Propositional Reasoning that Tracks Probabilistic Reasoning

Hanti Lin
Kevin T. Kelly*

May 12, 2011

Abstract

This paper concerns the extent to which propositional reasoning can be made consonant with probabilistic reasoning. A conditional acceptance rule (Leitgeb 2010) maps Bayesian credal state p and new information E to a propositional belief state $\mathsf{B}_{p,E}$ that should aptly reflect the underlying Bayesian belief state $p|_E$ that results from conditioning p on E. Consonance requires that $\mathsf{B}_{p,E} = \mathsf{B}_{p|_E,\top}$: i.e., that acceptance, followed by qualitative belief revision should end up in the same place as probabilistic conditioning followed by acceptance. Standard proposals for acceptance and belief revision do not fare well in that respect. The "Lockean" rule that accepts propositions above a probability threshold is subject to the familiar lottery paradox, and we show that it is also subject to new and more stubborn paradoxes when consonance with probabilistic conditioning is taken into account. Moreover, we show that the familiar AGM approach to belief revision cannot sensibly be made consonant with probabilistic conditioning. Finally, we present a plausible, alternative approach that avoids all of the paradoxes. It combines an odds-based acceptance rule proposed originally by Levi (1996) with a non-AGM belief revision rule proposed originally by Shoham (1987).

1 An Old Riddle of Uncertain Acceptance

There are two widespread practices for modeling the belief state of a doxastic subject—as a probability measure over propositions or as a single proposition corresponding to the conjunction of all propositions the subject believes. The two ideas fit together perfectly if propositional belief is understood to delimit the space of all propositions

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

over which probabilities are defined, but that restricts propositional belief to an arid, skeptical domain. It is natural, therefore, to view sufficiently high probability as a sufficient condition for propositional *acceptance*, so that accepted propositions can be uncertain. In honor of John Locke, who proposed something of that sort, the idea is referred to as the *Lockean acceptance rule*.

It is well known that, however the threshold for acceptance is set, the Lockean rule is inconsistent, 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 chance that a given ticket loses is 2/3, so it is accepted that each ticket loses. That implies that no ticket wins. But it is also accepted that some ticket wins because that proposition has probability one, so the conjunction of the accepted propositions is contradictory. In general, if t is the threshold, then a lottery with more than $1/(1-t)$ tickets suffices for acceptance of an inconsistent set of propositions.

One standard response to the lottery paradox is *contextualism*, the proposal to reset the threshold depending on the question at hand—if the question concerns a lottery of size n then the threshold should be set greater than $1 - 1/n$. That idea occasions its own difficulties and there is a literature of new objections and ad hoc responses (e.g. Pollock 1995).

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 standard of revision, namely, Bayesian updating. Mismatches between acceptance and Bayesian updating are another potential source of logical conundrums for non-skeptical acceptance. Such paradoxes can arise even when the threshold for acceptance is adjusted high enough to avoid the lottery paradox in the context of a given lottery.

For the first riddle, suppose that there are three tickets and the threshold is 3/4, so we are not even close to encountering the lottery paradox. But suppose 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 threshold that ticket 2 loses and that ticket 3 loses, from which it follows deductively that ticket 1 wins. Now entertain the new information that ticket 3 has been removed from the lottery, so it cannot win. It would be strange to retract one's acceptance of ticket 1 in that case, since ruling out a competing ticket seems only to provide further evidence that ticket 1 will win. But the Lockean rule does just that. By probabilistic 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 two thirds and the probability that ticket 2 wins is one third. Therefore, it is no longer accepted that ticket 1 wins, since that proposition has probability less than the threshold 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.

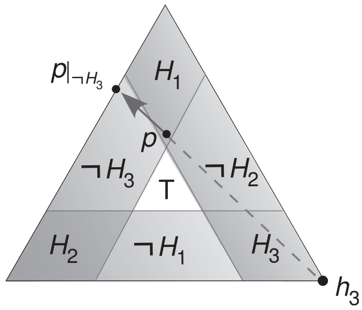

Figure 1: the first riddle

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 distrubtions that concentrate all probability on a single ticket. The assumed distribution p over tickets then corresponds to the point $(1/2, 1/4, 1/4)$ in the triangle. The conditional distribution $p|_{\neg H_3} = p(\cdot|\neg H_3)$ lies on a ray through p that originates from corner 3, which is the point $(2/3, 1/3, 0)$, because that is the line that holds the odds $H_1 : H_2$ constant. Each region in the triangle is labeled with the strongest proposition accepted by the probability measures it contains. 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 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, the paradox would not arise.

The riddle can be summarized by saying that the Lockean rule fails to satisfy the following, logical condition: if you already accept hypothesis H in light of background evidence K and you acquire new evidence E that is logically entailed by H, then you should continue to accept that H. Let $\mathsf{B}_p(H)$ abbreviate that hypothesis H is accepted in probabilistic credal state p. Then the principle violated may be stated succinctly as follows, where $p|_E$ denotes the conditional distribution $p(\cdot|E)$:

$$\mathsf{B}_p(H) \;\wedge\; H \models E \quad\Longrightarrow\quad \mathsf{B}_{p|_E}(H). \tag{1}$$

In the philosophy of science, one might call this principle *hypothetico-deductive monotonicity*, in the sense that observing a consequence of a hypothesis is usually thought to be evidence in favor of that hypothesis (the probability of the hypothesis can only go up by Bayesian conditioning) and the principle plausibly says not to drop acceptance of a hypothesis when one observes just what it would predict. In the logic literature, the principle is called *cautious monotonicity*, which is needlessly coy about the precise nature of the caution.

