{ is consistent with } E \implies \mathsf{B}(E) \models \mathsf{B}(\top) \wedge E. \tag{9}$$

These axioms are widely understood to be the least controversial axioms in the much-discussed AGM theory of belief revision, due to Harper (1975) and Alchourrn, Grdenfors, and Makinson (1985). A belief revision method is *accretive* if and only if it satisfies both Inclusion and Preservation. An acceptance rule is *accretive* if and only if each belief revision method B_p it assigns is accretive.

8. Sensible, Tracking Acceptance Cannot Be Accretive

Accretion sounds plausible enough when beliefs are certain, but it is not very intuitive when beliefs are accepted at probabilities less than 1. For example, suppose that we have two friends—Nogot and Havit—and we know for sure that at most one owns a Ford. Question: who owns a Ford? Three potential answers: "Nogot" vs. "Havit" vs. "nobody" (figure 4). Now, Nogot shows us car keys and his driver's license and Havit does nothing, so we think that it is pretty probable that Nogot

[7] Inclusion is equivalent to Case Reasoning, assuming the axiom called *Success*: $\mathsf{B}(E) \models E$.

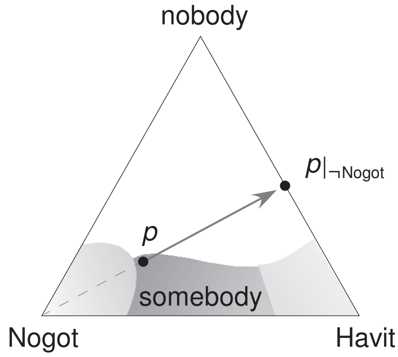

Figure 4. How Preservation may fail plausibly

has a Ford (i.e., credal state p is close to the corner for “Nogot”). Given the evidence we have, suppose that we want to be more judicious, only accepting the disjunction of “Nogot” with another answer. Suppose, further, that “Havit” is a bit more probable than “nobody” (i.e., credal state p is a bit closer to the “Havit” corner than to the “nobody” corner). So the strongest proposition we accept is the disjunction of “Nogot” with “Havit”, namely “somebody” (i.e., credal state p falls in the acceptance zone for “somebody”). Unfortunately, Nogot was only pretending to own a Ford. Suppose that now we learn the negation of “Nogot”. What would we accept then? Note that the new information “¬Nogot” undermines the main reason (i.e., “Nogot”) for accepting the old belief state “somebody”, in spite of that fact that the new information is compatible with the old belief state. So it seems plausible to drop the old belief in the new belief state, namely, to violate the Preservation axiom in this case. That intuition agrees with Bayesian conditioning: the posterior credal state $p|_{\neg\text{Nogot}}$ is almost half way between the two unrefuted answers, so it is plausible for the new belief state to be neutral between the two unrefuted answers.

If it is further stipulated that Havit actually owns a Ford, then we obtain Lehrer’s (1965) no-false-lemma variant of Gettier’s case (1963). At credal state p, we have justified, true, disjunctive belief that someone owns a Ford, which falls short of knowledge because the disjunctive belief’s reason relies so essentially on a false disjunct that, if the false disjunct were become doubtful, the disjunctive belief would be retracted. Any adequate theory of rational belief should be able to model this paradigmatic situation and, thus, should violate the Preservation axiom.

The preceding intuitions are vindicated by the following no-go theorem. First, we define some properties that a sensible acceptance rule should have. To begin with, we exclude skeptical acceptance rules that accept complete answers to $\mathcal{Q}$ at almost no credal state. This is less an axiom of rationality than a delineation of the topic under discussion, which is uncertain acceptance. Say that acceptance rule B is *non-skeptical* if and only if each answer to $\mathcal{Q}$ is accepted over some open neighborhood of credal states in $\mathcal{P}$. The idea is that the open neighborhood over which H_i is accepted should surround the credal state h_i that assigns probability 1 to H_i, but it is not necessary to assume that much. In a similar spirit, we exclude the extremely gullible or opinionated rules that accept complete answers to $\mathcal{Q}$ at almost every credal state. Say that B is *non-opinionated* if and only if there is some open subset of $\mathcal{P}$ over which some incomplete, disjunctive answer is accepted. Say that B is *consistent* if and only if the inconsistent proposition $\perp$ is accepted at no credal state. Say that B is *corner-monotone* if and only if acceptance of complete answer H_i at p implies acceptance of H_i at each point on the straight line segment from p to the corner h_i of the simplex at which H_i has probability one.[8] Aside from the intuitive merits of these properties, all proposed acceptance rules we are aware of satisfy them. When all four properties are satisfied by an acceptance rule, the rule is said to be *sensible.* Then we have:

THEOREM 1 (**no-go theorem for accretive acceptance**). *Let question $\mathcal{Q}$ have at least three answers. Then no sensible acceptance rule that tracks conditioning is accretive.*

Since AGM belief revision is accretive by definition, we also have:

COROLLARY 1 (**no-go theorem for AGM acceptance**). *Let question $\mathcal{Q}$ have at least three answers. Then no sensible acceptance rule that tracks conditioning is AGM.*

The theorem extends the preceding, informal misgivings about the Preservation property discussed above. One might attempt to force accretive belief revision to track Bayesian conditioning by never accepting what one would fail to accept after conditioning on compatible evidence, assuming that acceptance is sensible. But that comes with a high price: the acceptance rule will then fail to track conditioning,[9] and the resulting theory will be unable to handle the no-false-lemma Gettier cases, which should be handled by any adequate theory of rational beliefs.

[8] Analytically, the straight line segment between two probability measures p, q in $\mathcal{P}$ is the set of all probability measures of form $ap + (1 - a)q$, where a is in the unit interval $[0, 1]$.

[9] An implementation of this idea has been presented by Leitgeb (2010). Leitgeb is careful to point out that only one side of the the tracking property is satisfied by his rule.

