PhD Dissertation: Propositional Reasoning that Tracks Probabilistic Reasoning

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

Abstract

Bayesians model one's doxastic state by subjective probabilities. But in traditional epistemology, in logic-based artificial intelligence, and in everyday life, one's doxastic state is usually expressed in a qualitative, binary way: either one accepts (believes) a proposition or one does not. What is the relationship between qualitative and probabilistic belief? I show that, besides the familiar lottery paradox (Kyburg 1961), there are two new, diachronic paradoxes that are more serious. A solution to the paradoxes, old and new, is provided by means of a new account of the relationship between qualitative and probabilistic belief. I propose that propositional beliefs should crudely but aptly represent one's probabilistic credences. Aptness should include responses to new information so that propositional belief revision *tracks* Bayesian conditioning: if belief state B aptly represents degrees of belief p then the revised belief state $K * E$ should aptly represent the conditional degrees of belief $p(\cdot|E)$. I explain how to characterize synchronic aptness and qualitative belief revision to ensure the tracking property in the sense just defined. I also show that the tracking property is impossible if acceptance is based on thresholds or if qualitative belief revision is based on the familiar AGM belief revision theory of Alchourrón, Gärdenfors, and Makinson (1985).

Contents

1 Introduction

Bayesians model an agent's doxastic state as a *credal state* that assigns probabilistic degrees of belief or credences to all propositions in some sufficiently rich collection of propositions. But in mainstream epistemology, in logic-based artificial intelligence,[1] and in our everyday conversation, one's doxastic state is expressed in a qualitative, binary way: either one *accepts* a proposition or does not accept it. The qualitative viewpoint sometimes employs alternative phrases such as: "*believing* a proposition" or "*taking* a proposition as true". But how should the two models relate to one another? The present dissertation aims to provide a novel answer to that question.

2 Early Discussions

Extensive discussion about the relationship between probabilities and propositional beliefs was occasioned by Henry Kyburg's lottery paradox (1961). In this section, I focus on the views of three major contributors to that discussion during the 1960's: Henry Kyburg, Isaac Levi, and Richard Jeffrey.

2.1 Lottery Paradox

Here is a natural idea that connects acceptance and probability:

> (Sufficiency Principle) One may accept a proposition if its probability is greater than some fixed threshold t.

[1] See Thomason (2009) for an overview of logic-based artificial intelligence.

But that is inconsistent with the following two principles:

> (Conjunction Principle) One may accept any finite conjunction of acceptable propositions.
>
> (Consistency Principle) One may not accept a logical contradiction.

The proof of inconsistency proceeds by considering a lottery known to be fair with n tickets such that $(n-1)/n > t$. Then, by the the Sufficiency Principle, the agent is permitted to accept all of the following propositions (because each of them has probability greater than t):

> "Ticket 1 will not win."
>
> "Ticket 2 will not win."
>
> $\vdots$
>
> "Ticket n will not win."
>
> "One of the tickets among $1, 2, \ldots, n$ will win."

By n applications of the Conjunction Principle, one may accept the conjunction of all propositions listed in the above. But that conjunction is a logical contradiction. So the Consistency Principle is violated.

Each of the authors under discussion has his own response to the lottery paradox, informed by deeper opinions about the nature of acceptance itself. To these views we now turn.

2.2 Kyburg: Justified Acceptance and High Probability

For Kyburg, there are two senses in which one is justified to accept a proposition. In the first sense, one is justified to accept a proposition simply because there is no "serious question" about that proposition. Kyburg says little about which questions count as serious, except to comment that "frivolous questions don't count [as serious]" (1994: 10) and "[instead of 'frivolous'] some people would say 'philosophical'." (1994: n.16) The propositions that an agent is justified to accept in the first sense at a time t should be jointly consistent and, hence, are permitted to be closed under conjunction. Call justified acceptance in that sense *acceptance as evidence*.

The propositions accepted as evidence determine what Kyburg calls *evidential probabilities*. The basic idea is that, for a fixed agent at a given time t, every proposition A is supported to a certain degree p by the set of all propositions that the agent is justified to accept as evidence. In that case, p is said to be the evidential probability of A for the agent at t. Degree p of evidential support can be a real number in the unit interval $[0, 1]$ but, in general, Kyburg allows the value of p to be an interval $[a, b]$ included in $[0, 1]$. Some story for assigning evidential probability intervals to propositions in light of an arbitrary body of propositions accepted as evidence is required at this stage. Much of Kyburg's work focuses on that problem, but his views on the nature of acceptance can be outlined in abstraction from those details.[2]

Evidential probabilities determine justified acceptance in the second, *probabilistic* sense. An agent is justified to accept a proposition A in the second sense just in case the evidential probability of A (for the agent at the moment) is bounded from below by a threshold t. Threshold t is fixed by the agent's context at the time and is in general less than 1. Kyburg only gives examples to explain how t is to be fixed in light of the current context.[3] Kyburg is willing to tolerate mutual inconsistency among propositions accepted on probabilistic grounds, with the lottery paradox, discussed in detail above, as the star example.

The broad features of Kyburg's views on acceptance can be summarized as follows:

- There are two kinds of justified acceptance: acceptance as evidence and probabilistic acceptance.
- Propositions accepted as evidence are certain, mutually consistent, and closed under finite conjunction.

[2]According to Kyburg, the evidential probability of a proposition concerns especially how much it is supported by the frequencies or chances that are reported in what one accepts as evidence. For example, suppose that the body of eviednce Γ contains only the following propositions:

> "There are exactly 100 balls in urn U at time t",
>
> "There are more than 80 red balls in urn U at time t",
>
> "The balls in that urn are well mixed at time t",
>
> "Adam draws a ball without peeking at a time right after t".

Then the evidential probability of "the ball Adam draws is red" is $[80\%, 100\%]$. That is one of the simplest examples in which we have a clear idea about how evidential probabilities are to be assigned. It remains an on-going research program to determine evidential probabilities when the body of evidence is too complex. One of the most important problems is the so-called *reference class problem*. See Kyburg and Teng (2001) for discussion, especially chapters 9 and 10.

[3]Following Kyburg (1994), consider a context in which an agent is to be offered bets and the highest odds she can be offered to bet against a proposition are 99 to 1. (That may be because, for example, the agent cares only about winning and losing money, she have exactly 99 units of money in the pocket to lose, and she must win at least 1 unit of money if she wins at all). Kyburg assumes that, in that context, the agent would accept any bet against a proposition A if she accepts A. So Kyburg proposes that the probability threshold t for acceptance in that context be at least 99% to ensure that the agent maximizes expected utility.

- Propositions accepted probabilistically are uncertain, may be mutually inconsistent, and may fail to be closed under finite conjunction.

- Probabilistic acceptance depends on evidential probabilities which depend, in turn, on propositions accepted as evidence.

- Evidential acceptance is independent of probabilistic acceptance.

2.3 Levi: Acceptance and Relief from Doubt

Unlike Kyburg, Levi recognizes just one fundamental sense of acceptance: acceptance of a proposition is relief from all doubt about the truth of that proposition. So an agent's acceptance of a proposition A implies that upon acceptance the agent adopts subjective probability one for A. The converse does not hold because in a continuous probability measure every possibility has probability zero, so eliminating every possibility of probability zero would result in a contradiction.[4] Levi also requires that the set of accepted propositions be closed under finite conjunction.

Accepted propositions comprise just one part of one's doxastic state. According to Levi, one's doxastic state consists of two components, i.e. the propositions accepted and subjective probabilities (or some natural generalization thereof, such as sets of probabilities). The two components are irreducible to one another and they interact in at least two ways. (i) What an agent accepts at a time t imposes a constraint on the set of subjective probability functions that are permissible for the agent at t, which is called a *confirmational commitment*. If K is a set of accepted propositions, then the confirmational commitment C selects a probabilistic credal state $C(K)$ that makes each proposition in K certain. (ii) Levi conceives of acceptance as a voluntary act that, like any act, is subject to the Bayesian ideal of expected utility maximization (or some natural generalization thereof that represents the agent's credence in terms of sets of probabilities). But that decision may be undertaken either strategically (one decides to adopt a mechanical procedure to accept all inputs from some channel such as sense perception) or extensively (one decides to accept particular propositions, in context, one at a time). He refers to acceptance directed by a procedure as *routine* expansion (1980: 35-36). Routine expansion serves roughly the same function as acceptance as evidence in Kyburg's

[4]Levi (1980: 2) illustrates the point with the following example. Consider an agent who is thinking about the length L of a rod in meters. Suppose that she accepts that L is in the open interval $[1, 3]$ and has a uniform subjective probability distribution over the interval $[1, 3]$. Although the agent has probability one for the proposition that L is in $[1, 3] \setminus \{2\}$, she does not accept that proposition. In other words, the agent is not relieved from any doubt about the truth of the proposition that L is in $[1, 3] \setminus \{2\}$, because for the agent it might be false in the serious possibility "$L = 2$".

system, but after acceptance Levi does not distinguish routinely accepted propositions from deliberately accepted propositions. When it comes to deliberate acceptance of a particular proposition, Levi has his favored utilities, called epistemic utilities.[5] Regarding consistency, Levi has no blanket proscription against accepting inconsistent beliefs—in fact, acceptance of a contradiction is always an available act and it may occur at any time due to routine expansion. As a pragmatist, Levi's objection to inconsistent beliefs is not that they are fallible (accepted propositions are the standard for serious possibility, after all) but that they fail to direct action (since conditioning on a contradiction results in indifference over all possible

[5]For Levi, decision concerning which propositions to accept pursues the goals of an inquiry. An inquiry is motivated by a which-question: which of the hypotheses in $\mathcal{Q} = \{H_i : i \in I\}$ is true? (The hypotheses in $\mathcal{Q}$ are assumed to be mutually exclusive and jointly exhaustive.) For example, the agent may be interested in the following question: "Which of the following hypotheses are true: Shrödinger's non-relativistic equation (H_1), or Dirac's equation (H_2), or something else (H_3)?" Let $\mathcal{Q}$ be finite. Call the elements of $\mathcal{Q}$ *complete answers* to the question of interest. A disjunction of some but not all hypotheses in the set is called an *incomplete* answer. The goal of an inquiry is to accept true and informative answers to the question of interest. So, for each disjunction A of complete answers in $\mathcal{Q}$, the expected utility of taking A as the strongest answer to accept can be expressed as follows:

$$EU_p(\textsf{accepting } A \textsf{ as strongest}) \quad = \quad \sum_{i \in I} p(H_i) \cdot U(\textsf{accepting } A \mid H_i), \tag{1}$$

where $U(\textsf{accepting } A \mid H_i)$ is the utility of accepting A given that H_i is true. (It is assumed that the utilities depend solely on which answer A is accepted and on which complete answer H_i is true.) For Levi, the decision is to be made only with the goal of the inquiry in mind: accepting true and informative answers to the question of interest. So Levi proposes that the utility function $U(\textsf{accepting } A \mid H_i)$ take the following form:

$$U(\textsf{accepting } A \textsf{ as strongest} \mid H_i) \quad = \quad \alpha \cdot TruthValue(A \mid H_i) + (1-\alpha) \cdot Cont(A). \tag{2}$$

$TruthValue(A \mid H_i)$ equals 1 if A is true given H_i, 0 otherwise. $Cont(A)$ measures the amount of content that the agent takes A to possess. we can imagine that $Cont(A)$ is determined by "how many" possibilities are ruled out by A, i.e., "how many" possibilities are compatible with the negation of A, which is measured by $m(\neg A)$, where m is a normalized, finitely additive measure over the answers to the question of interest. (Normalization is the condition that $m(\top) = 1$. Finite additivity is the condition that $m(A \vee B) = m(A) + m(B)$ for each pair of logically incompatible propositions A and B.) Parameter α is a real number in the open interval $(0, 1)$ and the ratio of α to $1 - \alpha$ represents how much the agent values truth relative to how much she values informativeness. Levi (1967) solves for the propositions to be accepted for maximizing expected utility: acceptance of answer A as strongest maximizes expected utility if and only if A is the disjunction of certain complete answers such that:

- every complete answer H_i with $p(H_i) > \left(\frac{1-\alpha}{\alpha}\right) m(H_i)$ is included as a disjunct;
- no complete answer H_i with $p(H_i) < \left(\frac{1-\alpha}{\alpha}\right) m(H_i)$ is included as a disjunct.

