

Propositional Reasoning that Tracks Probabilistic Reasoning

Hanti Lin* Kevin T. Kelly^{†‡}

Carnegie Mellon University

March 19, 2012

Abstract

This paper concerns the extent to which uncertain propositional reasoning can track probabilistic reasoning, and addresses kinematic problems that extend the familiar Lottery paradox. An acceptance rule assigns to each Bayesian credal state p a propositional belief revision method \mathbb{B}_p , which specifies an initial belief state $\mathbb{B}_p(\top)$, that is revised to the new propositional belief state $\mathbb{B}(E)$ upon receipt of information E . An acceptance rule *tracks* Bayesian conditioning when $\mathbb{B}_p(E) = \mathbb{B}_{p|E}(\top)$, for every E such that $p(E) > 0$; namely, when acceptance by propositional belief revision equals Bayesian conditioning followed by acceptance. Standard proposals for uncertain 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 Alchourrón, Gärdenfors, and Makinson 1985) cannot be realized in a sensible way by any uncertain 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).

*e-mail: hantil@andrew.cmu.edu

†e-mail: kk3n@andrew.cmu.edu

‡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.

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 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. With probability one *some* ticket wins, so that proposition is also accepted. 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 similarly for H_2, H_3 . The




Figure 1: the first riddle

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 bottom corner of the diamond, which is so acute that conditioning carries p outside of the diamond. If the bottom corner of the diamond is made more blunt, to match the slope of the conditioning ray, then the paradox does 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 $B_p \models H$. Then the principle may

be stated succinctly as follows, where $p|_E$ denotes the conditional distribution $p(\cdot|E)$:

$$\mathbb{B}_p \models H \text{ and } H \models E \implies \mathbb{B}_{p|_E} \models H. \quad (1)$$

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 acceptance 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 a different and, in a sense, complemen-




Figure 2: second riddle

tary 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. The Lockean rule accepts that ticket 2 loses no matter whether one learns that ticket 3 wins (i.e. p moves to $p|_{H_3}$) or that ticket 3 loses (i.e. p moves to $p|_{\neg H_3}$), but the lockean rule refuses to accept that ticket 2 loses until one actually learns what happens with ticket 3. That violates the following principle:¹

$$\mathbb{B}_{p|_E} \models H \text{ and } \mathbb{B}_{p|_{\neg E}} \models H \implies \mathbb{B}_p \models H, \quad (2)$$

which we call *Case Reasoning*.

¹The 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 (2011), so non-conglomerability does not necessarily arise in the given partition.

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 thresholds that generate both riddles at once. Unlike the lottery paradox, which requires more 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 perfect, 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. Perfect harmony will be shown to be impossible to achieve if one insists on employing the popular AGM approach to propositional belief revision. Then, we exhibit a 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, \dots, H_i, \dots are the (*complete*) an-

²The claim is a special case of theorem 3 in Lin and Kelly (2011).

swers. Let \mathcal{A} denote the least σ -algebra containing \mathcal{Q} (i.e., the set of all disjunctions of complete answers together with the unsatisfiable proposition \perp). 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 we identify each belief state with the conjunction of all propositions believed. A *belief revision method* is a mapping $\mathbf{B} : \mathcal{A} \rightarrow \mathcal{A}$, understood as specifying the initial belief state $\mathbf{B}(\top)$, which would evolve into new belief state $\mathbf{B}(E)$ 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

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

Case Reasoning has a similar statement:⁵

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

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 function \mathbf{B} that assigns to each credal state p a belief revision method \mathbf{B}_p . Then $\mathbf{B}_p(\top)$ is the belief state *accepted* unconditionally at credal state p , and proposition H is *accepted* (unconditionally) by rule \mathbf{B} at credal state p if and only if $\mathbf{B}_p(\top) \models H$.⁶

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

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

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

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

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

$$E \vdash_p H \iff \mathbf{B}_p(E) \models H. \quad (6)$$

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 $B_p(\top)$ and then revises it propositionally to obtain the new propositional belief state $B_p(E)$ (i.e., the left-lower path in figure 3). According to the second course

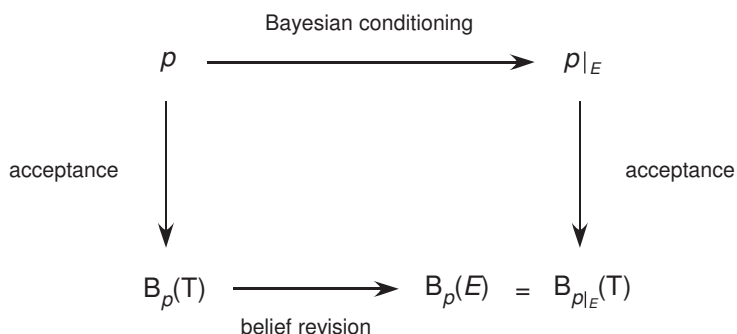


Figure 3: Belief revision tracks Bayesian conditioning

of action, she first conditions p to obtain the posterior credal state $p|_E$ and then accepts $B_{p|_E}(\top)$ (i.e., the upper-right path in figure 3). Pre-established harmony requires that the two processes should always agree (i.e., the diagram should always commute). Accordingly, say that acceptance rule B *tracks conditioning* if and only if:

$$B_p(E) = B_{p|_E}(\top), \quad (7)$$

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

7 Accretive Belief Revision

It is easy to achieve perfect tracking: just *define* $B_p(E)$ according to equation (7). To avoid triviality, one must 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