9. The Importance of Odds

From the no-go theorems, it is clear that any sensible rule that tracks conditioning must violate either Inclusion or Preservation. Another good bet, in light of the preceding discussion, is that any sensible rule that tracks Bayesian conditioning must pay attention to the odds between competing answers. Recall how Preservation fails at credal state p in figure 4, which we reproduce in figure 5. If, instead, one is

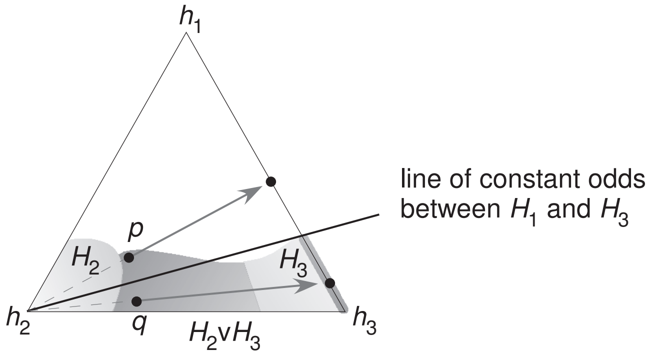

Figure 5. Line of constant odds

in credal state q, then one has a stable or robust reason for accepting $H_2 \vee H_3$ in the sense that each of the disjuncts has significantly high odds to the rejected alternative H_1, so Preservation holds. That intuition agrees with Bayesian conditioning. Since Bayesian conditioning preserves odds, H_3 continues to have significantly high odds to H_1 at the posterior credal state, where H_3 is indeed accepted. In general, the constant odds line depicted in figure 5 represents the odds threshold between H_1 and H_3 that determines whether Preservation holds or fails under new information $\neg H_2$.

We recommend, therefore, that the proper way to relax Preservation is to base acceptance on odds thresholds.

10. An Odds-Based Acceptance Rule

We now present an acceptance rule based on odds thresholds that illustrates how to sensibly track Bayesian conditioning (and to solve the two new riddles) by violating the counter-intuitive Preservation property. The particular rule discussed in this section motivates our general proposal.

Recall that an acceptance rule assigns a qualitative belief revision rule B_p to each Bayesian credal state p. Our proposed acceptance rule assigns belief revision rules of a particular form, proposed by Yoav Shoham (1987). On Shoham's approach, one begins with a well-founded strict partial order $\prec$ over some (not necessarily all) answers to $\mathcal{Q}$ that is interpreted as a *plausibility ordering*, where $H_i \prec H_j$ means that H_i is strictly more plausible than H_j with respect to order $\prec$.[10] Each plausibility order $\prec$ induces a belief revision method $\mathsf{B}_{\prec}$ as follows: given information E, let $\mathsf{B}_{\prec}(E)$ be the disjunction of all most plausible answers to $\mathcal{Q}$ with respect to $\prec$ that are logically compatible with E. Namely, we first restrict $\prec$ to the answers that are compatible with new information E to obtain new plausibility order $\prec|_E$, and then disjoin all most plausible answers therein to obtain new belief state (see figure 7.b for an example). Shoham revision always satisfies axiom Hypothetico-deductive Monotonicity, Case Reasoning, and Inclusion (Kraus, Lehmann, and Magidor 1990). But Shoham revision may violate the Preservation axiom, as shown in figure 7.b. To obtain an acceptance rule B, it suffices to assign to each credal state p a plausibility order $\prec_p$, which determines belief revision method B_p by:

$$\mathsf{B}_p \;=\; \mathsf{B}_{\prec_p}. \tag{10}$$

In light of the earlier discussion, it should come as little surprise that we define $\prec_p$ in terms of odds.[11] In particular, let t be a constant greater than 0 and define:

$$H_i \prec_p H_j \iff \frac{p(H_i)}{p(H_j)} > t, \tag{11}$$

for all i, j such that $p(H_i), p(H_j) > 0$. For $t = 3$, the proposed acceptance rule can be visualized geometrically as follows. The locus of credal states at which $p(H_1)/p(H_2) = 3$ is a line segment that originates at h_3 and intersects the line segment from h_1 to h_3, as depicted in figure 6.a. To determine whether $H_1 \prec_p H_2$, simply check whether p is above or below that line segment. Follow the same construction for each pair of answers. Figure 6.a depicts some of the plausibility rankings assigned to various regions of the simplex of Bayesian credal states.

To see that the proposed rule is sensible, recall that the initial belief state $\mathsf{B}_p(\top)$ at p is the disjunction of the most plausible answers in $\prec_p$. So the zone for accepting a belief state is bounded by the constant odds lines, as depicted in figure 6.b.[12] From the figure, it is evident that the rule is sensible.

---

[10] A strict partial order $\prec$ is said to be *well-founded* if and only if it has no infinite descending chain, or equivalently, every subset of the order has a least element.

[11] In comparison, Shoham (1987) does not explicate relative plausibility in terms of any probabilistic notions.

[12] The rule so defined was originally proposed by Isaac Levi (1996: 286), who mentions and rejects it for want of a decision-theoretic justification.

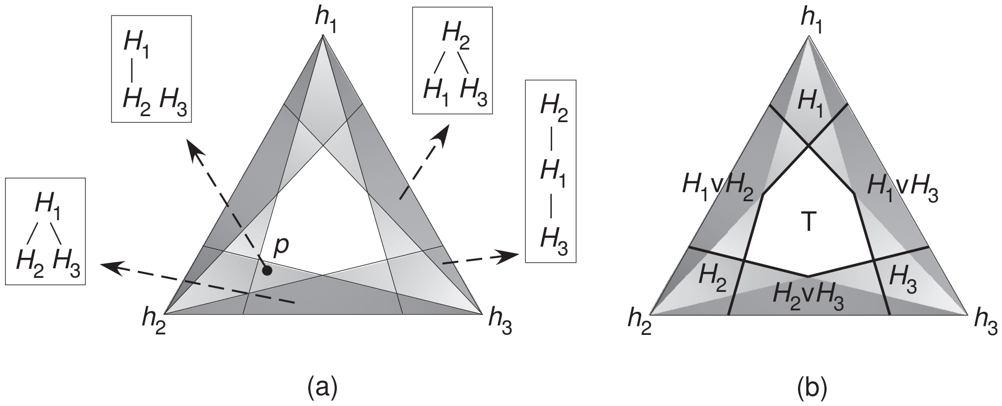

Figure 6. A rule based on odds thresholds

To see that the proposed rule tracks conditioning, take the credal state p depicted in figure 7 for example, with new information $E = H_1 \vee H_3$. To show that $\mathsf{B}_p(E) = \mathsf{B}_{p|_E}(\top)$, by the definition of Shoham