It does not matter whether a complete answer H_i with $p(H_i) = \left(\frac{1-\alpha}{\alpha}\right) m(H_i)$ is included as a disjunct or not.

actions). So one has a pragmatic motive to back out of the contradictions one enters into by routine expansion. For that reason, Levi has been a seminal figure in the development of propositional belief revision theory.[6] The lottery paradox does not involve routine expansion and its solution is tied to the particulars of epistemic utility maximization: i.e., deliberative acceptance of the lottery contradiction does not maximize expected utility if one's value for truth is sufficiently high compared to her value for informativeness.[7]

To summarize, Levi's account has the following features:

- There is just one fundamental notion of acceptance.
- Accepted propositions should be consistent and closed under finite conjunction.
- Credal probabilities depend on accepted propositions via rules called confirmational commitments.
- Acceptance depends on credal probabilities via expected utility maximization, since acceptance is always a voluntary act subject to the discipline of Bayesian rationality.
- The decision to accept may be undertaken either strategically (routine acceptance) or locally (maximization of expected epistemic utility).
- Routine acceptance corresponds roughly to Kyburg's evidential acceptance and the latter to Kyburg's probabilistic acceptance.

2.4 Jeffrey: the Bayesian, Pragmatist Challenge

Richard Jeffrey questions the very idea of accepting propositions in light of probabilities. He argues, from what may be called a pragmatist point of view, that the notion of acceptance is too "unclear" to be employed in rigorous areas of research such as philosophy and sciences:

> The notions of belief [i.e., acceptance] and disbelief are familiar enough but, I find, unclear. In contrast, I find the notion of subjective probability, for all its (decreasing) unfamiliarity, to be a model of clarity—a

[6]Levi's 1977 paper and 1980 book are among the earliest major contributions to belief revision theory.

[7]That is, according to Levi's solution to maximization of expected epistemic utility presented in footnote 6, deliberative acceptance of the lottery contradiction would not maximize expected utility if the ratio of α to $1-\alpha$ is suffiicently high in the sense that there is at least one complete answer H_i such that $p(H_i) > \left(\frac{1-\alpha}{\alpha}\right) m(H_i)$.

> clarity that it derives from its association with the concepts of utility and preference within the framework of Bayesian decision theory. (Jeffrey 1970: 183)

For that reason, he recommends that doxastic states be represented solely by subjective probabilities without recourse to accepted propositions.

Jeffrey in effect poses a *Bayesian challenge* to advocates of theories of acceptance:

> (The Bayesian Challenge) Clarify the notion of acceptance by explaining the role it plays in guiding (rational) action.

But that challenge raises a question: Why must one clarify the notion of acceptance *that way*? Jeffrey does not say much about that, but possible answers are not difficult to find. One way to justify the challenge is to evoke a pragmatist theory of meaning. Assuming that "there is no distinction of meaning so fine as to consist in anything but a possible difference of practice" (Peirce 1878), the notion of acceptance should be clarified by the difference it makes in the practice of those who are claimed to accept propositions—in fact, *any* notion should be clarified by the difference it makes in practice. More plausibly, one of the most important functions of doxastic states is for guiding action, and since acceptance as a propositional attitude reflects aspects of one's doxastic state, a theory of acceptance is sufficiently complete only if it contains a sub-theory about how accepted propositions guide action. The Bayesian challenge can be construed as a request for sketching what such a sub-theory would look like. It is not so much a criterion for abandoning a concept if it is too "unclear" as a methodological principle for questioning a research project if one of its explanatory components appears unpromising.

For example, that Levi's account involves elements directly relevant to the Bayesian challenge. For Levi, acceptance of a proposition requires that one modify one's credal probabilities to make all accepted propositions certain, which induces change in probabilities, which in turn makes a difference in Bayesian decision-making or any decision-making based on subjective probabilities.

3 Further Discussions

This section presents some more recent discussions concerning acceptance both in Bayesian and in traditional Epistemology.

3.1 Developments in Bayesian Epistemology

Bayesian epistemologists tend to agree with Levi's conception of acceptance as a rational act, but differ concerning the underlying aims of acceptance.

According to Kaplan (1981), to adopt A as the strongest proposition to accept is to take A as the strongest proposition to defend if one's only aim were to defend informative truth. Since one may defend a proposition without being fully certain that it is true, Kaplan rejects Levi's requirement that the revised credal probabilities make all accepted propositions certain. For Kaplan, subjective probabilities are more fundamental than accepted propositions and are unaffected by acts of acceptance.

For another example, Maher (1993) proposes that to accept proposition A is to be in a mental state that can be expressed by the sincere assertion of A. Since one can sincerely assert a proposition without being fully certain that it is true, Maher also rejects Levi's requirement that acceptance of a proposition results in full certainty of proposition. Like Kaplan but unlike Levi, Maher holds that acceptance has no impact on subjective probabilities.

Neither Kaplan nor Maher has much to say about the Bayesian challenge to account for the role of acceptance in decision making. They maintain that the state of accepting a proposition impinges only upon certain behaviors: defense of propositions for Kaplan, and sincere assertability for Maher. A determined response to the Bayesian challenge would still require, for example, a detailed decision theoretic analysis of what is at stake in defending a hypothesis or in sincere assertion freed from extraneous, pragmatic interests. To summarize their views:

- Following Levi, subjective probabilities influence which propositions are rational for an agent to accept, according to maximization of expected utility. But, pace Levi, acceptance of propositions has no impact whatever on subjective probabilities.

- Accepted propositions are permitted to be closed under conjunction.

- One may accept a proposition without being fully certain that it is true.

3.2 Developments in Traditional Epistemology: Defense of Joint Inconsistency

Kyburg's view remains influential in traditional epistemology. That is not surprising: Kyburg, like many traditional epistemologists, is concerned with theories of justified beliefs.

While Kyburg does not try to reduce justified beliefs to probabilities, Foley (1992, 2009) does. He maintains that one is justified to believe a proposition if and only if the agent is justified to have a sufficiently high degree of belief in that proposition given one's evidence. Although the biconditional sounds like a reduction of traditional epistemology to Bayesian epistemology, Foley points out that it actually gives traditional epistemology an indispensable role to play. The degrees of belief that one actually has seldom match exactly the degrees that one is justified to have. But it is a lot easier to have justified belief that A: it suffices that one does believe proposition A and, at the same time, the degree p of belief in A that one is justified to have is sufficiently high, no matter whether p is exactly one's actual degree of belief in A. Foley remarks that possession of justified propositional beliefs is usually what we require of an epistemic agent, and that secures an indispensable role for traditional epistemology.

Although Foley is less interested than Kyburg in working out in detail a theory of justified degrees of belief, he puts more effort into explaining how the threshold view is compatible with our intuitions. One of the most unintuitive results of the threshold view is that sometimes what one believes should not be closed under deduction, but that seems to clash with the intuition that deduction is one of the most important methods for extending what one is justified to believe. Foley explains that intuition away by noting that, when one provides a deductive argument from premises to conclusion, the propositional attitudes involved are not believing, but presuming, positing, assuming, supposing, or hypothesizing. Furthermore, the threshold view allows one to be justified in believing each element of a set of propositions that are jointly inconsistent. But that seems to contradict our intuition and make *reductio* arguments pointless. In reply, Foley notes that joint consistency is actually not a global requirement on the set of all beliefs that one has, but only a local requirement. For example, if one intends to use a specific set Γ of justified beliefs as evidence or as premises to argue for further propositions, then Γ is indeed required to be jointly consistent. But Γ typically contains only the propositions that have some privileged epistemic status among all of one's beliefs—privileged enough to make the elements of Γ appropriate as evidence for other propositions. Indeed, since both Kyburg and Foley employ probabilities *conditional on* the set of propositions that are accepted as evidence, that set has to be consistent in order to ensure that conditional probabilities exist. To summarize:

- Justified subjective probabilities influence which propositions are rational to accept. High justified probability is a necessary and sufficient condition for justified acceptance. But acceptance of propositions has no impact on subjective probabilities at all.

- Accepted propositions need not be jointly consistent and, in most cases, should not be closed under conjunction.

- One may accept a proposition without being fully certain that it is true.

3.3 Developments in Traditional Epistemology: the Unimportance of Subjective Probabilities

Some traditional epistemologists argue, pace Kyburg and Foley, that probability does not have a primary role in epistemic justification. Pollock (1983, 1995), for example, argues that what plays an primary role in justification is, instead, what counts as a *prima facie* reason that one can provide for a proposition, and how such a reason might be *defeated*. For him, probabilities may be cited in the reasons and, hence, only play a secondary role in justification. For example, Pollock formulates the following probabilistic rule for acceptance:

> Pollock's Acceptance Rule: Suppose that B is a *projectible* predicate with respect to A in the sense that "all A's that have been observed are B's and many A's have been observed" is a prima facie reason for believing "all A's are B's". Then, "$A(c)$ and $p(Bx/Ax)$ is sufficiently high" (understood as about chance or observed frequency) is a prima facie reason for believing $B(c)$, where c is a singular term.

A prima facie reason may be defeated and loses its power for justifying other claims, and the lottery paradox provides such an example. According to Pollock's acceptance rule, for each ticket i in the lottery, we do have a prima facie reason to believe "ticket i will lose" (L_i), and the reason is "ticket i is in the lottery and the chance for a ticket to lose given that it is in the lottery is extremely high" (R_i). But that probability claim is defeated as a reason for believing "ticket i will lose" because of what we may call *collective defeat*:

(a) There is a set $\{R_i : 1 \leq i \leq n\}$ of propositions that we are justified to believe.

(b) Each element R_i is a prima facie reason for believing L_i.

(c) So, by (b), the propositions that elements of $\{R_i : 1 \leq i \leq n\}$ are prima facie reasons for are jointly inconsistent.

(d) Furthermore, $\{R_i : 1 \leq i \leq n\}$ is a minimal set that leads to (c). To be specific, for each k such that $1 \leq k \leq n$, the propositions that elements of $\{R_i : 1 \leq i \leq n\} \setminus \{R_k\}$ are prima facie reasons for are *not* jointly inconsistent.

(e) It is (c) and (d) that constitute a collective defeat for R_i as a reason for L_i.[8]

[8] Condition (c) alone does not suffice for defeat, because inclusion of (d) is necessary for preventing a triviality result. Without (d), we can defeat any proposition R^* as a prima facie reason as follows: consider the set $\{R_i : 1 \leq i \leq n\} \cup \{R^*\}$, and the propositions that elements of that set are prima facie reasons for include $L_1, \ldots, L_n$, which are jointly inconsistent.

According to Pollock, there are many ways to defeat a reason and collective defeat is the sort of defeat that solves the lottery paradox.

To summarize the salient features of Pollock's account:

- There is only a very limited way in which objective probabilities such as observed frequencies and chances influence what count as the propositions that one is justified to accept. Subjective probabilities are irrelevant.
- Accepted propositions should be jointly consistent and are permitted to be closed under conjunction.
- One may accept a proposition without being fully certain that it is true.

3.4 Developments in Traditional Epistemology: Response to the Bayesian Challenge

Traditional epistemology does not lack resources that can be employed or developed in response to the Bayesian challenge: the kind of practical reasoning that we perform on a daily basis is an evident candidate. The role of propositional beliefs in practical reasoning is quite obvious, at least in folk talk: by checking what time it is now, I come to accept the proposition that the grocery store is still open, and employ that as one of the *premises* to argue to myself that I should go grocery shopping today. "The grocery store is still open. It will be closed in two hours. So I should go now." Bratman (1987) and Cohen (1989) develop that idea, although their primary interests are in philosophy of action rather than in the Bayesian challenge. Furthermore, planning theorists in artificial intelligence have taken steps to formalize how propositional beliefs may serve as premises for practical reasoning (the literature is huge; see Thomason (2009) for a review).

It is interesting to see how far the approach can be developed, and I will propose some such steps in section 9.

3.5 Developments in Nonmonotonic Logic and Belief Revision Theory

As a result of Kyburg's influence, some logical systems have been developed to allow that a set of premises may have conclusions that are not closed under conjunction (Hawthorne (1996) and Hawthorne and Bovens (1999)). The idea is that, relative to a probability function p defined over the language L in question, and relative to a probability threshold t, ϕ is a defeasible (i.e., nonmonotonic) consequence of a finite set Γ of premises, written $\bigwedge \Gamma \mathrel{|\!\sim}_{p,t} \phi$ just in case the standard conditional probability $p(\phi \mid \bigwedge \Gamma)$ is no less than t or (in the limiting case) $p(\bigwedge \Gamma) = 0$. It is possible that we have $\phi \mid\!\sim_{p,t} \psi$, $\phi \mid\!\sim_{p,t} \psi'$ but $\phi \not\mid\!\sim_{p,t} \psi \wedge \psi'$, as expected. All nonmonotonic consequence relations $\mid\!\sim$ thus defined have the following properties:

(**Reflexivity**) $\phi \mid\!\sim \phi$

(**Right Weakening**) Whenever $\phi \mid\!\sim \psi$ and ψ entails ψ' according to classical logic, then $\phi \mid\!\sim \psi'$.