We are indebted to Hannes Leitgeb (2010) for the idea of framing our discussion in terms of conditional acceptance, which he presented at the Opening Celebration of the Center for Formal Epistemology at Carnegie Mellon University. Our own approach (Lin and Kelly 2011), prior to seeing his work, was to formulate the issues in terms of conditional logic, via a probabilistic Ramsey test, which involves more cumbersome notation and an irrelevant commitment to an epistemic interpretation of conditionals.

into two parts: belief revision method B satisfies *Inclusion* if and only if:⁷

$$B(\top) \wedge E \models B(E). \quad (8)$$

Method B satisfies *Preservation* if and only if:

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

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 Alchourrón, Gärdenfors, 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. The question is: who owns a Ford? There are three potential answers: “Nogot” vs. “Havit” vs. “nobody” (figure 4). Now, Nogot shows us car keys and his driver’s




Figure 4: How Preservation may fail plausibly

license and Havit does nothing, so we think that it is pretty probable that Nogot has a Ford (i.e., credal state p is close to the acceptance zone for “Nogot”). 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

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

accept is the disjunction of “Nogot” with “Havit”, namely “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 “somebody”, in spite of the fact that the new information is still compatible with the old belief state. So it seems plausible to drop the old belief “somebody” in the new belief state, i.e., to violate the Preservation axiom. 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 (1963) celebrated counterexample to justified true belief as an analysis of knowledge. 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 theory of rational belief that models this paradigmatic situation must 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. That is less an axiom of rationality than a delineation of the topic under discussion, which is *uncertain* acceptance. Say that acceptance rule \mathbf{B} is *non-skeptical* if and only if each complete answer to \mathcal{Q} is accepted over some non-empty, open subset of \mathcal{P} . Think of the non-empty, open subset as a ball of non-zero diameter, so acceptance of H_i over a line or a scattered set of points would not suffice. Of course, it is natural to require that the ball include h_i , itself, but that follows from further principles. Open sets are understood to be unions of balls with respect to the standard Euclidean metric, according to which the distance between p, q in \mathcal{P} is just:⁸

$$\|p - q\| = \sqrt{\sum_{H_i \in \mathcal{Q}} (p(H_i) - q(H_i))^2}.$$

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 \mathbf{B} is *non-opinionated* if and only if there is some non-empty, open subset of \mathcal{P} over which some incomplete, disjunctive answer is accepted. Say that \mathbf{B} is *consistent* if and only if the inconsistent proposition \perp is accepted at no credal state. Say that \mathbf{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

⁸The sum over \mathcal{Q} is defined over \mathcal{P} and assumes maximum value $\sqrt{2}$.

one.⁹ Aside from the intuitive merits of these properties, all proposed acceptance rules we are aware of satisfy them. Rules that satisfy all four properties are said to be *sensible*. Then we have:

Theorem 1 (no-go theorem for accretive acceptance). *Let question Q have at least three complete 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 Q have at least three complete answers. Then no sensible acceptance rule that tracks conditioning is AGM.*

In light of the theorem, 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. But that comes with a high price: no such rule is sensible.¹⁰

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 should 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 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 in the posterior credal state, at which 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.

⁹Analytically, 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 $[0, 1]$.

¹⁰Leitgeb (2010) shows that a sensible AGM rule can satisfy one side of the tracking equivalence: $\mathbb{B}_p(E)$ is entailed by $\mathbb{B}_{p|E}(\top)$.

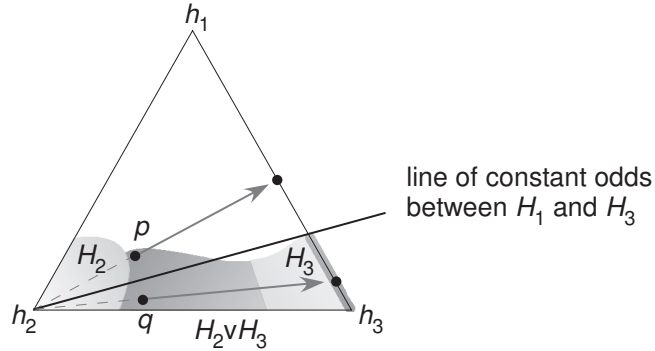


Figure 5: Line of constant odds

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) complete 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 .¹¹ Each plausibility order \prec induces a belief revision method B_{\prec} as follows: given information E in \mathcal{A} , let $B_{\prec}(E)$ be the disjunction of the most plausible answers to \mathcal{Q} with respect to \prec that are logically compatible with E . More precisely, we first restrict \prec to the answers that are compatible with new information E to obtain the new plausibility order $\prec|_E$, and then disjoin the most plausible answers compatible with E according to $\prec|_E$ to obtain our new belief state (see figure 7.b for an example). Shoham revision always satisfies axioms 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:

$$B_p = B_{\prec_p}. \quad (10)$$

¹¹A 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.

