

A Joint Characterization of Belief Revision Rules¹

Franz Dietrich & Christian List & Richard Bradley
CNRS & UEA LSE LSE

7 September 2012 (first version October 2010)

Abstract

This paper characterizes different belief revision rules in a unified framework: Bayesian revision upon learning some event, Jeffrey revision upon learning new probabilities of some events, Adams revision upon learning some new conditional probabilities, and ‘dual-Jeffrey’ revision upon learning an entire new conditional probability function. Though seemingly different, these revision rules follow from the same two principles: *responsiveness*, which requires that revised beliefs be consistent with the learning experience, and *conservativeness*, which requires that those beliefs of the agent on which the learning experience is ‘silent’ (in a technical sense) do not change. So, the four revision rules apply the same revision policy, yet to different kinds of learning experience.

Keywords: Subjective probability, Bayes’s rule, Jeffrey’s rule, axiomatic foundations, unawareness

1 Introduction

Belief revision rules capture changes in an agent’s subjective probabilities. The most commonly studied example is *Bayesian revision*. Here, the agent learns that some event B has occurred. In response, he (or she) raises the probability of B to 1, while retaining all probabilities conditional on B . Other revision rules have also been studied. Under *Jeffrey revision*, the agent learns a new probability of some event, for instance a 90% probability that someone stands at the end of the corridor, prompted by vaguely seeing something or hearing a noise; or, more generally, he learns a new probability distribution of some random variable such as the level of

¹We are grateful for helpful feedback received from the audiences at presentations of earlier versions of this paper at D-TEA 2010 (HEC & Polytechnique, Paris, France, June 2010), LSE Choice Group Seminar (London School of Economics, UK, September 2010), Pluralism in the Foundations of Statistics (University of Kent, UK, September 2010), and Decisions, Games & Logic 2012 (Ludwig Maximilians University of Munich, Germany, June 2012). Although this paper presents a jointly authored project, Christian List and Richard Bradley wish to note that the bulk of the mathematical credit should go to Franz Dietrich.

rainfall or GDP. In response, he assigns the new distribution to the random variable, while retaining all probabilities conditional on the random variable (e.g., Jeffrey 1957, Shafer 1981, Diaconis and Zabell 1982, Grunwald and Halpern 2003). Jeffrey revision generalizes Bayesian revision, where the agent learns a probability of 1 of some event. Under a further revision rule, *Adams revision*, the agent learns a new conditional probability of some event given another, or, more generally, a new distribution of some random variable given another random variable, for instance a new distribution of the weather given the weather forecast, or of GDP given inflation (e.g., Bradley 2005, 2007, Douven and Romeijn 2012).² An excellent treatment of various forms of probabilistic belief and belief revision can be found in Halpern’s (2003) handbook.

Standard economic models rarely refer to non-Bayesian belief revision, but this is at the cost of an artificial modelling move. To achieve a Bayesian representation of a wide range of belief changes, they define an agent’s subjective probability function on a potentially very complex algebra of events: one that is constructed to contain an event for each possible ‘learning experience’ that might lead to a belief change. Suppose we wish to model an Olympic sprinter who raises his subjective probability of winning gold from 25% to 75% after experiencing an overwhelming feeling of strength before the race. If we define the sprinter’s subjective probabilities on a simple algebra consisting of all the subsets of the binary set $\Omega = \{winning, losing\}$, we cannot represent the sprinter’s belief change in Bayesian terms. The sprinter’s initial probability measure p on 2^Ω is given by $p(\{winning\}) = \frac{1}{4}$ and his new one p' by $p'(\{winning\}) = \frac{3}{4}$. The change from p to p' is not Bayesian, since there is no event $B \subseteq \Omega$ such that $p' = p(\cdot|B)$. This is due to the sparseness of Ω , which does not allow one to identify an event in 2^Ω representing the ‘observation’ leading to the belief change, i.e., the feeling of strength before the race. The Bayesian modeller would therefore re-define Ω more richly, for instance as $\Omega = \{winning, losing\} \times \{feeling strong, not feeling strong\}$. The new algebra 2^Ω contains not only the event of victory, $A = \{winning\} \times \{feeling strong, not feeling strong\}$, but also the event of the feeling of strength, $B = \{winning, losing\} \times \{feeling strong\}$. One can therefore model the belief change as Bayesian conditionalization on B , namely by specifying an initial probability measure $p : 2^\Omega \rightarrow [0, 1]$ and a new one $p' : 2^\Omega \rightarrow [0, 1]$ such that $p(A) = \frac{1}{4}$, $p'(A) = \frac{3}{4}$, and $p' = p(\cdot|B)$.