One Lockean response to the preceding riddle is to adopt a higher threshold for disjunctions than for conjunctions (figure 2). But now one encounters a different and,

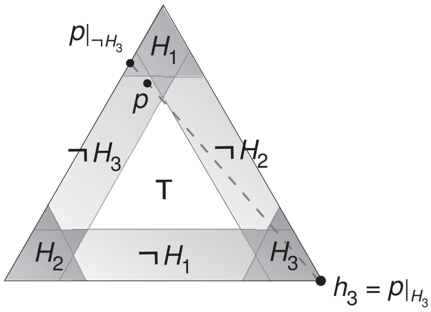

Figure 2: second riddle

in a sense, complementary riddle. 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 given that ticket 3 wins and also accepts that ticket 2 loses given that ticket 3 loses, but fails to accept that ticket 2 loses until one actually learns what happens with ticket 3. That violates the following rule of *reasoning by cases.*[1]

$$\mathsf{B}_{p|_E}(H) \;\wedge\; \mathsf{B}_{p|_{\neg E}}(H) \quad\Longrightarrow\quad \mathsf{B}_p(H). \tag{2}$$

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—the first riddle picks up where the second riddle leaves off and there are even thresholds that generate both riddles at once.[2] Unlike the lottery paradox, which must increase the number

[1]The principle is also related 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.

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

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. 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 revert constantly to the underlying probabilities. In the first riddle, one cannot rely on the entailment of E by H in order to determine whether to retain or reject 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 a desired objective in any case). Although one accepts propositions, 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 rules. So the riddles are are symptoms of a mismatch between purely propositional reasoning and the underlying evolution of probabilistic credal states.

The riddles would be met by a way of correlating propositional reasoning with probabilistic reasoning that would allow one to enter the propositional system, kick away the underlying probabilities, and still end up exactly where a persistent peeker at the underlying probabilities would end up—i.e., by a pre-established harmony between propositional and probabilistic reasoning. The realization of such a perfect harmony, without peeking, is far more challenging than merely to avoid acceptance of mutually inconsistent propositions. We show that there is no possible way to achieve that ideal by means of the popular AGM approach to propositional belief revision. Then we exhibit a broad collection of rules that do achieve perfect harmony with probabilistic 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 called a *question* and let the *complete answers* to $\mathcal{Q}$ be $H_1, \ldots, H_i, \ldots$. The set of all (possibly incomplete) answers to $\mathcal{Q}$ is $\mathcal{A}$, the least 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.* For example, in the three ticket lottery, $\mathcal{Q} = \{H_1, H_2, H_3\}$, where H_i says that ticket i wins and $\mathcal{P}$ is the triangle (simplex) of probability distributions over the three answers.

5 Conditional Acceptance

Recall that $\mathsf{B}_p(H)$ means that proposition H is accepted in credal state p. When acceptance is viewed as a rule that associates accepted propositions with probabilistic credal states, one can conveniently fold together the problem of acceptance with the problem of revising what has been accepted in light of new information by letting $\mathsf{B}_p(H|E)$ mean that H is *accepted conditionally* on E at p.[3] Unconditional acceptance of H at p is naturally defined by:

$$\mathsf{B}_p(H) \quad = \quad \mathsf{B}_p(H|\top), \tag{5}$$

where $\top$ is the tautology. Cautious monotonicity, for example, can now be stated in terms of conditional acceptance, rather than in terms of Bayesian conditioning. Think of K as information already received before E.

$$\mathsf{B}_p(H|K) \ \wedge \ H \models E \quad \Longrightarrow \quad \mathsf{B}_p(H|E \wedge K), \tag{6}$$

Reasoning by cases has a similar statement.

$$\mathsf{B}_p(H|E \wedge K) \ \wedge \ \mathsf{B}_p(H|\neg E \wedge K) \quad \Longrightarrow \quad \mathsf{B}_p(H|K). \tag{7}$$

It is convenient re-express the above idea in terms of belief states. Assuming that the set of all conditionally accepted propositions is closed under conjunction, there exists the strongest proposition conditionally accepted on E at p:

$$\mathsf{B}_{p,E} = \bigwedge\{H \in \mathcal{A} : \mathsf{B}_p(H|E)\}. \tag{8}$$

Refer to $\mathsf{B}_{p,E}$ as the *conditional belief state* on E that one has at p. Then one can naturally define *unconditional belief state* at p as follows:

$$\mathsf{B}_p \quad = \quad \mathsf{B}_{p,\top}\,. \tag{9}$$

Then, to be accepted is to be entailed by the belief state:

$$\mathsf{B}_p(H|E) \quad \Longleftrightarrow \quad \mathsf{B}_{p,E} \models H; \tag{10}$$

$$\mathsf{B}_p(H) \quad \Longleftrightarrow \quad \mathsf{B}_p \models H. \tag{11}$$

[3] For readers interested in epistemic conditionals and in default reasoning, the following, conditional acceptance *Ramsey tests* suffice to relate the results in this paper to their concerns:

$$p \Vdash E \Rightarrow H \quad \Longleftrightarrow \quad \mathsf{B}_p(H|E); \tag{3}$$

$$E \mid\!\sim_p H \quad \Longleftrightarrow \quad \mathsf{B}_p(H|E). \tag{4}$$

We are indebted to Hannes Leitgeb for the idea of framing our discussion in terms of conditional acceptance, 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 the probabilistic Ramsey test, which involves more cumbersome notation and an irrelevant commitment to an epistemic interpretation of conditionals.