We define \prec_p in terms of odds.¹² In particular, let t be a constant greater than 1 and define:

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

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




Figure 6: A rule based on odds thresholds

p is above or below that line segment. Follow the same construction for each pair of complete answers. Figure 6.a depicts some of the plausibility orders 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 $\mathbb{B}_p(\top)$ 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.¹³ From the figure, it is evident that the rule is sensible.

To see that the proposed rule tracks conditioning, consider the credal state p depicted in figure 7, with new information $E = H_1 \vee H_3$. To show that $\mathbb{B}_p(E) = \mathbb{B}_{p|E}(\top)$, it suffices to restrict the plausibility order at p to information $H_1 \vee H_3$, and to 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

¹²Shoham (1987) does not explicate relative plausibility in terms of any probabilistic notions.

¹³The rule so defined was originally proposed by Isaac Levi (1996: 286), who mentions and rejects it for want of a decision-theoretic justification.

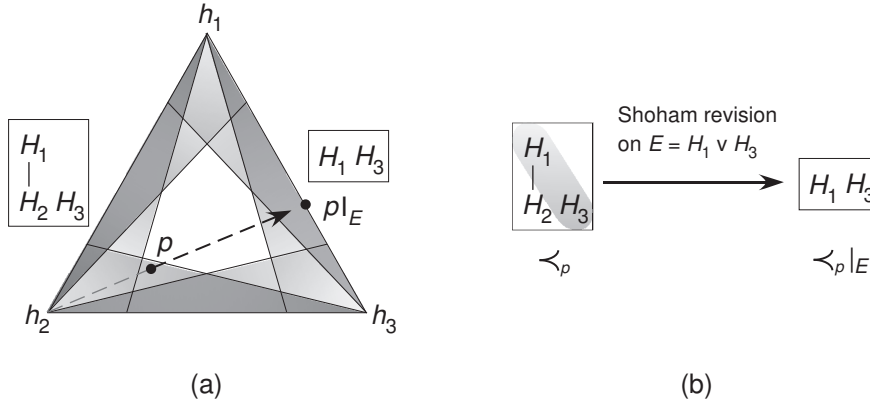


Figure 7: How the rule tracks conditioning

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 cannot be accused of “peeking” at the underlying probabilities at each qualitative revision. Whereas full specification of the position of p requires infinitely precise information, belief revision depends only on which discrete plausibility order is assigned to p , which amounts to just nineteen discrete possibilities in the case of three answers.

Furthermore, the proposed rule avoids the two new riddles (i.e., it satisfies Hypothetico-deductive Monotonicity (1) and Case Reasoning (2)). Although that claim follows in general from Proposition 2 below, it can be illustrated geometrically for the case at hand by drawing lines of conditioning on figure 6.b, as we did on figures 1 and 2.

The Preservation axiom (9) is violated (figure 8), for reasons similar to those discussed in the preceding section (figure 5). Preservation is 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 epistemically crucial. For p can model Lehrer’s Gettier case without false lemmas and q cannot (compare figure 8 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 \prec 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-




Figure 8: Preservation and odds

probability answers (i.e., \prec is a mapping $\prec(\cdot)$ that sends p to \prec_p). An acceptance rule \mathbf{B} is called *Shoham-driven* if and only if it is generated by some assignment $\prec(\cdot)$ of plausibility orders in the sense of equation (10). Recall 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 define an independent, propositional “space of reasons” that does not presuppose full probabilistic reasoning.

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

$$H_i \prec_p H_j \iff p(H_i)/p(H_j) > 3. \quad (12)$$

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

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

which is a special case of proposition 4 below. Equation (13) says that answer H_i is to be rejected if and only if its odds ratio against the 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 of the 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 (1) and Case Reasoning (2).*

Furthermore:

Theorem 2. *Suppose that acceptance rule \mathbf{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:*

$$\prec_p |E = \prec_{p|E}, \quad (14)$$

$$\mathbf{B}_{\prec_p |E} = \mathbf{B}_{\prec_{p|E}}. \quad (15)$$

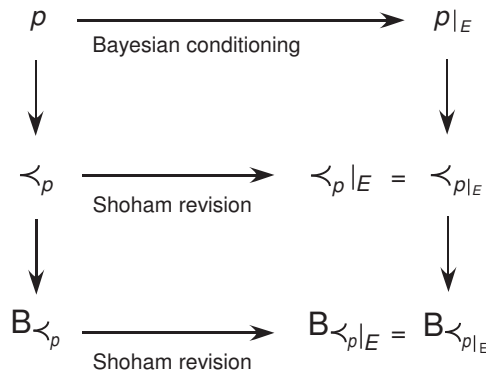


Figure 9: Shoham revision commutes with Bayesian conditioning

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 employ a single, uniform threshold. The idea can be generalized by allowing

each complete 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:

$$H_i \prec_p H_j \iff p(H_i)/p(H_j) > t_j. \quad (16)$$

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

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

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:

$$H_i \prec_p H_j \iff w_i p(H_i)/w_j p(H_j) > t_j. \quad (18)$$

The range of t_i and w_i should be restricted appropriately:

Proposition 3. *Suppose that $1 < t_i < \infty$ and $0 < w_i \leq 1$, 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 \mathbf{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 $\mathbf{B}_p(\top)$ is given by:

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

for all p in \mathcal{P} . When all weights w_i are equal, order (18) and belief state (19) reduce 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 \mathbf{B} is driven by the assignment of plausibility orders based on $(t_i : i \in I)$ and $(w_i : i \in I)$. Then:*

1. \mathbf{B} is a weighted odds threshold rule based on $(t_i : i \in I)$ and $(w_i : i \in I)$.
2. \mathbf{B} tracks conditioning.
3. \mathbf{B} is sensible if \mathcal{Q} contains at least two complete answers and there exists positive integer N such that for each i in I , $t_i \leq N$.

Rule **B** is not sensible if the antecedent of the preceding statement is false.¹⁴ 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:

$$H_i \prec_p H_j \iff w_i p(H_i)/w_j p(H_j) > t_j, \quad (20)$$

$$\iff p(H_i)/p(H_j) > t_j(w_j/w_i), \quad (21)$$

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

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

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 rule that tracks conditioning.

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

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

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

$$\triangleright = (\triangleright_{ij} : i, j \in I \text{ and } i \neq j), \quad (24)$$

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

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

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.*

¹⁴If \mathcal{Q} contains only one complete answer, then the rule is trivially opinionated. If the odds thresholds t_i are unbounded, say $t_i = i$ for each positive integer i , then every non-empty Euclidean ball at corner h_1 of \mathcal{P} fails to be contained in the acceptance zone for H_1 .

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 to Horacio Arló-Costa and Arthur Paul Pedersen for discussions.

Bibliography

- Alchourrón, C.E., P. Gärdenfors, 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.
- 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. (2011), “A Geo-logical Solution to the Lottery Paradox”, *Synthese*, DOI: 10.1007/s11229-011-9998-1.
- Schervish, M.J., Seidenfeld, T., and Kadane, J.B. (1984), “The Extent of Non-Conglomerability of Finitely Additive Probabilities,” *Zeitschrift fur Wahrscheinlichkeitstheorie 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 Kaufman.
- van Fraassen, B.C. (1984) “Belief and the Will”, *The Journal of Philosophy*, 81(5): 235-256.

A Proof of Theorem 1

To prove theorem 1, let \mathcal{Q} have at least three complete answers. Suppose that rule \mathbf{B} is consistent, corner-monotone, accretive (i.e. satisfies axioms Inclusion and Preservation), and tracks conditioning. Suppose further that \mathbf{B} is not skeptical. It suffices to show that \mathbf{B} is opinionated, which is accomplished by the following series of lemmas.

Lemma 1. *Let O be a non-empty open subset of \mathcal{P} , and H_i, H_j be distinct complete answers to \mathcal{Q} . Then O contains a credal state that assigns nonzero probabilities to both H_i and H_j .*

Proof. Since O is open (in Euclidean metric topology), p is the center of an open sphere S in Euclidean metric with some non-zero radius r that is contained in O . If p assigns non-zero probability to both H_i and H_j , we are done. If p assigns zero probability to exactly one of the two answers, say, H_i , then move probability mass $0 < q < \min(r/\sqrt{2}, p(H_j))$ from H_j to H_i to form p' . Then computing the Euclidean

distance between p, p' yields:

$$\|p - p'\| = \sqrt{\sum_{H_n \in \mathcal{Q}} (p(H_n) - p'(H_n))^2} \quad (26)$$

$$= \sqrt{(p(H_i) - p'(H_i))^2 + (p(H_j) - p'(H_j))^2} \quad (27)$$

$$< \sqrt{\left(\frac{r}{\sqrt{2}}\right)^2 + \left(\frac{r}{\sqrt{2}}\right)^2} = r. \quad (28)$$

If p assigns zero probability to both H_i and H_j , then remove probability mass $0 < q < \min(2r/\sqrt{6}, p(H_k))$ from some H_k (since p is a probability distribution) and assign equal amounts to H_i and H_j to form p' . Then:

$$\|p - p'\| < \sqrt{\left(\frac{r}{\sqrt{6}}\right)^2 + \left(\frac{r}{\sqrt{6}}\right)^2 + \left(\frac{2r}{\sqrt{6}}\right)^2} = r, \quad (29)$$

where the first two terms under the radical are for H_i, H_j and the last is for H_k . So p' assigns non-zero probability both to H_i and to H_j and is in $S \subseteq O$. \square