Many authors have raised concerns about this modelling practice, for example Jeffrey (1957), Shafer (1981), and Diaconis and Zabell (1982), who call the ascription of prior subjective probabilities to ‘many classes of sensory experiences [...] forced, unrealistic, or impossible’ (p. 823). The importance of non-Bayesian belief revision rules can be illustrated by considering two phenomena that call for them: *incomplete beliefs* and *unawareness*.

1. **Incomplete belief.** One drawback of the Bayesian re-modelling is that we must assume that the agent is able to assign prior probabilities to many complex events: our illustrative sprinter must assign subjective probabilities to the event that he will experience the feeling of strength, to the event that he will

²An important example of Adams revision is the learning of an equation $X = f(Y) + \epsilon$, where X and Y are two (possibly vector-valued) random variables, f is a deterministic function, and ϵ is a random error independent of Y . Learning this equation is equivalent to learning that X has a particular conditional distribution given Y .

experience it *and* lose the race, and so on. To give Bayesian accounts of further belief changes, for instance during and after the race, we must ascribe beliefs to the agent over an even more refined algebra of events, whose size grows exponentially with the number of belief changes to be modelled. This is not very plausible, since typical real-world agents have either no beliefs about such events or only imprecise ones.³ If we restrict the complexity of the event algebra, on the other hand, we may have to introduce non-Bayesian belief revision to capture the agent’s belief dynamics adequately within the smaller algebra.

2. **Unawareness.** The literature on unawareness suggests that a belief in an event (the assignment of a subjective probability to it) presupposes awareness of this event, where ‘awareness’ is understood, not as knowledge of the event’s occurrence (indeed, the agent may believe its non-occurrence), but as conceptualization, mental representation, imagination, or consideration of its possibility (e.g., Dekel *et al.* 1998; Heifetz *et al.* 2006; Modica and Rustichini 1999). But our Olympic sprinter may have experienced the overwhelming feeling of strength for the first time. All his past experiences may have been different in kind or intensity, so that he could not have imagined such a feeling before. He lacked not only *knowledge* but also *awareness* of the event. Arguably, many real-life belief changes – notably the more radical ones – involve the ‘observation’ or ‘experience’ of something that was previously not just unknown, but even beyond awareness or imagination. A Bayesian modelling of such belief changes involves an unnatural ascription of subjective probabilities to events beyond the agent’s awareness.

In sum, the modeller faces a choice between (i) ascribing *simple* Bayesian revision of *sophisticated* beliefs and (ii) ascribing *more complex* non-Bayesian revision of *simpler* beliefs. This choice is not just a matter of taste. The two alternatives are not merely different ways of saying the same thing, but different models of genuinely different phenomena, with distinct behavioural implications.⁴

Our paper and the literature. We analyse four salient belief revision rules, namely the above-mentioned Bayesian, Jeffrey, and Adams rules, and what we will call the dual-Jeffrey rule, which is a simpler variant of the Adams rule and which stands out for its duality to Jeffrey revision.⁵ In searching for axiomatic foundations for the first

³Even under a pure ‘as if’ interpretation of ascribed beliefs, highly sophisticated beliefs are dubious given the complexity of their behavioural implications (which may be hard to test empirically).

⁴An exact characterization of the behavioural differences between Bayesian and other belief revision models is beyond the scope of this paper.

⁵These four revision rules are of course not the only possible methods of belief revision; the literature contains several alternatives. Many of them depart from our assumption that beliefs are given by probability measures; see in particular (revision within) (i) the theory of Dempster-Shafer belief functions (e.g., Dempster 1967, Shafer 1976, Fagin and Halpern 1991a, Halpern 2003), (ii) theories with general non-additive probabilities (e.g., Schmeidler 1989, Wakker 1989, 2001, 2010, Sarin and Wakker 1994), (iii) theories of beliefs as *sets of* probability measures (e.g., Gilboa and Schmeidler 1989, Fagin and Halpern 1991b, Grove and Halpern 1998), and (iv) the theory of case-based qualitative beliefs (e.g., Gilboa and Schmeidler 2001). The theory of opinion pooling (e.g., Hylland and Zeckhauser 1979, McConway 1981, Genest *et al.* 1986, Genest and Zidek 1986, Dietrich 2010) is also sometimes interpreted as a theory of belief revision, by assuming that the agent learns