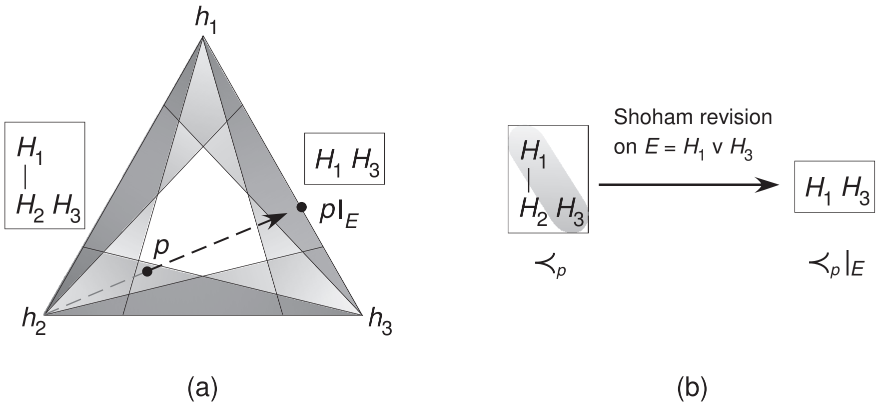

Figure 7. How the rule tracks conditioning

revision it suffices to take the plausibility order at p, restrict to information $H_1 \vee H_3$, and check that the resulting order (figure 7.b) equals the plausibility order at the posterior credal state $p|_{(H_1 \vee H_3)}$ (figure 7.a). Such equality is no accident: the relative plausibility between H_1 and H_3 at both credal states—prior and posterior—is defined by the same odds threshold, and conditioning on $H_1 \vee H_3$ always preserves the odds between H_1 and H_3. So the proposed rule tracks conditioning due to a simple principle of design: define relative plausibility by quantities preserved under conditioning. That principle falls far short of "peeking" at the underlying probabilities at each qualitative revision. After the partial order $\prec_p$ is constructed, p can be ignored through any number of qualitative revisions.

Furthermore, the proposed rule avoids the two new riddles (i.e., it satisfies Hypothetico-deductive Monotonicity (3) and Case Reasoning (4)). That can be verified by drawing lines of conditioning on figure 6.b, as we did on figures 1 and 2. Although this rule avoids the riddles because it exemplifies the general proposition 2 stated below, a simple line-drawing verification illustrates the geometric reason why no riddle occurs: the boundaries of the acceptance zones follow the lines of conditioning.

The Preservation axiom (9) is violated (figure 8), for reasons similar to those discussed in the preceding section (figure 5). Preservation is

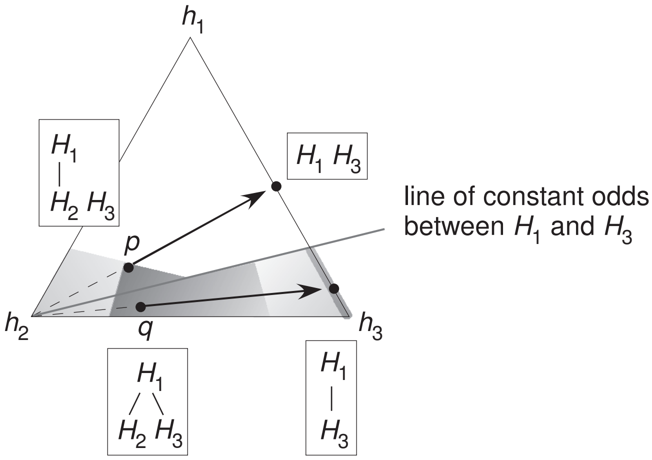

Figure 8. Preservation and odds

violated at p when $\neg H_2$ is learned, because acceptance of $H_2 \vee H_3$ depends mainly on H_2, as described above. In contrast, the acceptance of $H_2 \vee H_3$ at q is robust in the sense that each of the disjuncts is significantly more plausible than the rejected alternative H_1, so Preservation does hold at q. Indeed, the distinction between the two cases, p and q, is an intuitive one that any theory of propositional belief should be capable of drawing. For q cannot play a role that p can: modeling Lehrer's Gettier case without false lemmas (compare figure 5 with figure 4).

11. Shoham-driven Acceptance Rules

The ideas and examples in the preceding section anticipate the following theory.

An *assignment of plausibility orders* is a mapping that assigns to each credal state p a plausibility order $\prec_p$ defined on the set $\{H_i \in \mathcal{Q} : p(H_i) > 0\}$ of nonzero-probability answers, which we abbreviate as $\prec$ by abusing notation in the way we abuse B. (Think of $\prec$ as a mapping $\prec_{(\,\cdot\,)}$ that sends p to $\prec_p$). An acceptance rule B is called *Shoham-driven* if and only if it is generated by some assignment $\prec$ of plausibility orders in the sense of equation (10). Note that in the case of Shoham-driven rules, propositional belief revision is defined in terms of qualitative, plausibility orders and logical compatibility. So belief revision based on Shoham revision does seem to define an independent, propositional "space of reasons" that is not directly parasitic on probabilistic reasoning.

The example developed in the preceding section can be expressed algebraically as follows, where the question has countably many answers. Let the plausibility order $\prec_p$ assigned to p be defined by odds threshold 3:

$$H_i \prec_p H_j \iff p(H_i)/p(H_j) > 3. \tag{12}$$

Let assignment $\prec$ of plausibility orders drive acceptance rule B. Then B is sensible and tracks conditioning, due to proposition 4 below. The initial belief state $\mathsf{B}_p(\top)$ at p can be expressed by:

$$\mathsf{B}_p(\top) \;=\; \bigwedge \left\{ \neg H_i : \frac{p(H_i)}{\max_k p(H_k)} < \frac{1}{3} \right\}, \tag{13}$$

which is a special case of proposition 4 below. Equation (13) says that answer H_i is to be rejected if and only if the odds of it to the most probable alternative is "too low".