For arbitrary points p_1, p_2, p_3 in \mathcal{P} , let $\overline{p_1 p_2}$ denote the convex hull of p_1, p_2 , and let $\Delta p_1 p_2 p_3$ denote the convex hull of p_1, p_2, p_3 :

$$\begin{aligned} \overline{p_1 p_2} &= \left\{ \sum_{k=1}^2 a_k p_k : \sum_{k=1}^2 a_k = 1, a_k \geq 0 \text{ for } k = 1, 2 \right\}; \\ \Delta p_1 p_2 p_3 &= \left\{ \sum_{k=1}^3 a_k p_k : \sum_{k=1}^3 a_k = 1, a_k \geq 0 \text{ for } k = 1, 2, 3 \right\}. \end{aligned}$$

For each complete answer H_i to \mathcal{Q} , let h_i be the credal state in which H_i has probability 1, which we call a *corner* of \mathcal{P} . For each pair of distinct complete answers H_i, H_j to \mathcal{Q} , let $\overline{h_i h_j}$ denote the set of credal states in which $H_i \vee H_j$ has probability 1, which we call an *edge* of \mathcal{P} . For each edge $\overline{h_i h_j}$ of \mathcal{P} , define the following set:

$$L_{ij} = \{p \in \overline{h_i h_j} : \mathbf{B}_p(\top) = H_i\}.$$

Lemma 2. *For each edge $\overline{h_i h_j}$ of \mathcal{P} , L_{ij} is a connected line segment in $\overline{h_i h_j}$ that contains h_i but not h_j , and contains at least one point distinct from h_i, h_j .*

Proof. Let $\overline{h_i h_j}$ be an arbitrary edge of \mathcal{P} . By non-skepticism, there exists non-empty open subset O of \mathcal{P} over which \mathbf{B} accepts H_i as strongest. Since O is non-empty and open, lemma 1 implies that there exists p in O that assigns nonzero probabilities to both H_i and H_j . So $p|_{H_i \vee H_j}$ is defined, which also assigns nonzero probabilities to both H_i and H_j and, thus, is distinct from corners h_i, h_j . Since p is in O , \mathbf{B} accepts H_i at p . Then \mathbf{B} also accepts H_i at $p|_{H_i \vee H_j}$, by Preservation and conditioning-tracking.

Furthermore, \mathbb{B} accepts H_i as strongest at $p|_{H_i \vee H_j}$, since \mathbb{B} is consistent and H_i is a complete answer. Then \mathbb{B} accepts H_i as strongest at h_i , since \mathbb{B} is corner-monotone and consistent. So L_{ij} contains two distinct points $p|_{H_i \vee H_j}$ and h_i . The set L_{ij} is connected because \mathbb{B} is corner-monotone and consistent. To see that L_{ij} does not contain corner h_j , note that L_{ij} and L_{ji} are disjoint (by definition), so it suffices to show that L_{ji} contains h_j . That follows from permuting i and j in the preceding argument that L_{ij} contains h_i . \square

For each triple of distinct corners h_i, h_j, h_m of \mathcal{P} , consider two-dimensional simplex $\Delta h_i h_j h_m$ (figure 10.a), relative to which points a, b, c, d are defined as follows. Let a




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

be the endpoint of L_{im} that is closest to h_m ; namely, a is the credal state in $\overline{h_i h_m}$ such that $a(H_m) = \sup\{p(H_m) : p \in L_{im}\}$. Similarly, let b be the endpoint of L_{jm} that is closest to h_m . By the preceding lemma, a and b are in the interiors of $\overline{h_i h_m}$ and $\overline{h_j h_m}$, respectively. Let credal state d be the intersection point of lines ah_j and bh_i . Let c be the unique credal state in $\overline{h_i h_j}$ such that $c h_m$ contains d .

Lemma 3. *Let h_i, h_j, h_m be distinct corners of \mathcal{P} . Consider two-dimensional simplex $\Delta h_i h_j h_m$, relative to which points a, b, c, d are defined as above. Then \mathbb{B} accepts H_i as strongest over the interior of $\Delta a d h_i$. Furthermore, \mathbb{B} accepts H_j as strongest over the interior of $\Delta b d h_j$.*

Proof. Consider an arbitrary point p in the interior of $\Delta a d h_i$ (figure 10.b). Argue as follows that \mathbb{B} accepts H_i as strongest at p . Since posterior state $p|_{-H_j}$ exists and falls inside L_{im} , $\mathbb{B}_{p|_{-H_j}}(\top) = H_i$. So, since \mathbb{B} tracks conditioning, $\mathbb{B}_p(\neg H_j) = H_i$. Then, since \mathbb{B} satisfies Inclusion, $\mathbb{B}_p(\top) \wedge \neg H_j \models H_i$. So we have only three possibilities for $\mathbb{B}_p(\top)$:

$$\mathbb{B}_p(\top) = \text{either } H_i, \text{ or } H_j, \text{ or } H_i \vee H_j,$$

since \mathbf{B} is consistent and the complete answers are mutually exclusive. Rule out the last two possibilities as follows. Suppose for reductio that $\mathbf{B}_p(\top) = H_j$ or $H_i \vee H_j$. Then, since \mathbf{B} satisfies Preservation, $\mathbf{B}_p(\neg H_i)$ must entail the conjunction of the prior belief state and new information, i.e.:

$$\mathbf{B}_p(\neg H_i) \models \mathbf{B}_p(\top) \wedge \neg H_i.$$

The left-hand side equals $\mathbf{B}_{p|\neg H_i}(\top)$ by conditioning-tracking, and the right-hand side equals H_j by the reductio hypothesis. So $\mathbf{B}_{p|\neg H_i}(\top) \models H_j$. But \mathbf{B} is consistent and H_j is a complete answer, so $\mathbf{B}_{p|\neg H_i}(\top) = H_j$. Since $p|\neg H_i$ is in $\overline{h_j h_m}$, $p|\neg H_i$ is in L_{jm} by the definition of L_{jm} . But that is impossible according to the choice of p as an interior point of $\triangle a d h_i$ (figure 10.b). Ruling out the last two possibilities for $\mathbf{B}_p(\top)$, we conclude that $\mathbf{B}_p(\top) = H_i$. So we have established the first statement. The second statement follows by symmetry. \square

Lemma 4. *Continuing from the preceding lemma, \mathbf{B} accepts H_i as strongest at h_i and over the interior of $\overline{c h_i}$. Furthermore, \mathbf{B} accepts H_j as strongest at h_j and over the interior of $\overline{c h_j}$.*

Proof. By lemma 2, $\mathbf{B}_{h_i} = H_i$. Let q be an arbitrary point in the interior of $\overline{c h_i}$. Then $q = p|_{H_i \vee H_j}$, for some point p in the interior of $\triangle a d h_i$ (figure 10.c). So $\mathbf{B}_p(\top) = H_i$, by the preceding lemma. Then, since \mathbf{B} satisfies Preservation,

$$\mathbf{B}_p(H_i \vee H_j) \models \mathbf{B}_p(\top) \wedge (H_i \vee H_j).$$

The left-hand side equals $\mathbf{B}_{p|_{H_i \vee H_j}}(\top)$ by conditioning-tracking, and the right-hand side equals H_i (since $\mathbf{B}_p(\top) = H_i$). So $\mathbf{B}_{p|_{H_i \vee H_j}}(\top) \models H_i$. Hence, $\mathbf{B}_{p|_{H_i \vee H_j}}(\top) = H_i$, since \mathbf{B} is consistent and H_i is a complete answer. Then, since $p|_{H_i \vee H_j} = q$, we have that $\mathbf{B}_q(\top) = H_i$, as required. So we have established the first statement. The second statement follows by symmetry. \square

Lemma 5. *Continuing from the preceding lemma, $\overline{h_i h_j}$ contains at most one point at which \mathbf{B} accepts $H_i \vee H_j$ as strongest.*

Proof. By the preceding lemma, for every point p in $\overline{h_i h_j}$, if $\mathbf{B}_p(\top) = H_i \vee H_j$, then $p = c$ (figure 10.c). \square

Lemma 6. *Every edge $\overline{h_i h_j}$ of \mathcal{P} contains at most one point at which \mathbf{B} accepts $H_i \vee H_j$ as strongest.*

Proof. Let $\overline{h_i h_j}$ be an arbitrary edge of \mathcal{P} . Then, since \mathcal{Q} contains at least three complete answers, there exists a third, distinct corner h_m of \mathcal{P} . Apply the preceding lemma to the simplex $\triangle h_i h_j h_m$, and we are done. \square

The preceding lemma establishes opinionation only for each one-dimensional edge of the simplex. The next step extends opinionation to the whole simplex.

Lemma 7. *B is opinionated.*

Proof. Suppose for reductio that B is not opinionated. Then, for some disjunction $H_i \vee H_j \vee X$ of at least two distinct answers H_i, H_j , and for some non-empty open subset O of \mathcal{P} , we have that B accepts $H_i \vee H_j \vee X$ as strongest over O . Since O is non-empty and open, lemma 1 implies that there exists credal state p in O that assigns nonzero probabilities to both H_i and H_j . Then there exists an Euclidean ball B of radius $r > 0$ centered on p that is contained in O . Transfer probability mass x from H_i to H_j to obtain credal state q , where $0 < x < \min(r/\sqrt{2}, p(H_i))$. Then, as in the proof of Lemma 1, q is in $B \subseteq O$, q assigns nonzero probabilities to H_i, H_j , and $\frac{p(H_i)}{p(H_j)} \neq \frac{q(H_i)}{q(H_j)}$. It follows that $p|_{H_i \vee H_j}$ and $q|_{H_i \vee H_j}$ are defined and distinct. Since p, q are in O , we have that B accepts $H_i \vee H_j \vee X$ as strongest at p, q . Hence, B accepts $(H_i \vee H_j \vee X) \wedge (H_i \vee H_j)$ as strongest at $p|_{H_i \vee H_j}, q|_{H_i \vee H_j}$, since B tracks conditioning and satisfies both Inclusion and Preservation. Note that $(H_i \vee H_j \vee X) \wedge (H_i \vee H_j) = H_i \vee H_j$. So B accepts $H_i \vee H_j$ as strongest at two distinct points in edge $\overline{h_i h_j}$, which contradicts the preceding lemma. \square

To conclude the proof of Theorem 1, recall that it suffices to derive that B is opinionated from the suppositions made in the beginning of the present section. So we are done.