three rules, the literature has focused on a ‘distance-based’ approach. This consists in showing that a given revision rule is a *minimal* revision rule, which generates new beliefs that deviate as little as possible from initial beliefs, subject to certain constraints (given by the learning experience) and relative to some notion of ‘distance’ between beliefs (probability measures). In Bayesian revision, the constraint is that a particular event is assigned probability 1; in Jeffrey or Adams revision, it is that a particular random variable acquires a given distribution or conditional distribution. Bayesian and Jeffrey revision have been characterized as minimal revision relative to either the *variation* distance (defined by the maximal absolute difference in probability, over all events in the algebra), or the *Hellinger* distance, or the *relative entropy* distance (e.g., Csiszar 1967, 1977, van Fraassen 1981, Diaconis and Zabell 1982, Grunwald and Halpern 2003). The third notion of distance does not define a proper metric, as it is asymmetric in its two arguments. Douven and Romeijn (2012) have recently characterized Adams revision as minimal revision relative to yet another measure of distance, the *inverse relative entropy* distance (which differs from ordinary relative entropy distance in the inverted order of its arguments).

As elegant as these characterization results may be, they convey a non-unified picture of belief revision and a sense of arbitrariness. Different notions of distance are invoked to justify different revision rules, and their interpretation and relative advantages are controversial. We propose novel axiomatic foundations, which are not distance-based and lead to a unified axiomatic characterization of all four revision rules. In essence, we replace the requirement of distance-minimization from initial beliefs with the requirement of *conservativeness*, i.e., the preservation of those parts of a belief state (specific beliefs) on which the learning experience is ‘silent’. While the distance-based approach suggests that different revision rules differ in their underlying notions of distance, our main theorem shows that the four rules follow from the same underlying requirement of conservativeness. The real difference between the four rules consists in the learning experience prompting the belief change, not in the agent’s way of responding to it.⁶

2 Four revision rules in a single framework

A general framework for studying attitude revision (or more broadly, change in an agent’s state) can be obtained by specifying (i) a set \mathcal{P} of possible *states* in which the agent can be, and (ii) a set \mathcal{E} of possible (*learning*) *experiences* which can influence that state (see also Dietrich 2012). A *revision rule* is a function that maps pairs (p, E) of an initial state p in \mathcal{P} and an experience E in \mathcal{E} to a new state $p' = p_E$ in \mathcal{P} . Here, the pair (p, E) belongs to some domain $\mathcal{D} \subseteq \mathcal{P} \times \mathcal{E}$ containing those state-experience pairs that are admissible under the given revision rule. The revision rule is thus a

information in the form of opinions of other agents and merges them with his own initial opinion using an opinion pooling operator.

⁶Our conservativeness-based approach can be related to the rigidity-based approach (see Jeffrey 1957 for Bayesian and Jeffrey revision, and Bradley 2005 for Adams revision). For instance, Bayesian revision is *rigid* in the sense of preserving the conditional probability of any event given the learned event. The rigidity-based approach is so far not unified. One may interpret our conservativeness condition as a unified rigidity condition, applicable to any belief revision rule.

function from \mathcal{D} to \mathcal{P} .

Since we focus on *belief* revision, states are subjective probability functions. Specifically, we consider a fixed, non-empty set Ω of *worlds* which for expositional simplicity is finite or countably infinite.⁷ Subsets of Ω are called *events*. Let \mathcal{P} be the set of probability measures over Ω , i.e., countably additive functions $p : 2^\Omega \rightarrow [0, 1]$ with $p(\Omega) = 1$. We call any p in \mathcal{P} a *belief state*. The complement of any event $A \subseteq \Omega$ is denoted \overline{A} ($= \Omega \setminus A$), and $p(\omega)$ is an abbreviation for $p(\{\omega\})$. By a *partition*, we mean a partition of Ω into finitely many non-empty events. The *support* of a belief p is $\text{Supp}(p) := \{\omega \in \Omega : p(\omega) \neq 0\}$.

Before defining ‘experiences’, we consider informally the four revision rules to be studied. Suppose the agent is initially in belief state p in \mathcal{P} .