Shoham-driven rules suffice to guard against the old riddle of acceptance:

PROPOSITION 1 (**no Lottery paradox**). *Each Shoham-driven acceptance rule is consistent.*

To guard against all riddles—old and new—it suffices to require, further, that the rules track conditioning:

PROPOSITION 2 (**riddle-free acceptance**). *Each Shoham-driven acceptance rule that tracks conditioning is consistent and satisfies Hypothetico-deductive Monotonicity (3) and Case Reasoning (4).*

Furthermore:

THEOREM 2. *Suppose that acceptance rule* B *tracks conditioning and is Shoham-driven—say, by assignment* $\prec$ *of plausibility orders. Then for each credal state* p *and each proposition* E *such that* $p(E) > 0$*, it is the case that:*

$$\prec_p |_E \;=\; \prec_{p|_E}, \tag{14}$$
$$\mathsf{B}_{\prec_p|_E} \;=\; \mathsf{B}_{\prec_{p|_E}}. \tag{15}$$

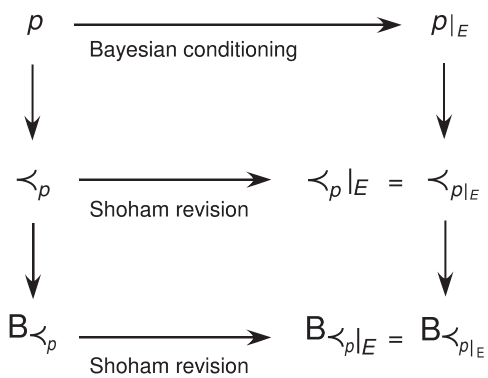

Figure 9. Shoham revision commutes with Bayesian conditioning

That is, Bayesian conditioning on E followed by assignment of a plausibility order to $p|_E$ (the upper-right path in figure 9) leads to *exactly the same result* as assigning a plausibility order to p and Shoham revising that order on E (the left-lower path in figure 9).

12. Shoham-Driven Acceptance Based on Odds

It is no accident that every Shoham-driven rule we have examined so far is somehow based on odds, as established by the main theorem of this section.

The assignment (12) of plausibility orders and the associated assignment (13) of belief states have a single, uniform threshold. The rules can be generalized by allowing each answer to have its own threshold. Let $(t_i : i \in I)$ be an assignment of odds thresholds t_i to answers H_i. Say that assignment $\prec$ of plausibility orders is *based on* assignment $(t_i : i \in I)$ of odds thresholds if and only if:

$$H_i \prec_p H_j \iff p(H_i)/p(H_j) > t_j. \tag{16}$$

Say that acceptance rule B is an *odds threshold* rule based on $(t_i : i \in I)$ if and only if the initial belief state $\mathsf{B}_p(\top)$ at p is given by:

$$\mathsf{B}_p(\top) \;=\; \bigwedge \left\{ \neg H_i : \frac{p(H_i)}{\max_k p(H_k)} < \frac{1}{t_i} \right\}, \tag{17}$$

for all p in $\mathcal{P}$. Still more general rules can be obtained by associating weights to answers that correspond to their relative *content* (Levi 1967)—e.g., quantum mechanics has more content than the catch-call hypothesis "anything else". Let $(w_i : i \in I)$ be an assignment of *weights* w_i to answers H_i. Say that assignment $\prec$ of plausibility orders is *based on* assignment $(t_i : i \in I)$ of odds thresholds and assignment $(w_i : i \in I)$ of weights if and only if:

$$H_i \prec_p H_j \iff w_i\, p(H_i)/w_j\, p(H_j) > t_j. \tag{18}$$

The range of t_i and w_i should be restricted appropriately:

PROPOSITION 3. *Suppose that that $1 < t_i < \infty$ for all i in I, and that $0 < w_i \leq 1$ for all i in I. Then for each p in $\mathcal{P}$, the relation $\prec_p$ defined by formula (18) is a plausibility order.*

Say that B is a *weighted odds threshold* rule based on $(t_i : i \in I)$ and $(w_i : i \in I)$ if and only if the unrevised belief state $\mathsf{B}_p(\top)$ is given by:

$$\mathsf{B}_p(\top) \;=\; \bigwedge \left\{ \neg H_i : \frac{w_i\, p(H_i)}{\max_k w_k\, p(H_k)} < \frac{1}{t_i} \right\}, \tag{19}$$

for all p in $\mathcal{P}$. When all weights w_i are equal, order (18) and belief state (19) are reduced to order (16) and belief state (17). Then we have:

PROPOSITION 4 (**sufficient condition for being sensible and tracking conditioning**). *Continuing proposition 3, suppose that acceptance rule B is driven by the assignment of plausibility orders based on $(t_i : i \in I)$ and $(w_i : i \in I)$. Then:*

1. *B is a weighted odds threshold rule based on $(t_i : i \in I)$ and $(w_i : i \in I)$.*

2. *B is sensible.*

3. *B tracks conditioning.*

So a Shoham-driven rule can easily be sensible and conditioning-tracking (and thus riddle-free, by proposition 2): it suffices that the plausibility orders encode information about odds and weights in the sense defined above.

Here is the next and final level of generality. The weights in formula (18) can be absorbed into odds without loss of generality:

$$H_i \prec_p H_j \iff w_i p(H_i)/w_j p(H_j) > t_j, \tag{20}$$
$$\iff p(H_i)/p(H_j) > t_j(w_j/w_i), \tag{21}$$

So we can equivalently work with double-indexed odds thresholds t_{ij} defined by:

$$t_{ij} = t_j(w_j/w_i), \tag{22}$$

where $i \neq j$. Now, allow double-indexed odds thresholds t_{ij} that are *not* factorizable into single-indexed thresholds and weights by equation (22); also allow double-indexed inequalities, which can be strict or weak. This generalization enables us to express every Shoham-driven, corner-monotone rules that tracks conditioning.

Specifically, an assignment t of *double-indexed odds thresholds* is of the form:

$$t = (t_{ij} : i, j \in I \text{ and } i \neq j), \tag{23}$$

where each threshold t_{ij} is in closed interval $[0, \infty]$. An assignment $\rhd$ of *double-indexed inequalities* is of the form:

$$\rhd = (\rhd_{ij} : i, j \in I \text{ and } i \neq j), \tag{24}$$

where each inequality $\rhd_{ij}$ is either strict $>$ or weak $\geq$. Say that assignment $\prec$ of plausibility orders is *based* on t and $\rhd$ if and only if each plausibility order $\prec_p$ is expressed by:

$$H_i \prec_p H_j \iff p(H_i)/p(H_j) \rhd_{ij} t_{ij} . \tag{25}$$

When an assignment $\prec$ of plausibility orders can be expressed in that way, say that it is *odds-based*; when an acceptance rule is driven by such assignment of plausibility orders, say again that it is *odds-based.*

THEOREM 3 (**representation of Shoham-driven rules**). *A Shoham-driven acceptance rule is corner-monotone and tracks conditioning if and only if it is odds-based.*

13. Conclusion

It is impossible for accretive (and thus AGM) belief revision to track Bayesian conditioning perfectly, on pain of failing to be sensible (theorem 1). But dynamic consonance is feasible: just adopt Shoham revision and an acceptance rule with the right geometry. When Shoham revision tracks Bayesian conditioning, acceptance of uncertain propositions must be based on the odds between competing alternatives (theorem 3). The resulting rules for uncertain acceptance solve the riddles, old and new (propositions 1 and 2). In particular, that approach solves the Lottery paradox.