(**Left Equivalence**) Whenever ϕ and ϕ' are equivalent according to classical logic and $\phi \mid\!\sim \psi$, then $\phi' \mid\!\sim \psi$.

(**Very Cautious Monotonicity**) Whenever $\phi \mid\!\sim \psi \wedge \gamma$, then $\phi \wedge \gamma \mid\!\sim \psi$.

(**Exclusive Or**) Whenever $\phi \mid\!\sim \psi$, $\phi' \mid\!\sim \psi$, and ϕ and ϕ' are incompatible according to classical logic, then $\phi \vee \phi' \mid\!\sim \psi$.

(**Weak And**) Whenever $\phi \mid\!\sim \psi$, $\phi \wedge \neg\psi' \mid\!\sim \psi'$, then $\phi \mid\!\sim \psi \wedge \psi'$.

A complete representation conjecture is proposed in Hawthorne (1996): every nonmonotonic consequence relation $\mid\!\sim$ that satisfies the above system, called O, is identical to $\mid\!\sim_{p,t}$ for some p and t. But that conjecture is false; see Hawthorne and Makinson (2007) and Paris and Simmonds (2009) and Makinson (2009) for discussion. It remains an open problem whether system O can be strengthened to obtain a complete representation in terms of $\mid\!\sim_{p,t}$.

An equally important open question concerns how a theory of practical reasoning can be based on a nonmonotonic logic that violates conjunctive closure of plausible conclusions. Formal studies of nonmonotonic logic were historically initiated in artificial intelligence through the need to develop a system of reasoning that helps an agent devise a plan regarding which sequence of actions to take to achieve a goal (see Thomason (2009) for a review). Perhaps conjunctive closure should not be applied to the totality of one's propositional beliefs, but it should be applicable to the propositional beliefs that are at least relevant to a planning task.

4 Proposed Approach

In this section, I propose a new research program for investigating the relationship between propositional beliefs and subjective probabilities. I will also explain how it differs from the approaches reviewed above, and its relative advantages. I also sketch how the present dissertation constitutes a crucial step in the proposed program.

4.1 A New Research program

Consider an agent who strives for the ideal postulated by the standard Bayesian theory. Namely, if the agent were given enough time for deliberation, she would have betting behaviors or preferences that define, or at least reflect, her underlying subjective probabilities (in a way that maximizes expected utility). Subjective probabilities take sharp values, and they are revised by Bayes conditioning on new information. In contrast, there is a more traditional view according to which the agent possesses, plans in accordance with, and revises uncertain *propositional* beliefs. It goes without saying that qualitative reasoning with propositional beliefs can never duplicate the fill refinement of Bayesian rationality. But perhaps propositional beliefs can nonetheless crudely but *aptly* reflect probabilistic beliefs in a way that approximates, tracks, or otherwise aligns with full probabilistic reasoning. To be specific, we want to see whether the Bayesian picture is compatible with the following claims:

(I) Propositional beliefs can serve as apt, albeit rough, representations of one's underlying credal probabilities.

(II) When the agent receives new information, she can reason directly with her propositional beliefs in order to revise them, without reasoning about probabilities, in such a way that the revised propositional beliefs can still aptly represent her updated credal probabilities (which are assumed to be obtained by Bayes conditioning).

(III) In routine decision-making or planning in everyday life, the agent can reason directly with her propositional beliefs to reach conclusions for action, without reasoning about probabilities and utilities. And it is possible that propositional practical reasoning can work in a way such that, if it yields a conclusion for action, it is consistent with maximization of expected utility.

Possible advantages of propositional reasoning over probabilistic reasoning include computational efficiency, cognitive plausibility, and communicability, but the question of compatibility between the two approaches precedes those considerations and will be the sole focus of the dissertation.

The approach I propose differs from the approaches reviewed above in a number of ways. Unlike the approach of Levi, acceptance of a proposition does not require being fully certain that the proposition is true, and does not require a momentous change in subjective probabilities. Unlike the approaches of Levi, Kaplan, and Maher, rational acceptance need not result from decision. The argument for an apt correspondence between probabilistic and propositional reasoning is compatible with acceptance as a brute, involuntary disposition. And even if we sometimes wish to decide which propositions to believe by maximization of expected utility, we need to figure out what *can* be done with propositional beliefs and how *well* the tasks can be done before we can know how to assign relevant utilities. I propose that the tasks described in (II) and (III) are among the most valuable things that we can do with propositional beliefs.

Second, unlike the approaches of Kyburg, Foley and Hawthorne, the program can be pursued with the requirement that one's propositional beliefs be jointly consistent and closed under conjunction—at least for the propositions that one cares about in a given context for answering a questions, making a decision, etc. In fact, joint consistency and conjunctive closure are almost compelled in the program, unless we are prepared to revolutionize decades of research in qualitative belief revision theory (regarding (II)) and in qualitative planning theory (regarding (III)). Qualitative belief revision theory is built upon the idea that, when the new information and one's old propositional beliefs are jointly inconsistent, one must revise the old beliefs rather than tolerate the inconsistency. Practical reasoning, as studied in classical planning theory, is built upon extensions of classical logic, which permits conjunctive closure.

Third, unlike the approach of Jeffrey, the program takes seriously that propositional beliefs have important roles to play in epistemology. And unlike the approach of Pollock, it takes seriously that subjective probabilities have important roles to play in epistemology. Their mutual intention to demolish one of the two sides appears to be motivated primarily by the ongoing failure to build a bridge between the two. The present dissertation is a step toward building that bridge.

4.2 Sketch of the Dissertation

Suppose that you have accepted a propositional belief state that aptly represents your underlying probabilistic credal state. Suppose that new information is provided. You now have an option how to revise your propositional belief state. According to the *epiphenomenal* plan, you revert back to your underlying credal state, incorporate the new information by Bayesian conditioning, and then accept a new propositional belief state that aptly represents the conditional probabilities. Your original, propositional belief state plays no role in the revision. Alternatively, one could directly revise one's propositional belief state in light of the new information without reverting back to one's probabilistic credal state. Since the propositional belief state is at best crudely but aptly represents the more detailed probabilistic credal state, one would expect repeated propositional revisions to eventually veer away from the properly Bayesian, epiphenomenal path. If that does not happen—if propositional revision forever maintains an apt relationship to the successive credal states that *would have* arisen on the epiphenomenal path, then say that propositional reasoning *tracks* probabilistic reasoning. The tracking condition is a joint constraint on the apt representation relation, propositional belief revision, and Bayesian conditioning. It defines what it means for propositional belief revision to aptly represent Bayesian conditioning. In the present dissertation, I will show the following:

- The tracking condition is unattainable, on pain of triviality, if one is restricted to the standard, *AGM* belief revision methods (Alchourrón, Gärdenfors, and Makinson 1985).
- And yet there is a strictly broader class of propositional belief revision methods that can satisfy the tracking condition and, *at the same time*, solve the paradoxes of uncertain acceptance, old and new, presented above.
- In order to obtain the preceding advantages, one must abandon the following, standard belief revision principle: "Do not give up beliefs if the new information is compatible with what you already believe."
- Rejection of that principle is also necessary for modeling Lehrer's (1965) no-false-lemma variant of Gettier's (1963) celebrated counterexample to justified true belief as an analysis of knowledge.

Those points will be made in two passes. For the sake of pictorial presentation, I begin by assuming that acceptance is a rule that picks out a *unique* propositional belief state for a given credal state (from section 5 to section 8.4). Then the assumption is relaxed: acceptance is generalized to a *relation* of apt representation between propositional belief states and credal states (section 8.5). Finally, I will explain how we may pursue the other parts of the program, especially those regarding the use of propositional beliefs in everyday practical reasoning (section 9).

5 Two New Paradoxes of Uncertain Acceptance Geometrized

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 this new riddle is *geometrical* rather than logical (figure 1). Let H_1 be the proposition that ticket 1 wins, and similarly for H_2,

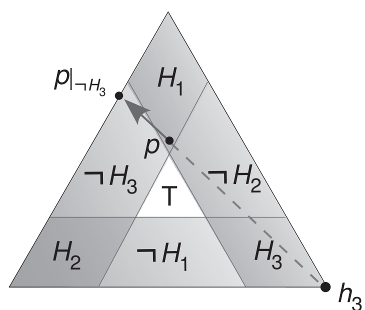

Figure 1: the first riddle

H_3. The space of all probability distributions over the three tickets consists of a triangle in the Euclidean plane whose corners have coordinates $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 1)$, which are the extremal distributions that concentrate all probability on a single ticket. The assumed distribution p over tickets then corresponds to the point $p = (1/2, 1/4, 1/4)$ in the triangle. The conditional distribution $p|_{\neg H_3} = p(\cdot|\neg H_3)$ is the point $(2/3, 1/3, 0)$, which lies on a ray through p that originates from corner 3, holding the odds $H_1 : H_2$ constant. Each zone in the triangle is labeled with the strongest proposition accepted at the probability measures inside. The acceptance zone for H_1 is a parallel-sided diamond that results from the intersection of the above-threshold zones for $\neg H_2$ and $\neg H_3$, since it is assumed that the accepted propositions are closed under conjunction. The rule leaves the inner triangle as the acceptance zone for the tautology $\top$. The riddle can now be seen to result from the simple, geometrical fact that p lies near the 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 $\mathsf{B}_p \models H$. Then the principle may be stated succinctly as follows, where $p|_E$ denotes the conditional distribution $p(\cdot|E)$:

$$\mathsf{B}_p \models H \text{ and } H \models E \implies \mathsf{B}_{p|_E} \models H. \tag{3}$$

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

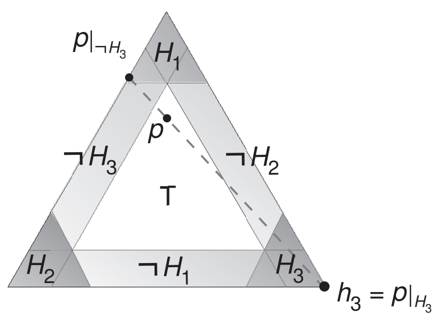

Figure 2: second riddle

sense, complementary riddle emerges. For suppose that the credal state is p, just inside the zone for accepting that either ticket 1 or 2 will win and close to, but outside of the zone for accepting that ticket 1 will win. 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:[9]

$$\mathsf{B}_{p|_E} \models H \text{ and } \mathsf{B}_{p|_{\neg E}} \models H \implies \mathsf{B}_p \models H, \tag{4}$$

which we call *Case Reasoning*.

The two new riddles add up to one big riddle: there is, in fact, *no* ad hoc manipulation of distinct thresholds for distinct propositions that avoids both riddles.[10] 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*.

6 Condition for Solving the Paradoxes

This section proposes a condition for solving the two new paradoxes presented in the preceding section.

6.1 The Propositional Space of Reasons

Part of what is jarring about the paradoxes 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 paradox, observed logical consequences H result in rejection of H. In the second paradox, 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 paradoxes witness that one has not really entered into a purely propositional "space of reasons" (Sellars 1956).

[9]The principle is analogous in spirit to the *reflection principle* (van Fraassen 1984), which, in this context, might be expressed by saying that if you know that you will accept a proposition regardless what you learn, you should accept it already. Also, a non-conglomerable probability measure has the feature that some B is less probable than it is conditional on each H_i. Schervish, Seidenfeld, and Kadane (1984) show that every finitely additive measure is non-conglomerable in some partition. In that case, any sensible acceptance rule would fail to satisfy reasoning by cases. Some experts advocate finitely additive probabilities and others view non-conglomerability as a paradoxical feature. For us, acceptance is relative to a partition (question), a topic we discuss in detail in Lin and Kelly (2011), so non-conglomerability does not necessarily arise in the given partition.

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

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 paradoxes 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, I exhibit a collection of rules that do achieve perfect harmony with Bayesian conditioning.

6.2 Questions, Answers, and Credal States

Let $\mathcal{Q} = \{H_i : i \in I\}$ be a countable collection of mutually exclusive and exhaustive propositions representing a *question* to which $H_1, \ldots, H_i, \ldots$ are the *(complete) answers*. Let $\mathcal{A}$ 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.