6 Bayesian Conditional Acceptance

According to the desired, pre-established harmony, the propositions conditionally accepted on E at prior state p should always agree with the unconditionally accepted propositions at posterior state $p|_E$. Accordingly, say that acceptance rule B is *Bayesian* if and only if:

$$\mathsf{B}_{p,E} \quad = \quad \mathsf{B}_{p|_E}, \tag{12}$$

for each proposition H, and for each credal state p and proposition E such that $p(E) > 0$ (so that $p|_E$ is defined). Consider two procedures for a subject in credal state p who entertains supposition or new information E and wants to determine whether $\mathsf{B}_p(H|E)$ (figure 3). According to the first, *probabilistic* procedure, she uses probabilistic

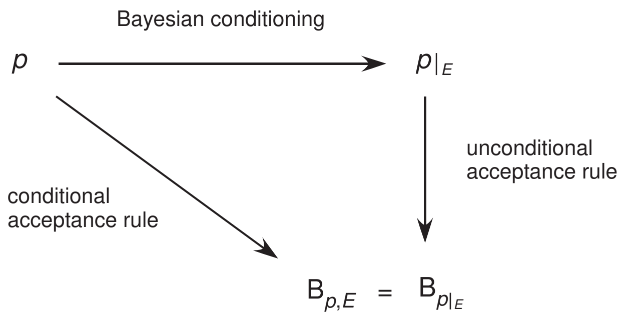

Figure 3: Bayesian acceptance

conditioning to compute $p|_E$ and and then revises her current belief state B_p to $\mathsf{B}_{p|_E}$ and checks whether $\mathsf{B}_{p|_E} \models H$. According to the second procedure, she just checks directly whether $\mathsf{B}_p(H|E)$, namely whether $\mathsf{B}_{p,E} \models H$. If B is Bayesian, the two procedures agree.

Conditional acceptance is closely related to the more familiar topic of belief *revision* (Alchourrón, Gärdenfors, and Makinson 1985). A *belief revision operator* $*$ for belief state B is a mapping from $\mathcal{A}$ to $\mathcal{A}$ that sends each proposition E to a proposition $B * E$ understood as the posterior belief state that results from revising prior belief state B in light of information E. A conditional acceptance rule B naturally associates each credal state p with a belief revision operator $*_p$ defined by:

$$\mathsf{B}_p *_p E \quad = \quad \mathsf{B}_{p,E}\,. \tag{13}$$

Then, by (12), (13), the commutative diagram in figure 3 becomes the more natural commutative diagram depicted in figure 4. In the new picture, acceptance rule B is Bayesian if and only if its associated belief revision operator commute with Bayesian conditioning.

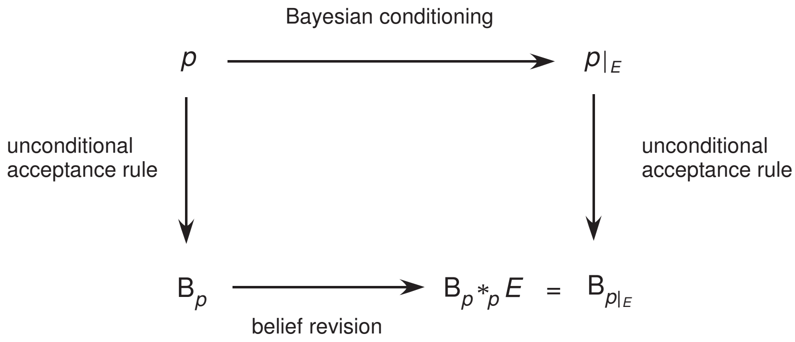

Figure 4: Bayesian belief revision

7 Accretive Acceptance

It remains to specify what would count as a propositional approach to conditional acceptance or belief revision that does not peek at probabilities to decide what to do. An obvious and popular idea is simply to conjoin new information with one's old beliefs, as long as no contradiction results. Say that belief revision operator $*$ satisfies *Inclusion* if and only if:

$$B \wedge E \models B * E.$$

Say that $*$ satisfies *Preservation* if and only if:

$$B \text{ is consistent with } E \implies B * E \models B \wedge E.$$

These axioms are widely understood to be the least controversial axioms of the popular AGM theory of belief revision due to Harper (1975) and Alchourrón, Gärdenfors, and Makinson (1985). They jointly imply the view just mentioned—that consistent belief revision should always proceed by conjunction. Say that a belief revision operator is *accretive* if and only if it satisfies Inclusion and Preservation and that a conditional acceptance rule is *accretive* if and only if each of its associated belief revision operators is accretive.