14. Acknowledgements

The authors are indebted to David Makinson and David Etlin for detailed comments. We are also indebted to Hannes Leitgeb for friendly discussions about his alternative approach and for his elegant concept of acceptance rules, which we adopt in this paper. We are indebted to Teddy Seidenfeld for pointing out that Levi already proposed a version of weighted odds threshold rules. Finally, we are indebted for discussions with Horacio Arl-Costa and Arthur Paul Pedersen.

Bibliography

Alchourrn, C.E., P. Grdenfors, and D. Makinson (1985), "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions", *The Journal of Symbolic Logic*, 50: 510-530.

Gettier, E. (1963) "Is Justified True Belief Knowledge?", *Analysis* 23: 121-123.

Grove, A. (1988), "Two Modellings for Theory Change", *Journal of Philosophical Logic*, 17(2): 157-170.

Harper, W. (1975) "Rational Belief Change, Popper Functions and Counterfactuals", *Synthese*, 30(1-2): 221-262.

Kraus, S., Lehmann, D. and Magidor, M. (1990) "Nonmonotonic Reasoning, Preferential Models and Cumulative Logics", *Artificial Intelligence* 44: 167-207.

Kyburg, H. (1961) *Probability and the Logic of Rational Belief*, Middletown: Wesleyan University Press.

Lehrer, K. (1965) "Knowledge, Truth, and Evidence", *Analysis* 25: 168-175.

Leitgeb, H. (2010) “Reducing Belief Simpliciter to Degrees of Belief”, presentation at the Opening Celebration of the Center for Formal Epistemology at Carnegie Mellon University in June 2010.

Levi, I. (1967) *Gambling With Truth: An Essay on Induction and the Aims of Science*, New York: Harper & Row. 2nd ed. Cambridge, Mass.: The MIT Press, 1973.

Levi, I. (1996) *For the Sake of the Argument: Ramsey Test Conditionals, Inductive Inference and Non-monotonic Reasoning*, Cambridge: Cambridge University Press.

Lin, H. and Kelly, T.K. (forthcoming), “A Geo-logical Solution to the Lottery Paradox”, *Synthese*.

Makinson, D. and P. Grdenfors (1991) “Relations Between the Logic of Theory Change and Nonmonotonic Logic”, in A. Fuhrmann and M. Morreau (eds.), *The Logic of Theory Change*, Springer-Verlag Lecture Notes in Computer Science 465, 183-205.

Schervish, M.J., Seidenfeld, T., and Kadane, J.B. (1984), “The Extent of Non-Conglomerability of Finitely Additive Probabilities,” Zeitschrift fur Wahrscheinlictkeitstheorie und verwandte Gebiete, 66: 205-226.

Sellars, W. (1956) “Empiricism and Philosophy of Mind”, in *The Foundations of Science and the Concepts of Psychoanalysis*, Minnesota Studies in the Philosophy of Science, Vol. I, ed. by H. Feigl and M. Scriven, University of Minnesota Press; Minneapolis, MN.

Shoham, Y. (1987) “A Semantical Approach to Nonmonotonic Logics”, in M. Ginsberg (ed.) *Readings in Nonmonotonic Reasoning*, Los Altos, CA: Morgan Kauffman.

van Fraassen, B.C. (1984) “Belief and the Will”, *The Journal of Philosophy*, 81(5): 235-256.

Appendix

A. Proof of Theorem 1

To prove theorem theorem 1, suppose that rule B is consistent, corner-monotone, accretive (i.e. satisfies axioms Inclusion and Preservation), and tracks conditioning. Suppose further that B is not skeptical. It suffices to show that B is opinionated, which we prove by the following series of lemmas. Let H_i, H_j be distinct answers to $\mathcal{Q}$. Choose an arbitrary, third answer H_m to $\mathcal{Q}$ (since $\mathcal{Q}$ is assumed to have at least three answers). Let h_i be the credal state in which H_i has probability 1, and similarly for h_j and h_m. Let $\triangle\, h_i h_j h_m$ denote the two dimensional space $\{p \in \mathcal{P} : p(H_i) + p(H_j) + p(H_m) = 1\}$ (figure 10.a). Let $\overline{h_i\, h_m}$

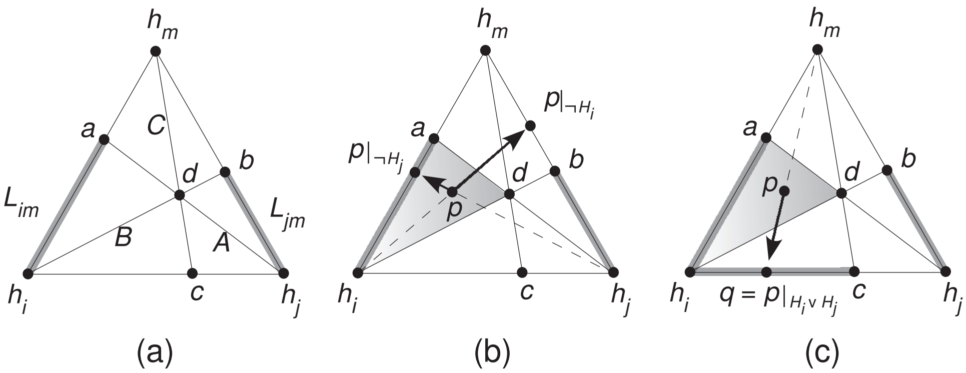

Figure 10. Why every accretive rule that tracks conditioning fails to be sensible

denote the one-dimensional subspace $\{p \in \mathcal{P} : p(H_i) + p(H_j) = 1\}$, and similarly for $\overline{h_i\, h_j}$ and $\overline{h_j\, h_m}$. Let L_{im} be the set of the credal states on line segment $\overline{h_i\, h_m}$ at which H_i is accepted by B as strongest; namely:

$$L_{im} \;=\; \{p \in \overline{h_i\, h_m} : \mathsf{B}_p = H_i\}.$$

LEMMA 1. *L_{im} is a connected line segment of nonzero length that contains h_i but does not contain h_m.*