6.3 Belief Revision

A (propositional) *belief state* is just a deductively closed set of propositions; but for the sake of convenience I identity each belief state with the conjunction of all propositions believed. A *belief revision method* is a mapping $\mathsf{B} : \mathcal{A} \to \mathcal{A}$, understood as specifying the initial belief state $\mathsf{B}(\top)$, which would evolve into new belief state $\mathsf{B}(E)$ upon revision on information E.[11] *Hypothetico-deductive Monotonicity*, for example, can now be stated in terms of belief revision, rather than in terms of

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

Bayesian conditioning:[12]

$$\mathsf{B}(\top) \models H \text{ and } H \models E \implies \mathsf{B}(E) \models H. \tag{5}$$

Case Reasoning has a similar statement:[13]

$$\mathsf{B}(E) \models H \text{ and } \mathsf{B}(\neg E) \models H \implies \mathsf{B}(\top) \models H. \tag{6}$$

6.4 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 B that assigns to each credal state p a belief revision method B_p. Then $\mathsf{B}_p(\top)$ is the belief state *accepted* unconditionally at credal state p, and proposition H is *accepted* (unconditionally) by rule B at credal state p if and only if $\mathsf{B}_p(\top) \models X$.[14]

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

$$\mathsf{B}_p(E) \quad = \quad \mathsf{B}_{p|_E}(\top), \tag{9}$$

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

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

[14] The following, conditional acceptance *Ramsey tests* translate the present framework into notation familiar in the logic of epistemic conditionals:

$$p \Vdash E \Rightarrow H \iff \mathsf{B}_p(E) \models H; \tag{7}$$
$$E \mathrel{|\!\sim}_p H \iff \mathsf{B}_p(E) \models H. \tag{8}$$

I am indebted to Hannes Leitgeb (2010) for the idea of framing the 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.

$$\begin{tikzcd}[column sep=huge, row sep=large]
p \arrow[r, "\text{Bayesian conditioning}"] \arrow[d, "\text{acceptance}"'] & p|_E \arrow[d, "\text{acceptance}"] \\
\mathsf{B}_p(\mathsf{T}) \arrow[r, "\text{belief revision}"'] & \mathsf{B}_p(E) = \mathsf{B}_{p|_E}(\mathsf{T})
\end{tikzcd}$$

Figure 3: Belief revision that tracks Bayesian conditioning

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 Acretive Belief Revision and An Impossibility Result

7.1 Accretive Belief Revision

It is easy to achieve perfect tracking: given the values of $\mathsf{B}_p(\top)$ for all p, just *define* the values of $\mathsf{B}_p(E)$ according to equation (9). 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 into two parts: belief revision method B satisfies *Inclusion* if and only if:[15]

$$\mathsf{B}(\top) \wedge E \models \mathsf{B}(E). \tag{10}$$

Method B satisfies *Preservation* if and only if:

$$\mathsf{B}(\top) \text{ is consistent with } E \implies \mathsf{B}(E) \models \mathsf{B}(\top) \wedge E. \tag{11}$$

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.

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

7.2 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

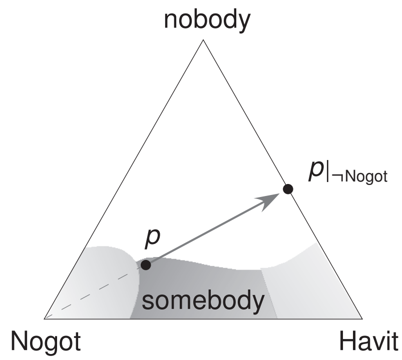

Figure 4: How Preservation may fail plausibly

driver's 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 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 to become doubtful, the disjunctive belief would be retracted. Any theory of rational belief that models this paradigmatic, espitemological situation must violate the Preservation axiom.

The preceding intuitions are vindicated by the following no-go theorem. First, let us define some properties that a sensible acceptance rule should have. To begin with, I exclude skeptical acceptance rules that refuses to accept complete answers to $\mathcal{Q}$ at almost every 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 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:[16]

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

In a similar spirit, I exclude the extremely gullible or opinionated rules that accept complete answers to $\mathcal{Q}$ at almost every credal state. Say that B is *non-opinionated* if and only if there is some non-empty, open subset of $\mathcal{P}$ over which some incomplete, disjunctive answer is accepted. Say that B is *consistent* if and only if the inconsistent proposition $\perp$ is accepted at no credal state. Say that B is *corner-monotone* if and only if acceptance of complete answer H_i at p implies acceptance of H_i at each point on the straight line segment from p to the corner h_i of the simplex at which H_i has probability one.[17] Aside from the intuitive merits of these properties, all proposed acceptance rules I am aware of satisfy them. Rules that satisfy all four properties are said to be *sensible*. Then we can have:

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

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

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

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.[18]

8 Shoham Belief Revision and A Possibility Result

8.1 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 is reproduced in figure 5. If, instead, one is in

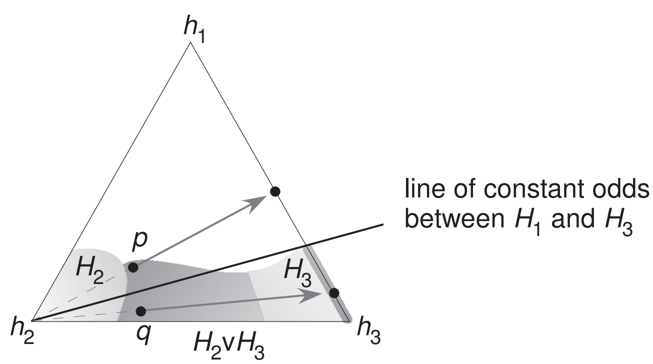

Figure 5: Line of constant odds

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 $\neg H_2$.

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

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

8.2 An Odds-Based Acceptance Rule

I now present an acceptance rule based on odds thresholds that illustrates how to sensibly track Bayesian conditioning (and to solve the two new paradoxes) by violating the counter-intuitive Preservation property. The particular rule discussed in this section motivates the 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).[19] 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 $\prec$.[20] Each plausibility order $\prec$ induces a belief revision method $\mathsf{B}_\prec$ as follows: given information E in $\mathcal{A}$, let $\mathsf{B}_\prec(E)$ be the disjunction of the most plausible answers to $\mathcal{Q}$ with respect to $\prec$ that are logically compatible with E. More precisely: 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 the 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:

$$\mathsf{B}_p \quad = \quad \mathsf{B}_{\prec_p}. \tag{12}$$

Here is a way to construct $\prec_p$ in terms of odds.[21] In particular, let t be a constant greater than 1 and define:

$$H_i \prec_p H_j \quad \iff \quad \frac{p(H_i)}{p(H_j)} > t, \tag{13}$$

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_2, as depicted in figure 6.a. To determine whether $H_1 \prec_p H_2$,

[19] Shoham proposed the approach as a semantics for non-monotonic logic but, in light of Makinson and Gärdenfors' (1991) translation of non-monotonic logic into belief revision notation, it can be viewed as a theory of belief revision strictly weaker than the AGM theory.

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

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

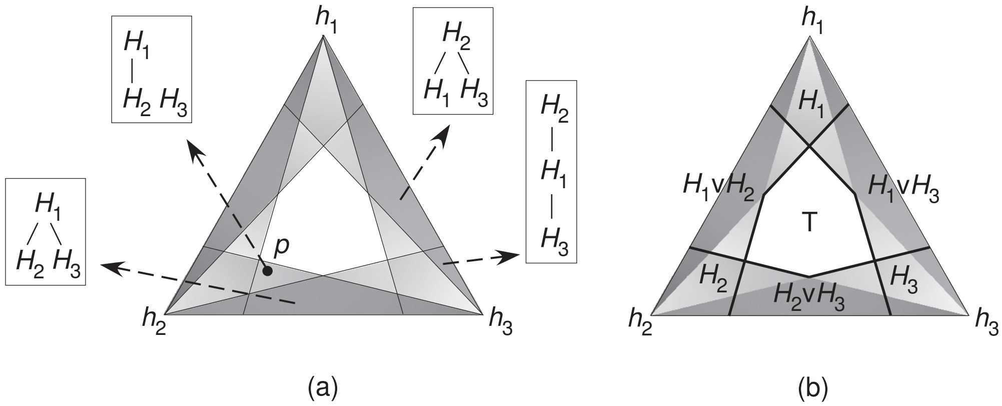

Figure 6: A rule based on odds thresholds

simply check whether 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 $\mathsf{B}_p(\top)$ at p is the disjunction of the most plausible answers in $\prec_p$. So the zone for accepting a belief state is bounded by the constant odds lines, as depicted in figure 6.b.[22] 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 $\mathsf{B}_p(E) = \mathsf{B}_{p|_E}(\top)$, 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). The equality is no accident: the relative plausibility between H_1 and H_3 at both credal states—prior and posterior—is defined by the same odds threshold, and conditioning on $H_1 \vee H_3$ always preserves the odds between H_1 and H_3. So the proposed rule tracks conditioning due to a simple principle of design: define relative plausibility by quantities preserved under conditioning. That principle 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 paradoxes (i.e., it satisfies

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

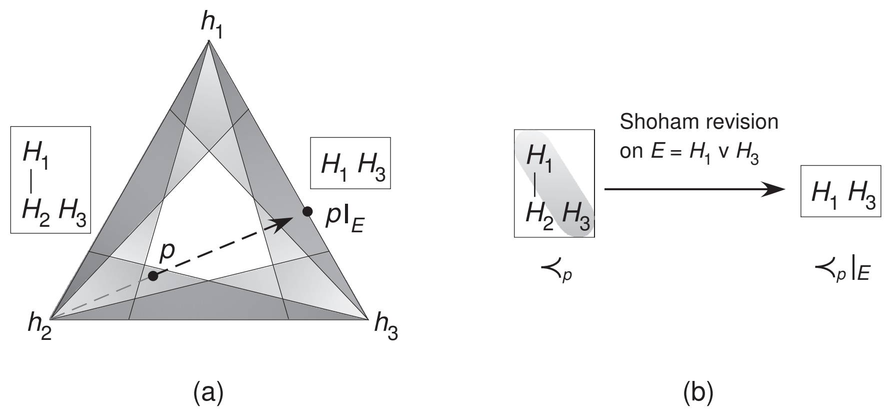

Figure 7: How the rule tracks conditioning

Hypothetico-deductive Monotonicity (3) and Case Reasoning (4)). 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 I did on figures 1 and 2.

The Preservation axiom (11) is violated (figure 8), for reasons similar to those discussed in the preceding section (figure 5). Preservation is violated at p when

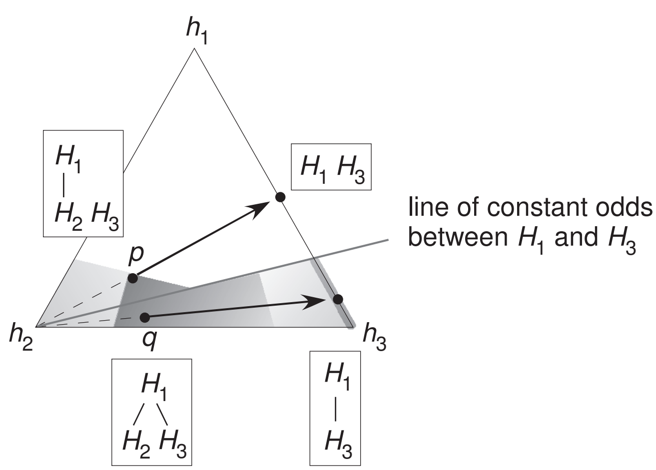

Figure 8: Preservation and odds

$\neg H_2$ 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).

8.3 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-probability answers (i.e., $\prec$ is a mapping $\prec_{(\,\cdot\,)}$ that sends p to $\prec_p$). An acceptance rule B is *Shoham-driven* if and only if it is generated by some assignment $\prec_{(\,\cdot\,)}$ of plausibility orders in the sense of equation (12). 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:

$$H_i \prec_p H_j \quad \Longleftrightarrow \quad p(H_i)/p(H_j) > 3. \tag{14}$$

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

$$\mathsf{B}_p(\top) \;\;=\;\; \bigwedge \left\{ \neg H_i : \frac{p(H_i)}{\max_k p(H_k)} < \frac{1}{3} \right\}, \tag{15}$$

which is a special case of proposition 4 below. Equation (15) 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 paradox of acceptance:

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

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

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

Furthermore:

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

$$\prec_p |_E \quad = \quad \prec_{p|_E}, \tag{16}$$

$$\mathsf{B}_{\prec_p|_E} \quad = \quad \mathsf{B}_{\prec_{p|_E}}. \tag{17}$$

$$\begin{tikzcd}
p \arrow[r, "\text{Bayesian conditioning}"'] \arrow[d] & p|_E \arrow[d] \\
\prec_p \arrow[r, "\text{Shoham revision}"'] \arrow[d] & \prec_p|_E = \prec_{p|_E} \arrow[d] \\
\mathsf{B}_{\prec_p} \arrow[r, "\text{Shoham revision}"'] & \mathsf{B}_{\prec_p|_E} = \mathsf{B}_{\prec_{p|_E}}
\end{tikzcd}$$

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).

8.4 Shoham-Driven Acceptance Based on Odds

This section explains why it is no accident that every Shoham-driven rule we have examined so far is somehow based on odds.