B 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 $p(H_i|E) > 0$, $p(H_j|E) > 0$ and that $H_i \vee H_j$ entails E . It follows that $p|_{(H_i \vee H_j)} = p|_{E \wedge (H_i \vee H_j)}$, and that 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:

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

C Proof of Theorem 3

Right-to-Left Side. Let \mathbf{B} be driven by an odds-based assignment ($\prec_p: p \in \mathcal{P}$) of plausibility orders. The corner-monotonicity of \mathbf{B} follows from algebraic verification of the following fact: the odds of H_i to H_j increase monotonically if the credal state travels from p to corner h_i along the line $\overline{ph_i}$. To see that \mathbf{B} tracks conditioning (i.e. that $\mathbf{B}_p(E) = \mathbf{B}_{p|E}(\top)$), since \mathbf{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. \square

Left-to-Right Side. Suppose that \mathbf{B} is corner-monotone, tracks conditioning, and is Shoham-driven according to 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 indices i, j in I , define odds threshold $t_{ij} \in [0, \infty]$ and inequality $\triangleright_{ij} \in \{>, \geq\}$ by:

$$\text{Odds}_{ij} = \left\{ \frac{q(H_i)}{q(H_j)} : q \in \mathcal{P}, q(H_i) + q(H_j) = 1, H_i \prec_q H_j \right\}; \quad (30)$$

$$t_{ij} = \inf \text{Odds}_{ij}; \quad (31)$$

$$\triangleright_{ij} = \begin{cases} \geq & \text{if } t_{ij} \in \text{Odds}_{ij}, \\ > & \text{otherwise.} \end{cases} \quad (32)$$

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

$$H_i \prec_q H_j \iff q(H_i)/q(H_j) \triangleright_{ij} t_{ij}. \quad (33)$$

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 , equation (25) holds with respect to odds thresholds (31) and inequalities (32):

$$H_i \prec_p H_j \iff p(H_i)/p(H_j) \triangleright_{ij} t_{ij}. \quad (34)$$

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:

$$\begin{aligned} H_i \prec_p H_j &\iff H_i \prec_p |_{(H_i \vee H_j)} H_j \\ &\iff H_i \prec_{p|_{(H_i \vee H_j)}} H_j && \text{by theorem 2;} \\ &\iff H_i \prec_q H_j && \text{by defining } q \text{ as } p|_{(H_i \vee H_j)}; \\ &\iff q(H_i)/q(H_j) \triangleright_{ij} t_{ij} && \text{by (33);} \\ &\iff p(H_i)/p(H_j) \triangleright_{ij} t_{ij} && \text{since } q = p|_{(H_i \vee H_j)}. \end{aligned}$$

\square

D Proof of Propositions 1-4

Proof of Proposition 1. Consistency follows from the well-foundedness of plausibility orders. \square

Proof of Proposition 2. Consistency is an immediate consequence of proposition 1. So it suffices to show, for each p , that the relation $\mathbb{B}_{p|E}(\top) \models H$ between E and H satisfies Hypothetico-deductive Monotonicity (1) and Case Reasoning (2). That relation is equivalent to relation $\mathbb{B}_p(E) \models H$ between E and H (by tracking conditioning). Since \mathbb{B} is Shoham-driven, the relation is defined for fixed p 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). \square

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 $H_i \succ_p H_j \succ_p H_k \succ_p \dots$. Since $t_i > 1$ for all i in I , we have that $w_i p(H_i) < w_j p(H_j) < w_k p(H_k) < \dots$. So the sum is unbounded. But each weight is assumed to be no more than 1, so the sum of (unweighted) probabilities $p(H_i) + p(H_j) + p(H_k) + \dots$ is also unbounded—which contradicts the fact that p is a probability measure. \square

Proof of Proposition 4. To see that \mathbb{B} is a weighted odds threshold rule, argue as follows:

$$\mathbb{B}_p(\top) = \bigvee \{H_j \in \mathcal{Q} : H_j \text{ is minimal in } \prec_p\} \quad (35)$$

$$= \bigvee \{H_j \in \mathcal{Q} : \max_k w_k p(H_k)/w_j p(H_j) \not> t_j\} \quad (36)$$

$$= \bigwedge \{\neg H_i \in \mathcal{Q} : \max_k w_k p(H_k)/w_i p(H_i) > t_i\} \quad (37)$$

$$= \bigwedge \left\{ \neg H_i \in \mathcal{Q} : \frac{w_i p(H_i)}{\max_k w_k p(H_k)} < \frac{1}{t_i} \right\}. \quad (38)$$

Part 2, that the rule tracks conditioning, is an immediate consequence of theorem 3, because the rule is a special case of odds-based rules. To see that the rule is sensible, recall that the parameters are assumed to be restricted as follows: $1 < t_i \leq N$ and