8 Accretive, Bayesian Acceptance is Trivial

Accretion sounds plausible enough when beliefs are certain, but it is less intuitive when beliefs are accepted at probabilities less than 1. For example, suppose that one's current belief state is $H_2 \vee H_3$. That should be possible when H_2 is very probable but not quite probable enough to be accepted in its own right and H_3 is just enough more probable than H_1 to warrant acceptance of $H_2 \vee H_3$, as for figure 5. Now suppose one learns that

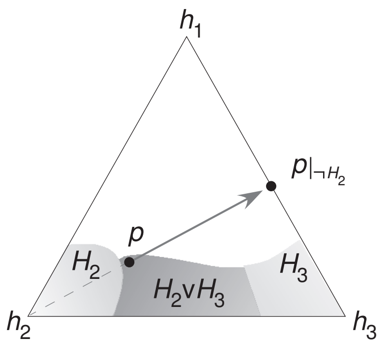

Figure 5: How Preservation may fail plausibly

H_2 is false. Revision by accretion insists that one *deduce* H_3, from $H_2 \vee H_3$ and $\neg H_2$, but that seems wrong—H_2 was the main reason for accepting $H_2 \vee H_3$, and that reason just evaporated. This intuitive doubt agrees with Bayesian conditioning. Bayesian conditioning preserves odds (probability ratios) between alternatives compatible with the new information—in this case H_1 and H_3. Therefore, since H_3 is assumed to be only slightly more probable than H_1, it remains only slightly more probable than H_1 after learning that H_2 is false (cf. the posterior $p|_{\neg H_2}$ in figure 5). So one's posterior credal state is more aptly represented by $H_1 \vee H_3$ than by H_3, whereas accretive revision requires the latter. Note that, in this case, it is the Preservation axiom that fails.

The preceding intuitions are vindicated by the following no-go theorem. Say that B is *non-skeptical* if and only if each answer to $\mathcal{Q}$ is accepted over some open neighborhood of credal states in $\mathcal{P}$. One would hope, at least, that each H_i is accepted over some open disk around the credal state that assigns unit probability to H_i. At the opposite extreme, say that B is *non-opinionated* if and only if there is an open subset of $\mathcal{P}$ over which an incomplete, disjunctive answer is accepted. Say that B is *consistent* if and only if the inconsistent proposition $\perp$ is accepted by B 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.[4] Aside from the intuitive merits of these properties, all proposed acceptance rules we are aware of satisfy all of them. Say that an acceptance rule is *sensible* if and only if it has all of the properties just described. Aside from the intuitive appeal of the sensibility requirement, all proposed acceptance rules we are aware of are sensible. Then we have:

Theorem 1 (no-go theorem for accretive acceptance). *Let question $\mathcal{Q}$ have at*

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

least three answers. Then no sensible, Bayesian, conditional acceptance rule is accretive.

Since AGM rules are accretive, we also have:

Corollary 1 (no-go theorem for AGM acceptance). *Let question Q have at least three answers. Then no sensible, Bayesian, conditional acceptance rule is AGM.*

The theorems extend the intuitive misgivings about accretiveness discussed above. For example, one might attempt to force accretivness to track Bayesian conditioning by never accepting what one would fail to accept after conditioning on compatible evidence. But it is still the case that no sensible such proposal is fully Bayesian.[5]

9 The Importance of Odds

From the no-go theorems, it is clear that any sensible rule that tracks Bayesian 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 5, which we reproduce in figure 6. If, instead, one is

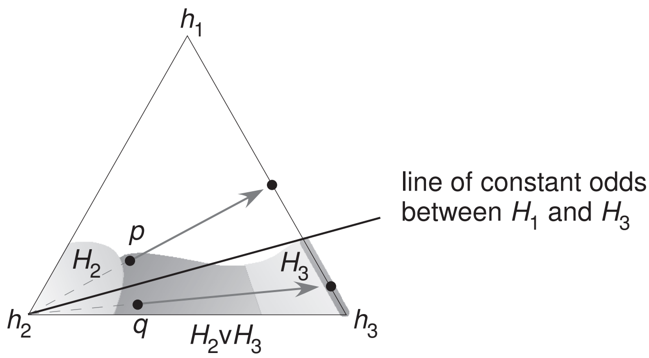

Figure 6: Line of constant odds

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 for sure. That intuition agrees with Bayesian conditioning, again. Since Bayesian conditioning preserves odds, H_3 continues to have significantly

[5] A sensible implementation of this idea has been presented by Leitgeb (2010). But it is sensible only because it fails to be fully Bayesian. Leitgeb is careful to point out that only one side of the the Bayesian property is satisfied by his rule.

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 6 gives the odds threshold between H_1 and H_3 that determines whether Preservation holds or fails under new information $\neg H_2$. Therefore, a sensible Bayesian acceptance rule that systematically relaxes Preservation should be based on odds thresholds.

Here is one *sensible* such rule. Say that answer H_i is (significantly) more plausible than answer H_j at credal state p if and only if the odds of the former to the latter at p passes a stipulated threshold, say 3. Answers of zero probability will be taken as too improbable to worth comparison in terms of plausibility. Draw the lines of constant odds 2, as in figure 7.a. Then the partial order of relative plausibility at p can be