The assignment (14) of plausibility orders and the associated assignment (15) 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:

$$H_i \prec_p H_j \quad \Longleftrightarrow \quad p(H_i)/p(H_j) > t_j. \tag{18}$$

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

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

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:

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

The range of t_i and w_i should be 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 (20) is a plausibility order.*

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

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

for all p in $\mathcal{P}$. When all weights w_i are equal, order (20) and belief state (21) reduce to order (18) and belief state (19). Then we have:

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

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

2. *B tracks conditioning.*

3. *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.[23] So a Shoham-driven rule can easily be sensible and conditioning-tracking (and thus

[23] If $\mathcal{Q}$ 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.

paradox-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 (20) can be absorbed into odds without loss of generality:

$$H_i \prec_p H_j \iff w_i p(H_i)/w_j p(H_j) > t_j, \tag{22}$$
$$\iff p(H_i)/p(H_j) > t_j(w_j/w_i), \tag{23}$$

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 by equation (24); 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:

$$t = (t_{ij} : i, j \in I \text{ and } i \neq j), \tag{25}$$

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

$$\rhd = (\rhd_{ij} : i, j \in I \text{ and } i \neq j), \tag{26}$$

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

$$H_i \prec_p H_j \iff p(H_i)/p(H_j) \rhd_{ij} t_{ij} . \tag{27}$$

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.*

8.5 Apt Representation Relations

I have assumed, for the sake of pictorial presentation, that acceptance is a rule that picks out a unique belief state for a given credal state. This section relaxes that assumption. Let R be a relation between probability measures and belief revision methods. Understand $\mathsf{R}(p, \mathsf{B})$ as saying that probability measure p is *aptly represented* by belief revision method B. Let A be a proposition; write $\mathsf{R}(p, A)$ to mean that $\mathsf{R}(p, \mathsf{B})$ for some belief revision method B such that $\mathsf{B}(\top) = A$. That is understood as saying that p is aptly represented by accepting A as one's belief state. The belief revision method $\mathsf{B}|_E$ that results from B by conditioning on information E is defined by:

$$\mathsf{B}|_E(E') \quad = \quad \mathsf{B}(E' \wedge E). \tag{28}$$

Say that R *tracks conditioning* if and only if, for every probability measure p, new information E such that $p(E) > 0$, and belief revision method B:

(i) if $\mathsf{R}(p, \mathsf{B})$, then $\mathsf{R}(p|_E, \mathsf{B}|_E)$;

(ii) if $\mathsf{R}(p|_E, \mathsf{B})$, then $\mathsf{R}(p, \mathsf{B}')$ for some belief revision method B' such that $\mathsf{B}'|_E = \mathsf{B}$.

The first clause says that, if one starts with a belief revision method B that aptly represents the prior probabilities, then the new belief revision method $\mathsf{B}|_E$ still aptly represents the posterior probabilities. The second clause says that every belief revision method B that aptly represents the posterior probabilities can be obtained from some belief revision method B' that aptly represents the prior probabilities.

The impossibility result presented above generalizes as follows. Say that R is *accretive* just in case we have that $\mathsf{R}(p, \mathsf{B})$ only if B is an accretive belief revision method. The four properties that jointly define sensibility are generalized as follows. Say that R is *non-skeptical* if and only if, for each complete answer H_i, there is an open neighborhood $\mathcal{O}$ of $\mathcal{P}$ such that, for each $p \in \mathcal{O}$, $\mathsf{R}(p, H_i)$. Say that R is *non-opinionated* if and only if there is an open neighborhood $\mathcal{O}$ of $\mathcal{P}$ and two distinct complete answers H_i, H_j such that every belief state that aptly R-represents some credal state in $\mathcal{O}$ is compatible with both H_i and H_j. Say that R is *corner-monotone* if and only if, for each probability measure p and complete answer H_i, if $\mathsf{R}(p, H_i)$, then $\mathsf{R}(q, H_i)$ for every q that lies on the straight line segment from p to the corner h_i of the simplex at which H_i has probability one. Say that R is *consistent* if and only if $\mathsf{R}(p, \perp)$ for no p. Say that R is *sensible* just in case it satisfies all those four properties. Then we have:

Theorem 4 (no-go theorem for accretive apt representation). *Let question $\mathcal{Q}$ have at least three complete answers. Then the following conditions are jointly inconsistent:*

1. R *is sensible;*

2. R *tracks conditioning;*

3. R *is accretive.*

It is not hard to construct a relation R that is sensible and tracks conditioning: just identify it with a sensible Shoham-driven acceptance rule $p \mapsto \mathsf{B}_p$. Namely, define R by: $\mathsf{R}(p, \mathsf{B})$ iff $\mathsf{B} = \mathsf{B}_p$. In general, we can use a set S of Shoham-driven acceptance rules to define an apt representation relation R: $\mathsf{R}(p, \mathsf{B})$ iff B is the belief revision method associated with p according to some acceptance rule in S.

8.6 Accetive Belief Revision Revisited

The preceding impossibility result says that there is no sensible, fixed relation that associates each Bayesian credal state with a qualitative belief revision method that aptly represents it. That result does not rule out the possibility that for each Bayesian credal state there exists a qualitative belief revision method that aptly represents it, if one is not concerned with the sensibility of the global relation between Bayesian credal states and belief revision methods.[24] Dropping the global, sensibility requirement makes some sense in light of a strict, first-person perspective in which the Bayesian individualist moves forward in noble isolation both from her compatriots and from counterfactual considerations concerning what she *might* have believed in similar, alternative situations. But social and counterfactual considerations do arise quite naturally in the context of belief change, and then it is awkward to drop the sensibility requirement as a social epistemic norm.

Recall that the sensibility requirement is defined as joint satisfaction of consistency, non-skepticism, non-opinionation, and corner-monotonicity. Consistency is a basic norm for propositional beliefs, and non-skepticism characterizes the topic under discussion, which is everyday, uncertain belief. Non-opinionation concerns the possibility of a state of judgment suspension that is at least minimally "stable". It should be at least possible that one is permitted to suspend judgment between two exclusive answers and to remain so suspended if one's subjective probabilities had been just slightly different from what they actually are. Similarly, it should remain possible that compatriots whose degrees of belief are almost indistinguishable from one's own opinions would not invariably believe either one of the two answers or the other. Corner-monotonicity can be understood as a very weak principle about counterfactual stability of belief: whenever you are permitted to believe a compete answer H_i to a question, then you would still be so permitted if your credal state *had been* more committed to H_i—in the sense that it lies on the line that connects

[24] I am endebted to Horacio Arlo-Costa and to Hannes Leitgeb for suggesting this alternative, weaker viewpoint on acceptance.

your actual credal state and the credal state that accords full probability to H_i. Corner-monotonicity can also be understood as a very weak but desirable social policy concerning acceptance. If you believe a compete answer H_i to a question, then you would like a person to agree with you if she does not need your persuasion at all—in the sense that her credal state already lies on the line that connects your actual credal state and the credal state that accords full probability to H_i. Otherwise, the debate concerning H_i would descend into an intractable dispute about standards.

Giving up the global, sensibility requirement carries a cost in terms of social and counterfactual implausibility. In contrast, giving up accretive belief revision only requires using partial orders instead of rankings for modeling plausibility—with an intuitive bonus: accommodating the no-false-lemma Gettier phenomenon. That is why I prefer to give up accretive belief revision in favor of the global sensibility requirement.

9 Prolegomea to A Decision Theory for Everyday Practical Reasoning

9.1 Four Open Questions

An important role for propositional beliefs can be found in the kind of practical reasoning that we perform on a daily basis, such as: "since today is Saturday, the grocery store is open today and will be closed tomorrow; so let's go today." It seems natural to say that the agent *believes* or *accepts* that today is Saturday, and that she reasons *directly* with the propositions that she believes but may not be fully certain of. Such scenarios exemplify everyday practical reasoning but, unfortunately, they are largely ignored in standard decision theory. Bayesian decision theory concerns global properties of rational preference and makes no claims about how particular preferences might be justified or generated in a conscious mind that can perform everyday practical reasoning. So Bayesian decision theory leaves some room for everyday practical reasoning to play a role in rational decision-making: everyday practical reasoning may serve as one of the deliberative means for Bayesian ends. But that raises some fudamental questions:

(Q1) Are there decision rules that connect propositional beliefs and qualitative desires to preferences over acts in everyday practical reasoning?

(Q2) What sort of logic should govern propositional beliefs in everyday practical reasoning? To what extent is that logic *necessary* for qualitative decision?

(Q3) What kinds of qualitative decisions are representable as results of everyday practical reasoning?

(Q4) Under what circumstances does everyday practical reasoning agree with the Bayesian ideal of expected utility maximization?

This section will sketch partial answers to (Q1) and (Q4), and will point to possible directions for obtaining answers to (Q2) and (Q3).

(Q1)-(Q4) are open questions in the literature to date. Qualitative decision-making has been studied extensively in logic-based artificial intelligence, but question (Q1) receives almost no treatment.[25] The role of propositional beliefs in decision-making has been studied in philosophy and economics. Bratman (1987) and Cohen (1989), for example, expound the idea that propositional beliefs can serve as premises in practical reasoning. Morris (1996) shows how the logic for beliefs may depend on one's possible preference relations, but he does so by defining propositional beliefs to be fully certain propositions, which excludes everyday beliefs such as "the grocery store is open today".[26] No such restriction is imposed in the following development. Some representation theorems for qualitative decisions have been provided in the literature, but none of them concerns everyday practical reasoning as illustrated in the grocery example. Brafman & Tennenholtz (1996), for instance, provide a representation theorem that completely characterizes the Maximin decision rule. Dubois et al. (2002) prove that a class of preference relations is representable by a qualitative decision rule they call the likely dominance rule, together with belief representations based on possibility measures, which is not intended to be about everyday practical reasoning.

The following explains how we may take some steps toward answering questions (Q1)-(Q4).

9.2 In Search of Decision Rules

To find a partial answer to (Q1), let us recall the grocery store example mentioned in the introduction. It can be understood in terms of the following, more explicit reasoning:

(1) Today is Saturday.

(2) So the grocery store is open today and will be closed tomorrow.

[25] For reviews, see Doyle and Thomason (1999) and Thomason (2009).

[26] To be more precise, Morris defines prpositional beliefs to be the negations of decision-theoretically null propositions.

(3) So going today would result in some food, and going tomorrow would result in no food.

(4) So going today would result in a better outcome than going tomorrow.

(5) Therefore, I prefer going today rather than tomorrow.

The agent believes (1), from which she defeasibly infers (2). Given (2), the agent reasons to (3), concerning which actions would produce which outcomes. Then, the agent's desirability order over outcomes leads to (4). The reasoning from (4) to (5) goes beyond logic; it rests on a decision rule so banal as to easily escape notice:

> (**Cliché Rule**) Prefer act a to act b if you believe that
> the outcome of a is more desirable than the outcome of b.

The Cliché Rule can be taken as a kind of dominance argument. Assuming the Hintikka-Kripke semantics for beliefs, one's propositional belief state can be modeled as a set B of possible worlds, so that one believes a proposition if and only if that proposition is true in every world in B. Then the Cliché Rule reads as follows. Prefer act a to act b if: the proposition "the outcome of a is more desirable than the outcome of b" is true in every world in your belief state B, namely, for each possible state s of the world in B, the outcome of a produced in s is more desirable to you than the outcome of b produced in s. So the Cliché Rule is like a dominance argument that requires quantification over, not all possible states of the world, but only those contained in your belief state B. Here I assume the Hintikka-Kripke semantics just for the sake of presentation. Propositional beliefs will be modeled in a more liberal way without assuming the defining features of the Hintikka-Kripke semantics (i.e., beliefs are closed under entailment with arbitrarily many premises).