Proof. By non-skepticism, there exists open subset O of $\mathcal{P}$ over which B accepts H_i as strongest. Let O' be the image $\{o|_{H_i \vee H_m} : o \in O\}$ of O under conditioning on $H_i \vee H_m$. Since O is open, O' is an open subset of $\overline{h_i\, h_m}$. Note that the conditioning proposition $H_i \vee H_m$ is consistent with the prior belief state H_i, so Preservation applies. Since B satisfies Preservation and tracks conditioning, B accepts old belief H_i over O'. It follows that B accepts H_i as strongest over O', because B is consistent and the only proposition strictly strongest than H_i in the algebra is the inconsistent proposition $\perp$. So L_{im} is nonempty. Then, since B is corner-monotone, L_{im} is a nonempty, connected line segment that contains h_i. It remains to show that L_{im} does not contain h_m. Suppose for reductio that L_{im} contains h_m, then L_{im} must be so large that it is identical to $\overline{h_i\, h_m}$, by corner-monotonicity. By the same argument for showing that there is an open subset O' of $\overline{h_i\, h_m}$ over which B accepts H_i, we have that there is an open subset O'' of $\overline{h_i\, h_m}$ over which B accepts H_m. So B accepts both H_m and H_i over O'', and hence by closure under conjunction, B accepts their conjunction, which is an inconsistent proposition. So B is not consistent—contradiction.

Let a be the endpoint of L_{im} that is closest to h_m; namely, probability measure a is such that:

$$a \in \overline{h_i\, h_m},$$
$$a(H_m) = \sup\{p(H_m) : p \in L_{im}\}.$$

By the lemma we just proved, point a lies in the interior of side $\overline{h_i\, h_m}$. Applying the above argument for pair (i, m) to pair (j, m), we have that the set L_{jm}, defined by

$$L_{jm} = \{p \in \overline{h_j\, h_m} : \mathsf{B}_p = H_j\},$$

is a connected line segment of nonzero length that contains h_j but does not contain h_m, with endpoint b that lies in the interior of side $\overline{h_j\, h_m}$. Since both points a, b lie in the interiors of their respective sides, we have the following constructions. Let A be the line that connects a to h_j, B be the line that connects b to h_i, and C be the line that connects h_m through the intersection d of A and B, to point c on side $\overline{h_i\, h_j}$.

LEMMA 2. B *accepts* H_i *as strongest over the interior of* $\triangle adh_i$.

Proof. Consider an arbitrary point p in the interior of $\triangle\, adh_i$ (figure 10.b). Argue as follows that B accepts $H_i \vee H_j$ at p. Take p as a prior state and consider $\neg H_j$ as the conditioning information. Note that credal state $p|_{\neg H_j}$ falls inside L_{im}, so B accepts H_i as strongest at the posterior credal state $p|_{\neg H_j}$. Then, since B tracks conditioning and satisfies Inclusion, we have that:

$$\mathsf{B}_p \wedge \neg H_j \models H_i$$

(namely the posterior belief state H_i is entailed by the conjunction of the the prior belief state and the conditioning information). Then, by the consistency of B and the mutual exclusion among the answers, we have only three possibilities for B_p:

$$\mathsf{B}_p \text{ is either } H_i, \text{ or } H_j, \text{ or } H_i \vee H_j.$$

Rule out the last two alternatives as follows. Suppose for reductio that the prior belief state B_p is H_j or $H_i \vee H_j$. Consider $\neg H_i$ as the conditioning information, which is consistent with the prior belief state and thus makes Preservation applicable. Then, since B tracks conditioning and satisfies Preservation, the posterior belief state $\mathsf{B}_{p|_{\neg H_i}}$ must entail $\mathsf{B}_p \wedge \neg H_i$ (i.e. the conjunction of the prior belief state and the information). But the latter proposition $\mathsf{B}_p \wedge \neg H_i$ equals H_j, by the reductio hypothesis. So $\mathsf{B}_{p|_{\neg H_i}} = H_j$, by the consistency of B. Hence $p|_{\neg H_i}$ lies on line segment L_{jm} by the construction of L_{jm}—but that is impossible (figure 10.b). Ruling out the last two alternatives for B_p, we conclude that $\mathsf{B}_p = H_i$.

LEMMA 3. B *accepts* H_i *as strongest over the interior of* $\overline{h_i\, c}$.

Proof. Let p be an arbitrary interior point of $\triangle\, adh_i$. So $\mathsf{B}_p = H_i$. Consider proposition $H_i \vee H_j$ as the conditioning information. Then, since B tracks conditioning and satisfies Preservation, the posterior belief state $\mathsf{B}_{p|_{H_i \vee H_j}}$ entails $\mathsf{B}_p \wedge (H_i \vee H_j)$ (i.e. the conjunction of the prior belief state and the information), which equals H_i. Then, by consistency, the posterior belief state is determined:

$$\mathsf{B}_{p|_{H_i \vee H_j}} \;=\; H_i.$$

Let q be an arbitrary point in the interior of $\overline{h_i\, c}$. Then q can be expressed as $q = p|_{H_i \vee H_j}$ for some point p in the interior of $\triangle\, adh_i$ (figure 10.c). So, by the formula we just proved, $\mathsf{B}_q = \mathsf{B}_{p|_{H_i \vee H_j}} = H_i$, as required.

LEMMA 4. *There is no open subset of* $\overline{h_i\, h_j}$ *over which* B *accepts* $H_i \vee H_j$ *as strongest.*

Proof. We have established in the last lemma that B accepts H_i as strongest over the interior of $\overline{h_i\, c}$. By the same argument, B accepts H_j as strongest over the interior of $\overline{h_j\, c}$ (figure 10.c). So if B accepts disjunction $H_i \vee H_j$ as strongest somewhere on $\overline{h_i\, h_j}$, B does so at some of the three points: h_i, h_j, and c. (We can rule out the first two alternatives; but the for the sake of the lemma, this result suffices.)