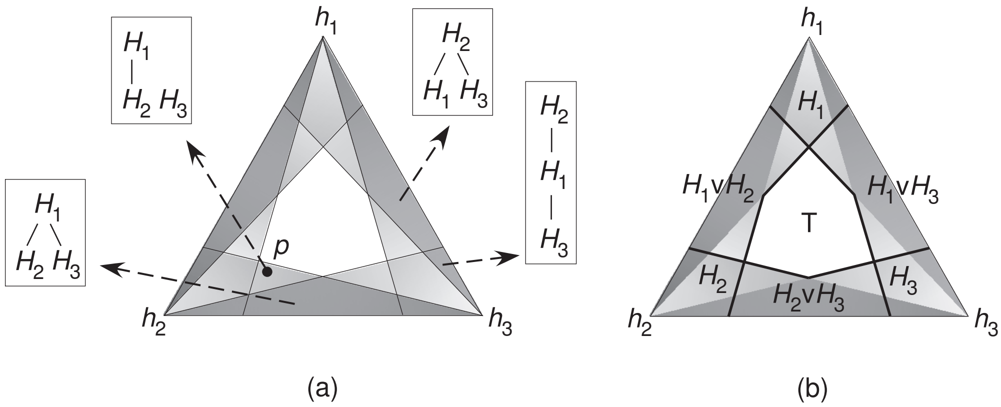

Figure 7: A rule based on odds thresholds

determined simply by checking whether p falls on one side or the other of a constant odds line (more plausible answers are depicted lower). Now define conditional acceptance rule B as follows. First, let propositional belief state B_p at p be the disjunction of the most plausible answers at p. So the zone for accepting a propositional belief state is bounded by the constant odds lines, as depicted in figure 7.b. From the figure, it is evident that the rule is sensible. To complete the definition of B, let conditional belief states $\mathsf{B}_{p,E}$ be defined by belief state B_p and formula (12), which ensures that the resulting rule is Bayesian.

Axiom Preservation is sometimes violated and sometimes satisfied by the acceptance rule just defined (figure 8), for reasons similar to those discussed in the preceding example (figure 6). At p, for example, Preservation is violated 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 holds at q.

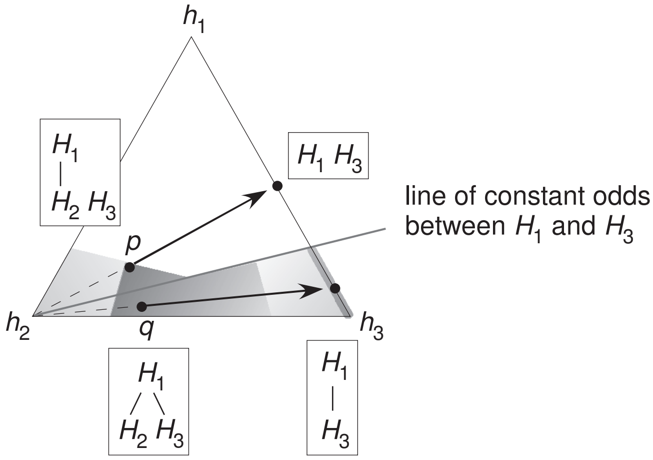

Figure 8: Preservation and odds

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 suppose that $H_2 \vee H_3$ is true because H_3 is true. Then acceptance of true belief $H_2 \vee H_3$ at credal state p is based mainly on false hypothesis H_2, which seems to be a Gettier (1963) case (justified true belief based on a materially defective justification), whereas robust acceptance at credal state q is not. It would be implausible to define propositional belief so robustly that Gettier cases are impossible, since it seems that almost every true, real-world belief might be a Gettier case based on hidden holograms or other standard, epistemological legerdemain.

The example just discussed contains all the essential ideas of our general theory of uncertain acceptance, which we now proceed to develop.

10 Shoham Acceptance

The concept of relative plausibility introduced in the preceding discussion can be formalized as follows. A *plausibility order* assigned to credal state p is a well-founded partial order $\mathcal{O}_p$ defined on the set $\{H_i \in \mathcal{Q} : p(H_i) > 0\}$ of nonzero-probability answers. We write $H_i \prec_{\mathcal{O}_p} H_j$ to mean that answer H_i is more plausible than answer H_j with respect to order $\mathcal{O}_p$. The posterior plausibility order $\mathcal{O}_{p,E}$ in light of information E is defined as the restriction of prior order $\mathcal{O}_p$ to the set of answers that are logically compatible with E (figure 9). That process is called *Shoham revision* on E (Shoham 1987). The posterior belief state, denoted by $\overline{\mathcal{O}}_{p,E}$, is defined as the disjunction of the answers that are most plausible with respect to $\mathcal{O}_{p,E}$. Subscript E is suppressed when

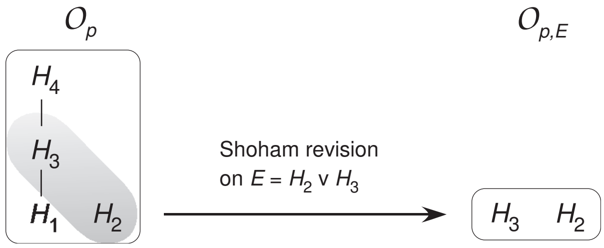

Figure 9: Shoham revision

E is the tautology.

It is routine to verify that Shoham revision always satisfies axiom Inclusion. But Shoham revision may violate the Preservation axiom. In figure 9, for example, the prior belief state $H_1 \vee H_2$ is consistent with the new information $H_2 \vee H_3$, but the the prior belief $H_1 \vee H_2$ is dropped in the posterior belief state $H_2 \vee H_3$.