To formulate the Cliché Rule in decision-theoretic terms, it is most convenient to adopt Savage's (1954) framework. Let a decision problem be modeled by an ordered pair (S, O), where S is a non-empty set of *states* that are mutually exclusive and jointly exhaustive and O is a non-empty set of *outcomes*. The states in S can be understood to correspond to what we call the answers to question $\mathcal{Q}$ when we discuss belief revision that tracks Bayes conditioning, but here let us stick to the more standard notation and terminology in decision theory. An *act* is a function a from S to O, where $a(s)$ denotes the outcome that a would produce if s were the actual state. An agent's preference over acts is modeled by a binary relation $\succeq$ over acts, where $a \succeq b$ means that act a is at least as preferable as act b to the agent. Call $\succeq$ a *preference relation*. Let $\succeq$ induce the "equally preferable" relation $\sim$ and the "strictly preferable" relation $\succ$ by the standard definition.[27] Preference is supposed

[27] Namely, $a \sim b$ iff $a \succeq b \succeq a$, and $a \succ b$ iff $a \succeq b \not\succeq a$.

to be determined by desire and belief. Let one's qualitative desire be modeled by a binary relation $\geq$ over outcomes, where $o \geq o'$ means that outcome o is at least as desirable as outcome o' to the agent. Call $\geq$ a *desirability order*. Let $\geq$ induce the "equally desirable" relation $\equiv$ and the "more desirable" relation $>$ by the standard definition.[28] A *proposition* is a subset of S, which is true in all and only the states it contains. So, for example, the proposition expressed by 'the outcome of a is more desirable than the outcome of b' is identified with the set $\{s \in S : a(s) > b(s)\}$, which is abbreviated as $[\![a > b]\!]$. In general, let the notation $[\![a\,R\,b]\!]$ be defined as follows, where a, b are acts and R is a binary relation between outcomes:[29]

$$\begin{aligned} [\![a\,R\,b]\!] &= \text{the proposition that the outcome produced by } a \text{ bears} \\ &\ \text{relation } R \text{ to the outcome produced by } b \\ &= \{s \in S : a(s) \text{ bears relation } R \text{ to } b(s)\}. \end{aligned}$$

The agent's propositional beliefs are modeled by a set Bel, which contains exactly the propositions that she believes. When proposition A belongs to Bel, write $\mathsf{Bel}(A)$ to emphasize what it means: the agent believes A. Then the Cliché Rule can be restated as follows:

(**Cliché Rule**) $a \succ b$ if $\mathsf{Bel}([\![a > b]\!])$.

That is most likely just a partial answer to (Q1), and we need to examine more examples in everyday practical reasoning to see whether there are other decision rules that cannot be reduced to the Cliché Rule.

9.3 Consistency with Bayesian Preference

Everyday practical reasoning is not meant to replace full Bayesian rationality. When, for example, the agent is considering complex stock investments, the situation is typically so uncertain that the agent has no ground for judging whether one investment will yield a better outcome than an alternative investment. In that case the Cliché Rule would say nothing about what to prefer. When everyday practical reasoning is silent, the agent can move to a more refined decision procedure, such as the Bayesian ideal of expected utility maximization. What is important in our daily life is, rather, that when everyday practical reasoning yields some recommendation, it may serve as a qualitative means for Bayesian ends.

[28] Namely, $o \equiv o'$ iff $o \geq o' \geq o$, and $o > o'$ iff $o \geq o' \not\geq o$.

[29] That notation is borrowed from statistics. Think of a as a random variable: proposition/event $[\![a > 3]\!]$ is a shorthand for $\{s \in S : a(s) > 3\}$.

Suppose, for example, that an agent is running out of food at home and believes that the grocery store is open today:

	s_1: Store Open Today	s_2: Store Closed Today	
a: Going Today	Satisfying Dinner	Energy Wasted	$\leftarrow$ the preferred act
b: Not Going Today	No Dinner	Energy Saved	
Comparison	$a(s) > b(s)$	$a(s) < b(s)$	

By applying the Cliché Rule, the agent prefers going today rather than not ($a \succ b$). That particular application maximizes expected utility if and only if:

$$\sum_{i=1,2} p(s_i)U(a(s_i)) \quad > \quad \sum_{i=1,2} p(s_i)U(b(s_i)),$$

where p and U are the agent's underlying probability measure and utility function, respectively.[30] So it suffices to require that:

$$p(\textsf{Store Open Today}) \quad > \quad \frac{1}{1 + \frac{U(\textsf{Satisfying Dinner}) - U(\textsf{No Dinner})}{U(\textsf{Energy Saved}) - U(\textsf{Energy Wasted})}}.$$

Hence, for expected utility to be maximized in the present case, it suffices that $p(\textsf{Store Open Today})$ be high enough; namely, the propositional belief to which the Cliché Rule applies is sufficiently probable. Or, it suffices that $U(\textsf{Satisfying Dinner}) - U(\textsf{No Dinner})$ is high enough; namely, the extra gain of utility for the preferred act a is sufficiently high given the truth of the propositional belief (s_1). Or, it suffices that $U(\textsf{Energy Saved}) - U(\textsf{Energy Wasted})$ is low enough; namely, the extra loss in utility for the preferred act a is sufficiently low given the falsity of the propositional belief (s_2).

A more stringent condition has to be met if *all* possible applications of the Cliché Rule are guaranteed to be consistent with Bayesian preference in a fixed decision problem. Let a decision scenario be defined by a quadruple $(p, U, \textsf{Bel}, \geq)$, which characterizes aspects of a given agent's mental state. (1) p is a probability distribution defined on S, which is assumed to contain only finitely many states. (2) $U : O \to \mathbb{R}$ is a utility function. (3) $\textsf{Bel}$ is a conditional belief set over the propositions over S. (4) $\geq$ is a desirability order on O. In the rest of this section, definitions, principles, and theorems are all relative to an arbitrary, fixed decision scenario $(p, U, \textsf{Bel}, \geq)$.

The agent's underlying Bayesian preference $\succeq_{\textsf{Bayes}}$ is defined as follows: for all acts a, b,

$$a \succeq_{\textsf{Bayes}} b \quad \text{iff} \quad \sum_{s \in S} P(s)U(a(s)) \;\geq\; \sum_{s \in S} P(s)U(b(s)).$$

[30] $P(s)$ abbreviates the strictly correct notation $P(\{s\})$.

Say that the Cliché Rule is *maximally* consistent with Bayesian preference if and only if:

> (**Maximal Consistency with the Bayesian Preference**) For *all* acts $a, b : S \to O$,
>
> $\mathsf{Bel}([\![a > b]\!]) \quad \Longrightarrow \quad a \succ_{\mathsf{Bayes}} b.$

A necessary and sufficient condition is provided in the following result:

Proposition 5. *Suppose that state space S and outcome space O are both finite, that there are outcomes $o, o' \in O$ such that $o > o'$, and that $\mathsf{Bel}(\top)$. Then,* (*Maximal Consistency with Bayesian Preference*) *holds iff the following two conditions hold:*

1. *for all outcomes $o, o' \in O$,*
$$o > o' \quad \Longrightarrow \quad U(o) > U(o');$$

2. *for all propositions $A \subseteq S$,*
$$\mathsf{Bel}(A) \quad \Longrightarrow \quad p(A) > \frac{\Delta U}{\Delta U + \delta U},$$
where:
$$\begin{aligned} \delta U &= \min\{\, |U(o) - U(o')| : o, o' \in O \text{ and } o > o' \,\}, \\ \Delta U &= \max\{\, |U(o) - U(o')| : o, o' \in O \,\}. \end{aligned}$$

It has long been speculated that propositional belief is related to high probability. In light of the second condition, each propositional belief has to be highly probable if the Cliché Rule is maximally consistent with Bayesian preference. Note that it is *not* required that each highly probable proposition be believed, which is one of the premises that lead to the lottery paradox (Kyburg 1961). Furthermore, the high probability threshold ($\Delta U/(\Delta U + \delta U)$) depends on contextual factors such as O (the set of outcomes that are relevant for the agent to consider in the present context) and U (the utility function of the agent in the present context).

9.4 Compatibility with the Tracking Condition

To ensure that application of the Cliché Rule is compatible with the Bayesian preference, it suffices that each propositional belief be highly probable. Is that requirement consistent with the desideratum that propositional belief revision tracks Bayes conditioning? The answer is positive. Recall that odds-based acceptance rules track conditioning. The following proposition shows how we may construct odds-based acceptance rules for which propositional beliefs are guaranteed to be highly probable.

Proposition 6. *Let* B *be a consistent, acceptance rule that is driven by odds-based assignment of plausibility orders. Let the plausibility order* $\prec_p$ *assigned to* p *be defined by*

$$H_i \prec_p H_j \quad \Longleftrightarrow \quad p(H_i)/p(H_j) > t_j. \tag{29}$$

Then, for each probability measure p *in* $\mathcal{P}$ *and for each proposition* A *in* $\mathcal{A}$*, if* A *is accepted by* B *at* p *(i.e.* $\mathsf{B}_p(\top) \models A$*), then:*

$$p(A) \quad > \quad \frac{1}{1 + \sum_{j \in I \setminus \{i\}} t_j^{-1}}. \tag{30}$$

If the number of complete answers to $\mathcal{Q}$ *(i.e., the number of states in* S*) is finite, say* n^**, and if the assignment of thresholds is constant, say* t^**, then the probability lower bound in the above can be simplified as follows:*

$$p(A) \quad > \quad \frac{1}{1 + \left(\frac{n^* - 1}{t^*}\right)}. \tag{31}$$

So, for each fixed finite n^*, we can adjust the constant threshold t^* within interval $[n^* - 1, \infty)$ to make the probability lower bound be any value in the interval $[0.5, 1)$.

9.5 What's Next?

It is likely that there are more decision rules to be discovered in everyday practical reasoning, so a more complete answer to (Q1) remains to be explored. But no matter what those rules would be, we can and should find the conditions under which they are compatible with maximization of expected utility, as we have done with the Cliché Rule, which illustrates the general strategy for answering (Q4). Question (Q3) is much harder but I am optimistic: perhaps (i) there is a set of axioms for preference that jointly capture the idea that the preference relation depends only on certain qualitative features of the decision problem and (ii) we can prove that each preference relation $\succeq$ that satisfies those axioms can be represented as the result of applying only the decision rules for everyday practical reasoning. Such a preference relation $\succeq$ is expected to be *incomplete* in the senses that there are acts a, b such that $a \not\succeq b$ and $b \not\succeq a$—in that case, application of the decision rules for everyday practical reasoning yields no recommendation for the choice between a and b. And, for answering (Q2), hopefully the representation result will yield a logic for propositional beliefs, just as Savage's representation yields the laws of probability as constraints on degrees of belief.

10 Conclusion

It is impossible for accretive (and thus AGM) belief revision to track Bayesian conditioning, on pain of failing to be sensible. But dynamic consonance is feasible: just adopt Shoham revision and an acceptance rule with the right geometry. The resulting theory for uncertain acceptance solve the paradoxes, old and new.

Earlier work on the relation between propositional belief revision (i.e. nonmonotonic logic) and Bayesian conditioning focuses primarily on the Kyburgian solution to the lottery paradox (Hawthorne 1996, Hawthorne and Makinson 2007, Paris and Simmonds 2009, and Makinson 2011). So their systems drop the conjunctive closure of conclusions and, hence, depart from most of the non-monotonic logical systems that have been developed for their major application: planning and practical reasoning in artificial intelligence.[31] Therefore, the present dissertation has the potential to establish a closer connection between Bayesian conditioning and some standard non-monotonic logics used in planning.

Since nonmonotonic logic is a formalism for expressing non-iterated conditionals, the proposed thesis also has potential applications to a new probabilistic semantics for conditionals that improves upon Adams' (1975) complicated definition of validity and that improves upon Pearl's (1989) unrealistic requirement that propositional beliefs have probabilities infinitesimally close to 1. Probabilistic semantics for conditionals is of considerable interest to philosophical logicians who are sympathetic to the thesis that indicative conditionals do not have truth conditions (e.g., Ramsey 1929, Adams 1975, Gibbard 1981, Edgington 1991, Bennett 2003).

The present thesis will also be of interest to those who are impressed by the Bayesian ideal, but who nonetheless wish to entertain propositional beliefs for various reasons. For example, propositions may be deemed instrumental due to limitations of cognitive architecture, or due to the need to communicate plans linguistically for the purpose of coordination. Or a committed Bayesian may simply find herself charged with the task of maintaining a large propositional database.

More ambitiously, one could continue the work in section 9 and seek a propositional decision theory that stands to propositional belief revision in the manner in which Bayesian decision theory stands to Bayes conditioning.

11 Acknowledgements

I am indebted to Kevin Kelly, my mentor and best friend in philosophy. *Each* part of the present work would be impossible without many enthusiastic discussions with him and his endlessly patient support. I am also indebted to Hannes Leitgeb for

[31] The literature is huge; for reviews see Shanahan (1997), Minker (2000) and Thomason (2009).

discussion and comparison of our views on acceptance and for his kind invitation to the *Round Table on Acceptance* at the Munich Center for Mathematical Philosophy (February 2012). I am indebted to Clark Glymour for discussions on philosophical issues in the dissertation. I am indebted to David Makinson for discussions on the logical foundations of the dissertation. I am also indebted to David Etlin for his comments in the 8th Formal Epistemology Workshop (2011). I am indebted to Teddy Seidenfeld for pointing out that Levi already proposed a version of weighted odds threshold rules. Finally, I am indebted to Horacio Arló-Costa and Arthur Paul Pedersen for discussions.