Since the choice of H_i and H_j is arbitrary, the last lemma generalizes to the following:

LEMMA 5. *For each pair of distinct answers* H_i, H_j *to* $\mathcal{Q}$*, there is no open subset of* $\overline{h_i\, h_j}$ *over which* B *accepts* $H_i \vee H_j$ *as strongest.*

The last lemma establishes opinionation for all edges of the simplex. The next step is to extend opinionation to the whole simplex.

LEMMA 6. *For each disjunction* D *of at least two distinct answers to the question, there is no open subset of* $\mathcal{P}$ *over which* B *accepts* D *as strongest.*

Proof. Suppose for reductio that some disjunction $H_i \vee H_j \vee X$ of at least two distinct answers is accepted by B as strongest over some open subset O of $\mathcal{P}$. Take $H_i \vee H_j \vee X$ as the prior belief state at each point in O and consider $H_i \vee H_j$ as the conditioning information. So the image O' $(= \{p|_{H_i \vee H_j} : p \in O\})$ of O under conditioning on $H_i \vee H_j$ is an open subset of 1-dimensional space $\overline{h_i\, h_j}$. Let p' be an arbitrary point in O'. Since B tracks conditioning and satisfies Inclusion, posterior belief state $\mathsf{B}_{p'}$ is entailed by $(H_i \vee H_j \vee X) \wedge (H_i \vee H_j)$ (i.e. the conjunction of the prior state and the new information), which equals $H_i \vee H_j$. But $\mathsf{B}_{p'}$ also entails $H_i \vee H_j$, for otherwise the process of conditioning p' on $\neg H_j$ to obtain h_i would violate the fact that B tracks conditioning, satisfies Inclusion, and accepts H_i at h_i. So $\mathsf{B}_{p'} = H_i \vee H_j$. Hence B accepts $H_i \vee H_j$ as strongest over open subset O' of $\overline{h_i\, h_j}$, which contradicts the last lemma.

Proof. [Proof of Theorem 1] Since the last lemma states that B is opinionated, we are done.

B. Proof of Theorem 2

Proof. [Proof of Theorem 2] The domains of $\prec_{p|_E}$ and $\prec_p|_E$ coincide, because each plausibility order $\prec_q$ is defined on the set of the answers to $\mathcal{E}$ that have nonzero probability with respect to q. Let H_i and H_j be arbitrary distinct answers in the (common) domain. Since both answers are in the domain of $\prec_{p|_E}$, we have that $p(H_i|E) > 0$, $p(H_j|E) > 0$. Since both answers are in $\prec_{p,E}$, we have that $H_i \vee H_j$ entails E. It follows that $p|_{(H_i \vee H_j)} = p|_{E \wedge (H_i \vee H_j)}$, where both terms are defined. Then it suffices to show that $H_i \prec_{p|_E} H_j$ if and only if $H_i {\prec_p}|_E H_j$, as follows:

$$
\begin{array}{lll}
 & H_i \prec_{p|_E} H_j & \\
\iff & \mathsf{B}_{p|_E}(H_i \vee H_j) = H_i & \text{by being Shoham-driven;} \\
\iff & \mathsf{B}_{p|_{(E \wedge (H_i \vee H_j))}}(\top) = H_i & \text{by tracking conditioning;} \\
\iff & \mathsf{B}_{p|_{(H_i \vee H_j)}}(\top) = H_i & \text{since } p|_{(H_i \vee H_j)} = p|_{E \wedge (H_i \vee H_j)}; \\
\iff & \mathsf{B}_p(H_i \vee H_j) = H_i & \text{by tracking conditioning;} \\
\iff & H_i {\prec_p}|_{(H_i \vee H_j)} H_j & \text{by being Shoham-driven;} \\
\iff & H_i {\prec_p}|_E H_j & \text{since } H_i \vee H_j \text{ entails } E.
\end{array}
$$

So $\prec_p |_E = \prec_{p|_E}$. Hence, $\mathsf{B}_{\prec_p|_E} = \mathsf{B}_{\prec_{p|_E}}$.

C. Proof of Theorem 3

Proof. [Proof of the Right-to-Left Side of Theorem 3] Let B be driven by an odds-based assignment $(\prec_p : p \in \mathcal{P})$ of plausibility orders. Then, *a fortiori*, B is Shoham-driven. That B is corner-monotone follows from routine, algebraic verification. To see that B tracks conditioning (i.e. that $\mathsf{B}_p(E) = \mathsf{B}_{p|_E}(\top)$), since B is Shoham-driven, it suffices to show that an answer is most plausible in $\prec_p|_E$ if and only if it is most plausible in $\prec_{p|_E}$, which follows from the odds-based definition of $\prec_p$ and preservation of odds by Bayesian conditioning.

Proof. [Proof of the Left-to-Right Side of Theorem 3] Suppose that B is corner-monotone, tracks conditioning, and is Shoham-driven, namely driven by assignment $(\prec_p : p \in \mathcal{P})$ of plausibility orders. It suffices to show that $(\prec_p : p \in \mathcal{P})$ is odds-based. For each pair of distinct indexes i, j in I, define odds threshold $t_{ij} \in [0, \infty]$ and inequality $\rhd_{ij} \in \{>, \geq\}$ by:

$$\mathsf{Odds}_{ij} = \{q(H_i)/q(H_j) : q \in \mathcal{P}, q(H_i) + q(H_j) = 1, \text{ and } H_i \prec_q H_j\} \tag{26}$$

$$t_{ij} = \inf \mathsf{Odds}_{ij}\,; \tag{27}$$

$$\rhd_{ij} = \begin{cases} \geq & \text{if } t_{ij} \in \mathsf{Odds}_{ij}\,, \\ > & \text{otherwise.} \end{cases} \tag{28}$$