Say that conditional acceptance rule B is *driven* by assignment $(\mathcal{O}_p : p \in \mathcal{P})$ of plausibility orders if and only if:

$$\mathsf{B}_{p,E} \quad = \quad \overline{\mathcal{O}}_{p,E}, \tag{14}$$

for all credal states p and propositions E. Rule B is called *Shoham-driven* if and only if B is driven by some assignment of plausibility orders. When each credal state is associated with a plausibility order, Shoham revision is defined in terms of logical compatibility. So belief revision based on Shoham revision does seem to define a propositional "space of reasons" that is not a mere epiphenomenon of probabilistic reasoning.

For example, let plausibility order $\mathcal{O}_p$ be defined by odds threshold 3 as follows:

$$H_i \prec_{\mathcal{O}_p} H_j \quad \iff \quad p(H_i)/p(H_j) > 3. \tag{15}$$

Let rule B be driven by assignment $(\mathcal{O}_p : p \in \mathcal{P})$. Then rule B is sensible and Bayesian, due to proposition 4 below. The unconditional belief state B_p can be expressed by:

$$\mathsf{B}_p \quad = \quad \bigwedge \left\{ \neg H_i \in \mathcal{Q} : \frac{p(H_i)}{\max_k p(H_k)} < \frac{1}{3} \right\}, \tag{16}$$

which is a consequence of proposition 4 below. So rule B rejects H_i at p if and only if the odds of H_i to the most probable answer is too low at p.

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

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

To guard against all riddles—old and new—it suffices to require further that the rules are Bayesian:

Proposition 2 (**riddle-free acceptance**). *Each Bayesian, Shoham-driven conditional acceptance rule is consistent and satisfies cautious monotonicity and reasoning by cases.*

We also have:

Theorem 2. *Suppose that conditional acceptance rule* B *is Bayesian and Shoham-driven—say, driven by assignment* $(\mathcal{O}_p)_{p \in \mathcal{P}}$ *of plausibility orders. Then for each credal state* p *and each proposition* E *such that* $p(E) > 0$*, we have:*

$$\mathcal{O}_{p|_E} \quad = \quad \mathcal{O}_{p,E}. \tag{17}$$

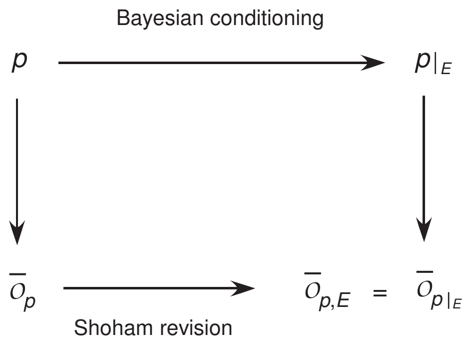

Figure 10: Shoham revision commutes with Bayesian conditioning

That is, probabilistic conditioning on E followed by assignment of a plausibility order to $p|_E$ leads to the same result as assigning a plausibility order to p and Shoham revising that order on E (figure 10).

11 Shoham-Driven Acceptance Based on Odds

The assignment (15) of plausibility orders and its associated acceptance rule (16) have a single, unform 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 $(\mathcal{O}_p : p \in \mathcal{P})$ of plausibility orders is *based on* assignment $(t_i : i \in I)$ of odds thresholds if and only if:

$$H_i \prec_{\mathcal{O}_p} H_j \quad \Longleftrightarrow \quad p(H_i)/p(H_j) > t_j. \tag{18}$$

Say that B is a *odds threshold* rule based on $(t_i : i \in I)$ if and only if the unconditional acceptance rule induced by B is defined by:

$$\mathsf{B}_p \quad = \quad \bigwedge \left\{ \neg H_i \in \mathcal{Q} : \frac{p(H_i)}{\max_k p(H_k)} < \frac{1}{t_i} \right\}. \tag{19}$$

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 "anything else". Let $(w_i : i \in I)$ be an assignment of *weights* w_i to answers H_i. Say that assignment $(\mathcal{O}_p : p \in \mathcal{P})$ 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:

$$H_i \prec_{\mathcal{O}_p} H_j \quad \Longleftrightarrow \quad w_i\, p(H_i)/w_j\, p(H_j) > t_j. \tag{20}$$

The range of t_i and w_i should be appropriately restricted:

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 the relation defined by formula (20), for each p in $\mathcal{P}$, 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 unconditional acceptance rule induced by B is defined by:[6]

$$\mathsf{B}_p \quad = \quad \bigwedge \left\{ \neg H_i \in \mathcal{Q} : \frac{w_i\, p(H_i)}{\max_k w_k\, p(H_k)} < \frac{1}{t_i} \right\}. \tag{21}$$

When all the weights w_i are equal, order (20) and rule (21) are reduced to order (18) and rule (19). Then we have:

Proposition 4 (sufficient condition for being sensible and Bayesian). *Continuing proposition 3, suppose that conditional acceptance rule B is driven by 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 is Bayesian.*

So a Shoham-driven rule can easily be sensible and Bayesian: it suffices that the plausibility orders encode information about odds and weights in the sense defined above.