Bibliography

Adams, E.W. (1975) *The Logic of Conditionals*, Dordrecht: D. Reidel.

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.

Bennett, J. (2003) *A Philosophical Guide to Conditionals*, Oxford: Clarendon Press.

Brafman, M. and M. Tennenholtz. (1996) "On the Foundations of Qualitative Decision Theory", In *Proceedings of the 13th National Conference on Artificial Intelligence* (AAAI-96, Portland, Ore.), 1291-1296.

Bratman, M. (1987) *Intention, Plans, and Practical Reason*, Cambridge, MA: Harvard University Press.

Cohen, L. J. (1989) "Belief and Acceptance", *Mind* 98: 367-389.

Dubois, D., H. Fargier, H. Prade, and P. Perny (2002) "Qualitative Decision Theory: From Savage's Axioms to Nonmonotonic Reasoning", *Journal of the ACM*, 49(4): 455-495.

Douven, I. (2002) "A New Solution to the Paradoxes of Rational Acceptability", British Journal for the Philosophy of Science 53: 391-410.

Doyle, J. and R. Thomason, (1999) "Background to Qualitative Decision Theory", *AI Magazine*, 20(2): 55-68.

Edgington, D. (1991) "The Mystery of the Missing Matter of Fact", *Proceedings of the Aristotelian Society Supplementary*, 65: 185-209.

Foley, R. (1992) "The Epistemology of Belief and the Epistemology of Degrees of Belief", *American Philosophical Quarterly* 29(2): 111-124.

Foley, R. (2009) "Beliefs, Degrees of Belief, and the Lockean Thesis", in *Degrees of Belief*, F. Huber, C. Schmidt-Petri (eds.), *Synthese Library* 342, Springer.

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

Gibbard, A. (1981) "Two Recent Theories of Conditionals" in Harper, Stalnaker and Pearce (eds.) 1981.

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

Harper, W., (1977) "Rational Conceptual Change", *PSA* 1976: 462âĂŞ494.

Hawthorne, J. (1996) "On the Logic of Nonmonotonic Conditionals and Conditional Probabilities", *Journal of Philosophical Logic* 25: 185-218.

Hawthorne, J. and L. Bovens (1999) "The Preface, the Lottery, and the Logic of Belief", *Mind* 108, 241-264.

Hawthorne, J. and D. Makinson (2007) "The Quantitative/Qualitative Watershed for Rules of Uncertain Inference" *Studia Logica* 86: 249-299.

Jeffrey, R. (1970) "Dracula Meets Wolfman: Acceptance vs. Partial Belief", in *Induction, Acceptance, and Rational Belief*, M. Swain (ed.): 157-185.

Kaplan, M. (1981) "A Bayesian Theory of Rational Acceptance", *Journal of Philosophy* 78(6): 305-330.

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.

Kyburg, H. (1994) "Believing on the Basis of the Evidence", *Computational Intelligence* 10: 3-20.

Kyburg H. and Teng, C.M. (2001) *Uncertain Inference*, Cambridge 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. (1980) *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability and Chance*, MIT.

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. (2012), “A Geo-logical Solution to the Lottery Paradox”, *Synthese*, 186(2), 531-575.

Maher, P. (1993) *Betting on Theories*, Cambridge: Cambridge University Press.

Makinson, D. (2011) “Conditional Probability in the Light of Qualitative Belief Change”, *Journal of Philosophical Logic* 40(2): 121-153.

Morris, S. (1996) “The Logic of Belief and Belief Change: A Decision Theoretic Approach”, *Journal of Economic Theory*, 29(1):1-23.

Paris, J. and R. Simmonds (2009) “O is Not Enough”, *Review of Symbolic Logic* 2: 298-309.

Pearl, J. (1989) “Probabilistic Semantics for Nonmonotonic Reasoning: A Survey”, In Proc. First International Conference on Principles of Knowledge Representation and Reasoning (KR ’89): 505-516. Reprinted in G. Shafer and J. Pearl (Eds.) (1990) *Readings in Uncertain Reasoning*, San Francisco: Morgan Kaufmann, 699-710.

Peirce, C.S. (1878) “How to Make Our Ideas Clear”, *Popular Science Monthly* 12: 286-302.

Pollock, J. (1983) “Epistemology and Probability”, *Synthese*, 55(2): 231-252.

Pollock, J. (1990) *Nomic Probability and the Foundations of Induction*, Oxford: Oxford University Press.

Ramsey, F.P. (1929) “General Propositions and Causality”, reprinted in F.P. Ramsey (1990) *Philosophical Papers* ed. by H. A. Mellor, Cambridge: Cambridge University 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.

Shanahan, M. (1997) *Solving the Frame Problem*, Cambridge, Massachusetts: The MIT Press.

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

Thomason, R. (2009) “Logic and Artificial Intelligence”, *The Stanford Encyclopedia of Philosophy* (Spring 2009 Edition), Edward N. Zalta (ed.), URL = `http://plato.stanford.edu/archives/spr2009/entries/logic-ai/`

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 B is consistent, corner-monotone, accretive (i.e. satisfies axioms Inclusion and Preservation), and tracks conditioning. Suppose further that B is not skeptical. It suffices to show that B is opinionated, which 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), let p be 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:

$$\begin{aligned}
\|p - p'\| &= \sqrt{\sum_{H_n \in \mathcal{Q}} (p(H_n) - p'(H_n))^2} && (32)\\
&= \sqrt{(p(H_i) - p'(H_i))^2 + (p(H_j) - p'(H_j))^2} && (33)\\
&< \sqrt{\left(\frac{r}{\sqrt{2}}\right)^2 + \left(\frac{r}{\sqrt{2}}\right)^2} = r. && (34)
\end{aligned}$$

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:

$$\|p - p'\| < \sqrt{\left(\frac{r}{\sqrt{6}}\right)^2 + \left(\frac{r}{\sqrt{6}}\right)^2 + \left(\frac{2r}{\sqrt{6}}\right)^2} = r, \qquad (35)$$

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 $\triangle\, p_1\, p_2\, p_3$ denote the convex hull of p_1, p_2, p_3:

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

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 $\mathcal{P}$. So, for each pair of distinct complete answers H_i, H_j to $\mathcal{Q}$, $\overline{h_i\, h_j}$ is 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:

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

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 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, B accepts H_i at p. Then B also accepts H_i at $p|_{H_i \vee H_j}$, by Preservation and conditioning-tracking. Furthermore, B accepts H_i as strongest at $p|_{H_i \vee H_j}$, since B is consistent and H_i is a complete answer. Then B accepts H_i as strongest at h_i, since 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 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 $\triangle\, h_i\, h_j\, h_m$ (figure 11.a), relative to which points a, b, c, d are defined as follows. Let a be the endpoint of L_{im} that is closest to h_m; namely, a is the credal state in

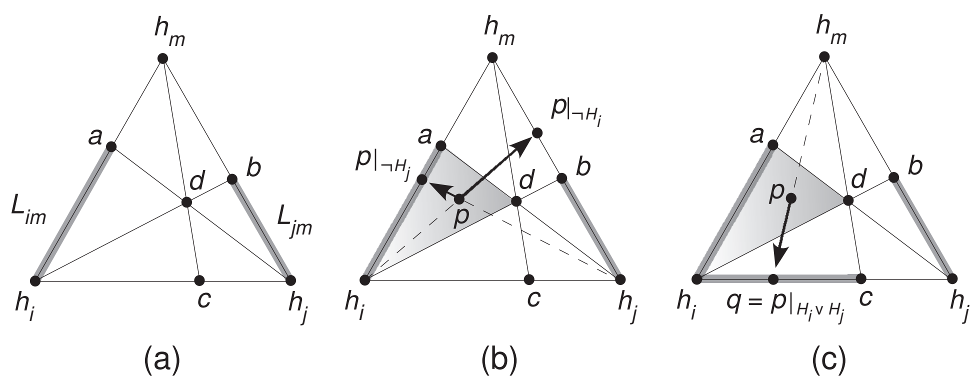

Figure 10: Why every accretive rule that tracks conditioning fails to be sensible

$\overline{h_i\, h_m}$ 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 $\overline{a\, h_j}$ and $\overline{b\, h_i}$. Let c be the unique credal state in $\overline{h_i\, h_j}$ such that $\overline{c\, h_m}$ contains d.

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

Proof. Consider an arbitrary point p in the interior of $\triangle\, a\, d\, h_i$ (figure 11.b). Argue as follows that B accepts H_i as strongest at p. Since posterior state $p|_{\neg H_j}$ exists and falls inside L_{im}, $\mathsf{B}_{p|_{\neg H_j}}(\top) = H_i$. So, since B tracks conditioning, $\mathsf{B}_p(\neg H_j) = H_i$. Then, since B satisfies Inclusion, $\mathsf{B}_p(\top) \wedge \neg H_j \models H_i$. So we have only three possibilities for $\mathsf{B}_p(\top)$:

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

since B is consistent and the complete answers are mutually exclusive. Rule out the last two possibilities as follows. Suppose for reductio that $\mathsf{B}_p(\top) = H_j$ or $H_i \vee H_j$. Then, since B satisfies Preservation and $\mathsf{B}_p(\neg H_i)$ is consistent with new information $\neg H_i$, we have that:

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

The left-hand side equals $\mathsf{B}_{p|_{\neg H_i}}(\top)$ by conditioning-tracking, and the right-hand side equals H_j by the reductio hypothesis. So $\mathsf{B}_{p|_{\neg H_i}}(\top) \models H_j$. But B is consistent and H_j is a complete answer, so $\mathsf{B}_{p|_{\neg H_i}}(\top) = H_j$. 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 11.b). Ruling out the last two possibilities for $\mathsf{B}_p(\top)$, we conclude that $\mathsf{B}_p(\top) = H_i$. So we have established the first statement. The second statement follows by symmetry. □

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

Proof. By lemma 2, $\mathsf{B}_{h_i}(\top) = H_i$. 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 11.c). So $\mathsf{B}_p(\top) = H_i$, by the preceding lemma. Then, since B satisfies Preservation,

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

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

Lemma 5. *Continuing from the preceding lemma,* $\overline{h_i\, h_j}$ *contains at most one point at which* 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 $\mathsf{B}_p(\top) = H_i \vee H_j$, then $p = c$ (figure 11.c). □

Lemma 6. *Every edge $\overline{h_i\,h_j}$ of $\mathcal{P}$ contains at most one point at which 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 $\mathcal{P}$. The present lemma follows immediately from applying the preceding lemma to the simplex $\triangle\, h_i\, h_j\, h_m$. □

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. □

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:

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

C Proof of Theorem 3

Right-to-Left Side. Let B be driven by an odds-based assignment $(\prec_p : p \in \mathcal{P})$ of plausibility orders. The corner-monotonicity of 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 $\overline{p\,h_i}$. To see that B tracks conditioning (i.e. that $\mathsf{B}_p(E) = \mathsf{B}_{p|_E}(\top)$), since B is Shoham-driven, it suffices to show that an answer is most plausible in ${\prec_p}|_E$ if and only if it is most plausible in $\prec_{p|_E}$, which follows from the odds-based definition of $\prec_p$ and preservation of odds by Bayesian conditioning. $\square$

Left-to-Right Side. Suppose that 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, for each $p \in \mathcal{P}$, $\prec_p$ is odds-based. For each pair of distinct indices i, j in I, define odds threshold $t_{ij} \in [0, \infty]$ and inequality $\rhd_{ij} \in \{>, \geq\}$ by:

$$\mathsf{Odds}_{ij} = \left\{ \frac{q(H_i)}{q(H_j)} : q \in \mathcal{P}, q(H_i) + q(H_j) = 1, H_i \prec_q H_j \right\}; \tag{36}$$

$$t_{ij} = \inf \mathsf{Odds}_{ij}\,; \tag{37}$$

$$\rhd_{ij} = \begin{cases} \geq & \text{if } t_{ij} \in \mathsf{Odds}_{ij}\,, \\ > & \text{otherwise.} \end{cases} \tag{38}$$

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

$$H_i \prec_q H_j \iff q(H_i)/q(H_j) \ \rhd_{ij}\ t_{ij}. \tag{39}$$

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 (27) holds with respect to odds thresholds (37) and inequalities (38):

$$H_i \prec_p H_j \iff p(H_i)/p(H_j) \ \rhd_{ij}\ t_{ij}\ . \tag{40}$$

Note that, since H_i and H_j are in the domain of $\prec_p$, $p(H_i \vee H_j) = p(H_i) + p(H_j) > 0$, so $p|_{(H_i \vee H_j)}$ is defined. Then:

$$\begin{array}{rcll} H_i \prec_p H_j & \iff & H_i \prec_p |_{(H_i \vee H_j)} H_j & \\ & \iff & H_i \prec_{p|_{(H_i \vee H_j)}} H_j & \text{by theorem 2;} \\ & \iff & H_i \prec_q H_j & \text{by defining } q \text{ as } p|_{(H_i \vee H_j)}; \\ & \iff & q(H_i)/q(H_j) \ \rhd_{ij} \ t_{ij} & \text{by (39);} \\ & \iff & p(H_i)/p(H_j) \ \rhd_{ij} \ t_{ij} & \text{since } q = p|_{(H_i \vee H_j)}. \end{array}$$