By corner-monotonicity, Odds_{ij} is closed upward, namely that $s \in \mathsf{Odds}_{ij}$ and $s < s'$ implies that $s' \in \mathsf{Odds}_{ij}$. So for each q in $\mathcal{P}$ such that $q(H_i) + q(H_j) = 1$,

$$H_i \prec_q H_j \iff q(H_i)/q(H_j) \ \rhd_{ij}\ t_{ij}. \tag{29}$$

It remains to check that for each credal state p and pair of distinct answers H_i and H_j in the domain of $\prec_p$ (i.e. $p(H_i) > 0$ and $p(H_j) > 0$), equation (25) holds with respect to odds thresholds (27) and inequalities (28):

$$H_i \prec_p H_j \iff p(H_i)/p(H_j) \ \rhd_{ij}\ t_{ij}\,. \tag{30}$$

Note that $p(H_i \vee H_j) = p(H_i) + p(H_j) > 0$, so $p|_{(H_i \vee H_j)}$ is defined. Then:

$$\begin{aligned}
& H_i \prec_p H_j \\
\iff\ & H_i \prec_p |_{(H_i \vee H_j)} H_j \\
\iff\ & H_i \prec_{p|_{(H_i \vee H_j)}} H_j && \text{by theorem 2;} \\
\iff\ & H_i \prec_q H_j && \text{by defining } q \text{ as } p|_{(H_i \vee H_j)}; \\
\iff\ & q(H_i)/q(H_j) \ \rhd_{ij}\ t_{ij} && \text{by (29);} \\
\iff\ & p(H_i)/p(H_j) \ \rhd_{ij}\ t_{ij} && \text{since } q = p|_{(H_i \vee H_j)}.
\end{aligned}$$

D. Proof of Propositions 1-4

Proof. [Proof of Proposition 1] Consistency follows from the well-foundedness of a plausibility order.

Proof. [Proof of Proposition 2] Consistency follows as an immediate consequence of proposition 1. So it suffices to show, for each p, that the relation $\mathsf{B}_{p|_E}(\top) \models H$ between E and H satisfies Hypothetico-deductive Monotonicity (3) and Case Reasoning (4). That relation is equivalent to relation $\mathsf{B}_p(E) \models H$ between E and H (by tracking conditioning). Since B is Shoham-driven, the latter relation is defined by the plausibility order $\prec_p$ assigned to p, which is a special case of the so-called *preferential models* that validate nonmonotonic logic system P (Kraus, Lehmann, and Magidor 1990). Then it suffices to note that system P entails Hypothetico-deductive Monotonicity (as a consequence of axiom Cautious Monotonicity) and Case Reasoning (as a consequence of axiom Or).

Proof. [Proof of Proposition 3] To show that $\prec_p$ is transitive, suppose that $H_i \prec_p H_j$ and $H_j \prec_p H_k$. So $w_i p(H_i)/w_j p(H_j) > t_j$ and $w_j p(H_j)/w_k p(H_k) > t_k$. Hence $w_i p(H_i)/w_k p(H_k) > t_j t_k$. But odds threshold t_j is assumed to be greater than 1, so $w_i p(H_i)/w_k p(H_k) > t_k$. So $H_i \prec_p H_k$, which establishes transitivity. Irreflexivity follows from the fact that $w_i p(H_i)/w_i p(H_i) = 1 \not> t_i$, by the assumption that $t_i > 1$. Asymmetry follows from the fact that if $w_i p(H_i)/w_j p(H_j) > t_j > 1$, then $w_j p(H_j)/w_i p(H_i)$ is less than 1 and thus fails to be greater than t_i. To establish well-foundedness, suppose for reductio that $\prec_p$ is not well-founded. Then $\prec_p$ has an infinite descending chain $H_i \succ_p H_j \succ_p H_k \succ_p \ldots$. Since $t_i > 1$ for all i in I, we have that $w_i p(H_i) < w_j p(H_j) < w_k p(H_k) < \ldots$. So the sum is unbounded. But each weight is assumed to be no more than 1, so the sum of (unweighted) probabilities $p(H_i) + p(H_j) + p(H_k) + \ldots$ is also unbounded—which contradicts the fact that p is a probability measure.

Proof. [Proof of Proposition 4] Part 1, that B is a weighted odds threshold rule, is established as follows:

$$\mathsf{B}_p(\top) = \bigvee \{H_j \in \mathcal{Q} : H_j \text{ is most plausible in } \prec_p\} \tag{31}$$

$$= \bigvee \{H_j \in \mathcal{Q} : \max_k w_k p(H_k)/w_j p(H_j) \not> t_j\} \tag{32}$$

$$= \bigwedge \{\neg H_i \in \mathcal{Q} : \max_k w_k p(H_k)/w_i p(H_i) > t_i\} \tag{33}$$

$$= \bigwedge \left\{\neg H_i \in \mathcal{Q} : \frac{w_i\, p(H_i)}{\max_k w_k\, p(H_k)} < \frac{1}{t_i}\right\}. \tag{34}$$

For part 2, that the rule is sensible, note that the parameters are assumed to be restricted as follows: $1 < t_i < \infty$ and $0 < w_i \leq 1$ for all i in I. Then the rule is consistent, because for each credal state p, the rule does not reject the answer H_k in $\mathcal{Q}$ that maximizes $w_k\, p(H_k)$. The rule is corner-monotone, because when the rule accepts H_k at p, moving p toward corner h_i only makes H_k have higher odds to all the other answers and hence the rule continues to accept H_k. To see that the rule is non-skeptical, here is a recipe for constructing, for each answer H_k, an open set of credal states over h_k (i.e. the credal state at which H_k has unit probability). For each answer H_i distinct from H_k, H_i has zero probability at h_k, and thus the condition for rejecting H_i in formula (34) is satisfied. Then, by the upper bound on odds thresholds t_i and the strict inequality in formula (34), each of the rejection conditions for H_i, where $i \neq k$, continues to hold over an open neighborhood O of h_k. So the rule accepts H_k over O. To establish that the rule is non-opinionated, it suffices to show that one particular disjunction, say $H_1 \vee H_2$, is accepted over an open neighborhood. Consider the credal state p such that $w_1\, p(H_1) = w_2\, p(H_2) > 0$ and $p(H_i) = 0$ for all $i \neq 1, 2$. Then, by formula (34) and the lower bound on thresholds t_i, the rule accepts $H_1 \vee H_2$ at p. By the lower bound on thresholds t_i, again, we can find an open neighborhood of p over which the rule accepts $H_1 \vee H_2$. So the rule is sensible, by having the above four properties. Part 3, that the rule tracks conditioning, is a special case of the left-to-right side of theorem 3.