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

For a representation theorem, we proceed to the next and final level of generality. Then weights in formula (20) can be absorbed into odds without loss of generality. That is, note that:

$$H_i \prec_{\mathcal{O}_p} H_j \iff w_i p(H_i)/w_j p(H_j) > t_j, \tag{22}$$

$$\iff p(H_i)/p(H_j) > t_j(w_j/w_i), \tag{23}$$

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

$$t_{ij} = t_j(w_j/w_i), \tag{24}$$

where $i \neq j$. Now, allow double-indexed odds thresholds t_{ij} that are *not* factorizable into single-indexed thresholds and weights according to equation (24), and to allow double-indexed inequalities, which can be strict or weak. Those enable us to express all the Bayesian, Shoham-drive, and corner-monotone rules. An assignment of *double-indexed odds thresholds* is of the form $t = (t_{ij} : i, j \in I \text{ and } i \neq j)$, where each t_{ij} is in interval $[0, \infty]$ (including the 0 and ∞). An assignment of *double-indexed inequalities* is of the form $\rhd = (\rhd_{ij} : i, j \in I \text{ and } i \neq j)$, where each inequality $\rhd_{ij}$ is either strict $>$ or weak $\geq$. Say that assignment $(\mathcal{O}_p : p \in \mathcal{P})$ of plausibility orders is *based* on t and $\rhd$ if and only if each plausibility order $\mathcal{O}_p$ is expressed by:

$$H_i \prec_{\mathcal{O}_p} H_j \iff p(H_i)/p(H_j) \rhd_{ij} t_{ij} . \tag{25}$$

When assignment $(\mathcal{O}_p : p \in \mathcal{P})$ of plausibility orders can be expressed in that way, say that it is *odds-based*.

It is no accident that all the Bayesian, Shoham-driven rules we have examined are odds-based:

Theorem 3 (**representation of Bayesian, Shoham-driven rules**). *Let* B *be a conditional acceptance rule. Then* B *is corner-monotone, Bayesian, and Shoham-driven if and only if* B *is driven by an odds-based assignment of plausibility orders.*

12 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 not hard to seek: just adopt Shoham revision. When Shoham revision tracks Bayesian conditioning perfectly, 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, the Lottery paradox as the old, static riddle has a new solution motivated by dynamic concerns.

13 Acknowledgements

The authors are indebted to Hannes Leitgeb for friendly discussion and for his elegant concept of conditional 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. We are also indebted to David Etlin for encouraging us to make explicit which AGM axioms are responsible for the no-go theorem. Finally, we are indebted for the discussion with Horacio Arló-Costa and Arthur Paul Pedersen.

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.

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.

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*.

Pollock, J. (1995) *Cognitive Carpentry*, Cambridge, MA: MIT Press.

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 Proc. Logics in Computer Science*, 275-279, Ithaca, N.Y.

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 theorem 1, suppose that rule B is consistent, corner-monotone, Bayesian, and pre-AGM-driven (i.e. satisfies axioms Inclusion and Preservation). 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 11.a). Let $\overline{h_i\, h_m}$ denote the one-dimensional subspace $\{p \in \mathcal{P} : p(H_i)+p(H_j) = 1\}$, and

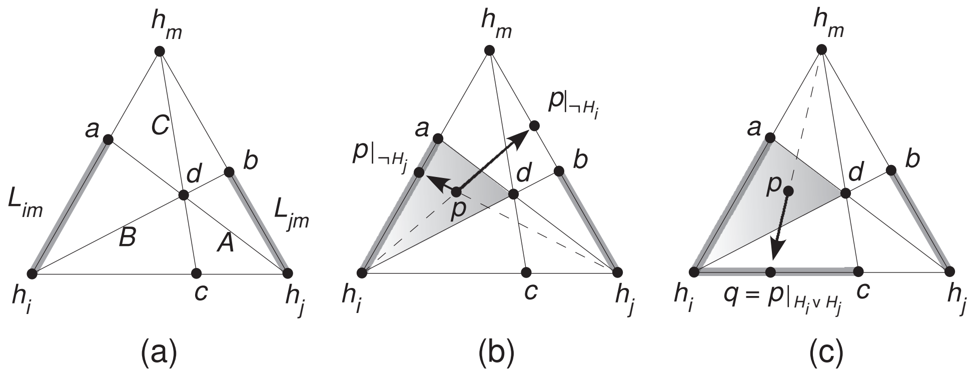

Figure 11: why Bayesian AGM is trivial

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:

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

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 is Bayesian, 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 $\bot$. 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:

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

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

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

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 11.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 is Bayesian and satisfies Inclusion, we have that:

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

(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 is Bayesian 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 11.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 is Bayesian 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:

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

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 11.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 11.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 is Bayesian 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 is Bayesian, 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 of Theorem 1. Since the last lemma states that B is opinionated, we are done. □

B Proof of Theorem 2

Proof of Theorem 2. The domains of $\mathcal{O}_{p|_E}$ and $\mathcal{O}_{p,E}$ coincide, because each plausibility order $\mathcal{O}_q$ 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 $\mathcal{O}_{p|_E}$, we have that $p(H_i|E) > 0$, $p(H_j|E) > 0$. Since both answers are in $\mathcal{O}_{p,E}$, 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 $H_i \prec_{\mathcal{O}_{p|_E}} H_j$

if and only if $H_i \prec_{\mathcal{O}_{p,E}} H_j$, as follows:

$$\begin{array}{lll}
 & H_i \prec_{\mathcal{O}_{p|_E}} H_j & \\
\iff & \overline{\mathcal{O}}_{p|_E,(H_i \vee H_j)} = H_i & \text{by the definition of Shoham conditioning;} \\
\iff & \mathsf{B}_{p|_E,(H_i \vee H_j)} = H_i & \text{by being Shoham-driven;} \\
\iff & \mathsf{B}_{p|_{(E \wedge (H_i \vee H_j))}} = H_i & \text{by being Bayesian;} \\
\iff & \mathsf{B}_{p|_{(H_i \vee H_j)}} = H_i & \text{since } p|_{(H_i \vee H_j)} = p|_{E \wedge (H_i \vee H_j)}; \\
\iff & \mathsf{B}_{p,(H_i \vee H_j)} = H_i & \text{by being Bayesian;} \\
\iff & \overline{\mathcal{O}}_{p,(H_i \vee H_j)} = H_i & \text{by being Shoham-driven;} \\
\iff & H_i \prec_{\mathcal{O}_{p,(H_i \vee H_j)}} H_j & \text{by the defintiion of Shoham conditoning;} \\
\iff & H_i \prec_{\mathcal{O}_{p,E}} H_j & \text{since } H_i \vee H_j \text{ entails } E.
\end{array}$$