□

D Proof of Propositions 1-4

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

Proof of Proposition 2. Consistency is an immediate consequence of proposition 1. So it suffices to show, for each p, that the relation $\mathsf{B}_{p|_E}(\top) \models H$ between E and H satisfies Hypothetico-deductive Monotonicity (3) and Case Reasoning (4). That relation is equivalent to relation $\mathsf{B}_p(E) \models H$ between E and H (by tracking conditioning). Since B is Shoham-driven, the 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). □

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

Proof of Proposition 4. To see that B is a weighted odds probability-threshold rule, argue as follows ($\max_k w_k p(H_k)$ exists because there is no infinite ascending chain of weighted probabilities, as shown in the proof of the preceding proposition):

$$\begin{aligned}
\mathsf{B}_p(\top) &= \bigvee\{H_j \in \mathcal{Q} : H_j \text{ is minimal in } \prec_p\} && (41)\\
&= \bigvee\{H_j \in \mathcal{Q} : \max_k w_k p(H_k)/w_j p(H_j) \not> t_j\} && (42)\\
&= \bigwedge\{\neg H_i \in \mathcal{Q} : \max_k w_k p(H_k)/w_i p(H_i) > t_i\} && (43)\\
&= \bigwedge\left\{\neg H_i \in \mathcal{Q} : \frac{w_i\, p(H_i)}{\max_k w_k\, p(H_k)} < \frac{1}{t_i}\right\}. && (44)
\end{aligned}$$

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, by equation (44), at 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

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

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 $\lim_{n\to\infty} \|h_i - p_n\| = 0$. So $\lim_{n\to\infty} p_n(H_i) = 1$ and $\lim_{n\to\infty} p_n(H_j) = 0$. 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, each complete answer H_k distinct from H_i is rejected by the rule over B_i. So 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 $\lim_{n \to \infty} \|p^* - p_n\| = 0$. So, $\lim_{n \to \infty} w_1 p(H_1)/w_2 p(H_2) = w_1 p^*(H_1)/w_2 p^*(H_2) = 1$. Then, since $t_2 > 1$, there exists sufficiently large m such that $w_1 p_m(H_1)/p_m(H_2) < t_2$. But that contradicts $w_1 p_m(H_1)/p_m(H_2) > t_2$, 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. □

E Proof of Theorem 4

To prove theorem 4 is proved in exact parallel to the proof of theorem 1. Let $\mathcal{Q}$ have at least three complete answers. Suppose that R is consistent, corner-monotone, accretive, and tracks conditioning. Suppose further that R is non-skeptical. It suffices to show that B is opinionated. Replace the original definition of L_{ij} by:

$$L_{ij} \quad = \quad \{p \in \overline{h_i\, h_j} : \mathsf{R}(p, H_i)\}.$$

Then lemmas still 1 and 2 hold—for exactly the same reason. Lemmas 3-7 have the following lemmas, 8-12, as their counterparts. We only need to prove lemma 8 (the counterpart of lemma 3), because it is the only lemma that makes use of the second, new part of the tracking property.

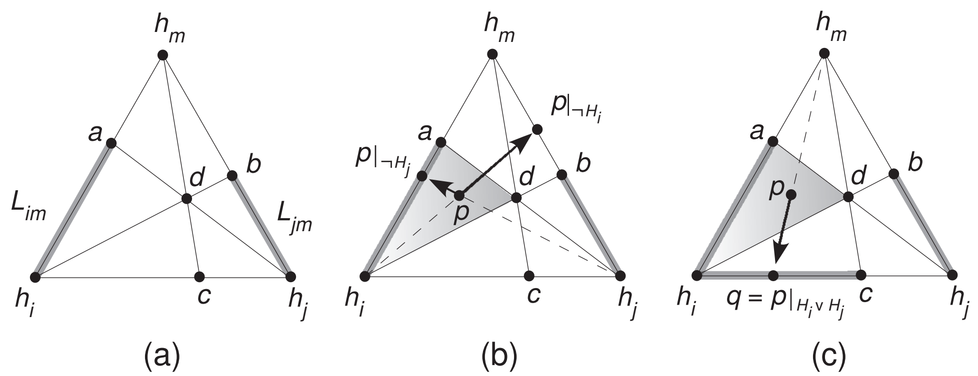

Figure 11: Reproduction of figure 10 for the sake of convenience

Lemma 8. *Let h_i, h_j, h_m be distinct corners of $\mathcal{P}$. Consider two-dimensional simplex $\triangle\, h_i\, h_j\, h_m$, relative to which points a, b, c, d are defined as above. Then* $\mathsf{R}(p, H_i)$ *for each point p in the interior of $\triangle\, a\, d\, h_i$. Furthermore,* $\mathsf{R}(p, H_j)$ *for each point p in the interior of $\triangle\, b\, d\, h_j$.*

Proof. Let p be a point in the interior of $\triangle\, a\, d\, h_i$. Argue as follows that $\mathsf{R}(p, H_i)$. Since $\mathsf{R}(p|_{\neg H_j}, H_i)$, it follows from the second part of the tracking property that there exists belief revision method B such that $\mathsf{R}(p, \mathsf{B})$ and $\mathsf{B}|_{\neg H_j}(\top) = H_i$. So $\mathsf{B}(\neg H_j) = H_i$. (The above is the only step that makes use of the second part of the tracking property.) Then, by the argument in lemma 3 (with B_p replaced by the present B), $\mathsf{B}(\top) = H_i$. So $\mathsf{R}(p, H_i)$. The second statement follows by symmetry. $\square$

Lemma 9. *Continuing from the preceding lemma,* $\mathsf{R}(p, H_i)$ *for each point p in the interior of $\overline{c\, h_i}$. Furthermore,* $\mathsf{R}(p, H_j)$ *for each point p in the interior of $\overline{c\, h_j}$.*

Lemma 10. *Continuing from the preceding lemma, $\overline{h_i\, h_j}$ contains at most one point p such that $\mathsf{R}(p, H_i \vee H_j)$.*

Lemma 11. *Every edge $\overline{h_i\, h_j}$ of $\mathcal{P}$ contains at most one point p such that $\mathsf{R}(p, H_i \vee H_j)$.*

Lemma 12. B *is opinionated.*

F Proof of Propositions 5 and 6

Proof of Proposition 5. If "if" side follows from the following calculation:

$$\begin{aligned} EU_p(a) - EU_p(b) &= \sum_{s \in S} p(s)\,[U(a(s)) - U(b(s))] \\ &= \sum_{s \in [\![a>b]\!]} p(s)\,[U(a(s)) - U(b(s))] \\ &\quad + \sum_{s \in [\![a \not> b]\!]} p(s)\,[U(a(s)) - U(b(s))] \\ &\geq \sum_{s \in [\![a>b]\!]} p(s)\,\delta U \;+\; \sum_{s \in [\![a \not\geq b]\!]} p(s)\,(-\Delta U) \\ &= p([\![a > b]\!])\,\delta U \;-\; p([\![a \not> b]\!])\,\Delta U \\ &> 0, \end{aligned}$$

where the last step holds because $p([\![a > b]\!]) > \frac{\Delta U}{\Delta U + \delta U}$. To prove the "only if" side, suppose that (Maximal Consistency with Bayesian Preference) holds. Suppose for reductio that condition 1 is false, namely that there exist outcomes o, o' such that $o > o'$ but $U(o) \leq U(o')$. Let f_o be the constant act that maps all states to o, and $f_{o'}$ the constant act that maps all states to o'. Since $o > o'$, $[\![f_o \succ f_{o'}]\!] = \top$. And since $\mathsf{Bel}(\top)$ by hypothesis, $\mathsf{Bel}([\![f_o \succ f_{o'}]\!])$. But $f_o \not\succ_{\mathsf{Bayes}} f_{o'}$, because the expected utility of f_o equals $U(o)$, the expected utility of $f_{o'}$ equals $U(o')$ and $U(o) \leq U(o')$. So we have both that $\mathsf{Bel}([\![f_o \succ f_{o'}]\!])$ and that $f_o \not\succ_{\mathsf{Bayes}} f_{o'}$, which contradicts (Maximal Consistency with Bayesian Preference). Suppose for reductio that condition 2 is false, namely that there exists proposition A such that $\mathsf{Bel}(A)$ and $p(A) \leq \frac{\Delta U}{\Delta U + \delta U}$. By hypothesis, there are only finitely many outcomes, so ΔU exists. By hypothesis, there are two outcomes o, o' such that $o > o'$, so δU exists. The existence of ΔU and δU guarantees that there exist outcomes o_1, o_2, o_3, o_4 in O such that $o_4 - o_1 = \Delta U$, $o_3 - o_2 = \delta U$, and $o_3 > o_2$. Let a be the act that maps all states in A to o_3 and all the other states to o_1. Let b be the act that maps all states in A to o_2 and all the other states to o_4. It follows that $[\![a \succ b]\!] = A$. And, since $\mathsf{Bel}(A)$, we have that $\mathsf{Bel}([\![a \succ b]\!])$. Since $p(A) \leq \frac{\Delta U}{\Delta U + \delta U}$, it is routine to verify that the expected utility of a is no more than that of b with respect to p and U and, hence, $a \not\succ_{\mathsf{Bayes}} b$. So we have both that $\mathsf{Bel}([\![a \succ b]\!])$ and that $a \not\succ_{\mathsf{Bayes}} b$, which contradicts (Maximal Consistency with Bayesian Preference). $\square$

Proof of Proposition 6. Let us first prove the following claim:

> Claim: For each complete answer H_i to $\mathcal{Q}$, if H_i is accepted by B at p, then $p(H_i) > \frac{1}{1+\sum_{j \in I \setminus \{i\}} t_j^{-1}}$.

Suppose that H_i is accepted by B at p. Since B is an odds-based acceptance rule, it is driven by an odds-based assignment of plausibility orders such that the plausibility order $\prec_p$ assigned to p is defined by:

$$H_m \prec_p H_n \iff p(H_m)/p(H_n) \ > \ t_n \ .$$

Since H_i is accepted by B at p and $\perp$ is not (because B is consistent), H_i is the uniquely minimal element in $\prec_p$. So we have: $H_i \prec_p H_j$ for each $j \in I \setminus \{i\}$. It follows, from the definition of $\prec_p$, that $p(H_i)/p(H_j) \ \rhd_j \ t_j$ for each $j \in I \setminus \{i\}$. So $p(H_i) \geq t_j \cdot p(H_j)$ for each $j \in I \setminus \{i\}$. Hence, $p(H_i) \geq 1/(1 + \sum_{j \in I \setminus \{i\}} t_j^{-1})$, which establishes the claim. Now prove the proposition as follows. Suppose that A is accepted by B at p. Since B is consistent, $\perp$ is accepted by B at p. It follows that the plausible order $\prec_p$ assigned to p has a minimal element, say H_i, that entails A. Let probability measure p' be defined as follows:

> $p'(H_i) = p(A)$ (note that H_i entails A);
>
> $p'(H_j) = 0$ for each H_j distinct from H_i that entails A;
>
> $p'(H_k) = p(H_k)$ for each H_k incompatible with A.

Argue as follows that H_i is accepted at p'. Let H_k be an arbitrary complete answer incompatible with A. Since H_k is rejected at p, there exists complete answer H_{j^*} compatible with A such that $H_{j^*} \prec_p H_k$, i.e. $p(H_{j^*})/p(H_k) \ <_k \ t_k$. By construction, $p(H_{j^*}) \leq p'(H_i)$ and $p(H_k) = p'(H_k)$. So $p'(H_i)/p'(H_k) \ \rhd_k \ t_k$ and, hence, $H_i \prec_{p'} H_k$. Furthermore, for each H_j distinct from H_i that entails A, $H_i \prec_{p'} H_j$ because $p'(H_j) = 0$. It follows that $H_i \prec_{p'} H_j$ for all H_j distinct from H_i. That is, H_i is a uniquely minimal element in plausibility order $\prec_{p'}$. So H_i is accepted at p'. Then:

$$p(A) = p'(H_i) > \frac{1}{1 + \sum_{j \in I \setminus \{i\}} t_j^{-1}},$$

where the quality follows from the construction of p' and the inequality follows from the claim we have proved. That completes the proof. $\square$