$0 < w_i \leq 1$ for all i in I , where N is a positive integer. 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 corner-monotonicity of the rule is an immediate consequence of theorem 3, because the rule is a special case of odds-based rules. Non-skepticism is established as follows. Suppose that $i \neq j$. Define

$$r_{ij} = \inf \{ \|h_i - p\| : p \in \mathcal{P}, w_i p(H_i)/p(H_j) \leq N \}.$$

The value of r_{ij} is independent of the choice of j because of the symmetry of \mathcal{P} , so let r_i denote the invariant value of r_{ij} . Argue as follows that $r_i > 0$. Suppose for reductio that $r_i = 0$. Then there exists sequence $(p_n)_{n \in \omega}$ of points such that for all $n \in \omega$, $w_i p_n(H_i)/p_n(H_j) \leq N$ and $\lim_{n \rightarrow \infty} \|h_i - p_n\| = 0$. So $\lim_{n \rightarrow \infty} p_n(H_i) = 1$ and $\lim_{n \rightarrow \infty} p_n(H_j) = 0$. Then, for some sufficiently large m , we have that $w_i p_m(H_i)/p_m(H_j) > N$. But that contradicts $w_i p_m(H_i)/p_m(H_j) \leq N$, which is guaranteed by the construction. Therefore, $r_i > 0$. Let B_i be the Euclidean ball centered at corner h_i with radius r_i . Suppose that $k \neq i$. Then:

$$\begin{aligned} p \in B_i &\implies \|h_i - p\| < r_i \\ &\implies w_i p(H_i)/p(H_k) > N \\ &\implies w_i p(H_i)/p(H_k) > t_k \quad \text{since } N \geq t_k; \\ &\implies w_i p(H_i)/w_k p(H_k) > t_k \quad \text{since } w_k \leq 1. \end{aligned}$$

Hence H_i is accepted by the rule over B_i . 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 set. Consider the unique credal state p^* such that $w_1 p^*(H_1)/w_2 p^*(H_2) = 1$ and $p^*(H_j) = 0$, for all $j \neq 1, 2$. Suppose that $j \neq 1, 2$. Define:

$$\begin{aligned} a_j &= \inf \{ \|p^* - p\| : p \in \mathcal{P}, w_1 p(H_1)/p(H_j) \leq N \}; \\ b_j &= \inf \{ \|p^* - p\| : p \in \mathcal{P}, w_2 p(H_2)/p(H_j) \leq N \}; \\ c &= \inf \{ \|p^* - p\| : p \in \mathcal{P}, w_1 p(H_1)/w_2 p(H_2) > t_2 \}; \\ d &= \inf \{ \|p^* - p\| : p \in \mathcal{P}, w_2 p(H_2)/w_1 p(H_1) > t_1 \}. \end{aligned}$$

By the symmetry of \mathcal{P} , a_j and b_j do not depend on j , so let a denote the invariant value of a_j and similarly for b . (If \mathcal{Q} has only two complete answers, so that H_j does not exist, then let $a = b = 1$.) It follows that $a, b > 0$, by the same argument as in the non-skeptical case. Argue as follows that $c > 0$. Suppose for reductio that $c = 0$. Then there exists sequence $(p_n)_{n \in \omega}$ of points such that for all $n \in \omega$, $w_1 p_n(H_1)/w_2 p_n(H_2) > t_2$ and $\lim_{n \rightarrow \infty} \|p^* - p_n\| = 0$. So, $\lim_{n \rightarrow \infty} w_1 p(H_1)/w_2 p(H_2) = 1$. Then, since $t_2 > 1$, there exists sufficiently large m such that $w_1 p(H_1)/w_2 p(H_2) < t_2$. But that contradicts $w_i p_m(H_i)/p_m(H_j) > t_2$, which is guaranteed by the construction. Therefore, $c > 0$. By

symmetry, $d > 0$. Since a, b, c, d are strictly greater than 0, the following quantity also exceeds 0:

$$r^* = \inf \{a, b, c, d\}.$$

Let B^* be the Euclidean ball centered at p^* with radius r^* . It suffices to show that $H_1 \vee H_2$ is accepted as strongest by the rule over B^* . So it suffices to show that for each p in B^* and for each complete answer H_k distinct from H_i :

$$\begin{aligned} w_1 p(H_1) / w_k p(H_k) &> t_k; \\ w_2 p(H_2) / w_k p(H_k) &> t_k; \\ w_1 p(H_1) / w_2 p(H_2) &\leq t_2; \\ w_2 p(H_2) / w_1 p(H_1) &\leq t_1. \end{aligned}$$

The first two statements follow by the same argument as in the non-skeptical case. To prove the third statement, argue as follows:

$$\begin{aligned} p \in B^* &\implies \|p^* - p\| < r^* \\ &\implies \|p^* - p\| < c \\ &\implies w_1 p(H_1) / w_2 p(H_2) \leq t_2. \end{aligned}$$

The fourth statement follows by symmetry. □