□

C Proof of Theorem 3

Proof of the Right-to-Left Side of Theorem 3. Let B be driven by an odds-based assignment $\{\mathcal{O}_p : p \in \mathcal{P}\}$ of plausibility orders. Then, *a fortiori*, B is Shoham-driven. That B is corner-monotone follows from routine, algebraic verification. To see that B is Bayesian (i.e. that $\mathsf{B}_{p,E} = \mathsf{B}_{p|_E}$), by (14) it suffices to show that $\overline{\mathcal{O}}_{p,E} = \overline{\mathcal{O}}_{p|_E}$, so it suffices to show that an answer is most plausible in $\mathcal{O}_{p,E}$ if and only if it is most plausible in $\mathcal{O}_{p|_E}$, which follows from the odds-based definition of $\mathcal{O}_p$ and preservation of odds by Bayesian conditioning. □

Proof of the Left-to-Right Side of Theorem 3. Suppose that B is corner-monotone, Bayesian, and Shoham-driven, namely driven by assignment $(\mathcal{O}_p)_{p \in \mathcal{P}}$ of plausibility orders. It suffices to show that $(\mathcal{O}_p)_{p \in \mathcal{P}}$ 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:

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

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

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

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 $q(H_i) + q(H_j) = 1$,

$$H_i \prec_{\mathcal{O}_q} H_j \iff q(H_i)/q(H_j) \rhd_{ij} t_{ij}. \tag{29}$$

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

$$H_i \prec_{\mathcal{O}_p} H_j \iff p(H_i)/p(H_j) \ \rhd_{ij} \ t_{ij} \ . \tag{30}$$

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:

$$\begin{array}{lll} & H_i \prec_{\mathcal{O}_p} H_j & \\ \iff & H_i \prec_{\mathcal{O}_{p,(H_i \vee H_j)}} H_j & \text{by Shoham conditioning;} \\ \iff & H_i \prec_{\mathcal{O}_{p|_{(H_i \vee H_j)}}} H_j & \text{since } \mathcal{O}_{p,(H_i \vee H_j)} = \mathcal{O}_{p|_{(H_i \vee H_j)}}\text{, by proposition 2;} \\ \iff & H_i \prec_{\mathcal{O}_q} H_j & \text{by defining } q \text{ as } p|_{(H_i \vee H_j)}; \\ \iff & q(H_i)/q(H_j) \ \rhd_{ij} \ t_{ij} & \text{by (29);} \\ \iff & p(H_i)/p(H_j) \ \rhd_{ij} \ t_{ij} & \text{since } q = p|_{(H_i \vee H_j)}. \end{array}$$

□

D Proof of Propositions 1-4

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

Proof of Proposition 2. Consistency follows as an immediate consequence of proposition 1. So it suffices to show, for each p, that the relation $\mathsf{B}_{p|_E}(H)$ between propositions E and H satisfies cautious monotonicity and reasoning by cases. That relation is equivalent to relation $\mathsf{B}_p(H|E)$ (by being Bayesian), which is equivalent to $\overline{\mathcal{O}}_{p,E} \models H$ (by being Shoham-driven). The latter is defined by plausibility order $\mathcal{O}_p$ assigned to p, which is a special case of the so-called preferential models and, thus, validates non-monotonic logic system P (Kraus, Lehmann, and Magidor 1990). Then it suffices to note that system P includes cautious monotonicity and reasoning by cases (which is a special case of axiom Or). □

Proof of Proposition 3. The specified range of odds thresholds t_i guarantees that the defined relation is a partial order. The specified range of weights w_i guarantees well-foundedness. □

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

$$\begin{aligned}
\mathsf{B}_p &= \overline{\mathcal{O}}_p & (31)\\
&= \bigvee\{H_j \in \mathcal{Q} : H_j \text{ is most plausible in } \mathcal{O}_p\} & (32)\\
&= \bigvee\{H_j \in \mathcal{Q} : \max_k w_k p(H_k)/w_j p(H_j) \not> t_j\} & (33)\\
&= \bigwedge\{\neg H_i \in \mathcal{Q} : \max_k w_k p(H_k)/w_i p(H_i) > t_i\} & (34)\\
&= \bigwedge\left\{\neg H_i \in \mathcal{Q} : \frac{w_i\, p(H_i)}{\max_k w_k\, p(H_k)} < \frac{1}{t_i}\right\}. & (35)
\end{aligned}$$

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 (35) is satisfied. Then, by the upper bound on odds thresholds t_i and the strict inequality in formula (35), 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 (35) 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 is Bayesian, is a special case of the left-to-right side of theorem 3. □