Bayesian revision: The agent learns some event B (with $p(B) \neq 0$) and adopts the new belief state p' given by

$$p'(A) = p(A|B) \text{ for all events } A \subseteq \Omega. \quad (1)$$

Jeffrey revision: The agent learns a new probability π_B for each event B in some partition \mathcal{B} (while keeping his conditional probabilities given B). He thus adopts the new belief state p' given by

$$p'(A) = \sum_{B \in \mathcal{B}} p(A|B) \pi_B \text{ for all events } A \subseteq \Omega. \quad (2)$$

The family of learned probabilities, $(\pi_B) \equiv (\pi_B)_{B \in \mathcal{B}}$, is assumed to be a *probability distribution on \mathcal{B}* , i.e., to consist of non-negative numbers with sum-total one.⁸ Often $|\mathcal{B}| = 2$. For instance, if the agent learns that it will rain with probability $\frac{1}{3}$, then partition \mathcal{B} contains the events of rain (B) and no rain (\overline{B}), where $\pi_B = \frac{1}{3}$ and $\pi_{\overline{B}} = \frac{2}{3}$. Jeffrey revision generalizes Bayesian revision since \mathcal{B} can contain a set B for which $\pi_B = 1$.

Dual-Jeffrey revision: The agent learns a new conditional probability function given any event C from some partition \mathcal{C} ; i.e., he learns that, given C , each event A has probability $\pi^C(A)$ (without learning a new probability of C). He thus adopts the new belief state p' given by

$$p'(A) = \sum_{C \in \mathcal{C}} \pi^C(A) p(C) \text{ for all events } A \subseteq \Omega. \quad (3)$$

The family $(\pi^C) \equiv (\pi^C)^{C \in \mathcal{C}}$ ($\in \mathcal{P}^{\mathcal{C}}$) is assumed to be a *conditional probability distribution given \mathcal{C}* , i.e., to consist of belief states $\pi^C \in \mathcal{P}$ with support $\text{Supp}(\pi^C) = C$. Often $|\mathcal{C}| = 2$. For instance, the agent might learn new distributions given the event C of a ‘rainy’ weather forecast and the event \overline{C} of a ‘dry’ forecast, so that $\mathcal{C} = \{C, \overline{C}\}$.

⁷Everything we say could be generalized to an arbitrary (measurable) set Ω .

⁸The revised belief state p' is only defined under the condition that no event B in \mathcal{B} has zero initial belief $p(B)$ but non-zero learnt probability π_B . This ensures that whenever in expression (2) a term $p(A|B)$ is undefined (because $p(B) = 0$) then this term does not matter (because it is multiplied by $\pi_B = 0$).

Dual-Jeffrey revision also captures the simple scenario of learning a new distribution given just one event, say the event of the ‘rainy’ forecast, without learning a new distribution given the ‘dry’ forecast. Here, \mathcal{C} contains the event C of a ‘rainy’ forecast and all the trivial singleton events $\{\omega\}$, where $\omega \in \overline{C}$, π^C is the newly learned conditional belief given C , and each $\pi^{\{\omega\}}$ is trivially given by $\pi^{\{\omega\}}(\omega) = 1$. The duality between Jeffrey and dual-Jeffrey revision consists in the fact that, if $\mathcal{B} = \mathcal{C}$, the two forms of revision concern complementary parts of the agent’s belief state: while the former affects probabilities of events in \mathcal{B} and leaves probabilities given these events unaffected, the latter does the reverse.⁹

Adams revision: The agent learns a new conditional probability π_B^C of any event B from a first partition \mathcal{B} given any event C from a second partition \mathcal{C} (without learning a new probability of C or new conditional probabilities given $B \cap C$). He thus adopts the new belief state p' given by

$$p'(A) = \sum_{B \in \mathcal{B}, C \in \mathcal{C}} p(A|B \cap C) \pi_B^C p(C) \text{ for all } A \subseteq \Omega. \quad (4)$$

The family $(\pi_B^C) \equiv (\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ is assumed to be a *conditional probability distribution on \mathcal{B} given \mathcal{C}* , i.e., a family of numbers indexed by both \mathcal{B} and \mathcal{C} such that $\sum_{B \in \mathcal{B}} \pi_B^C = 1$ for all $C \in \mathcal{C}$ and such that $\pi_B^C > (=) 0$ whenever $B \cap C \neq (=) \emptyset$.¹⁰ Often $|\mathcal{B}| = |\mathcal{C}| = 2$. For instance, if the agent learns that it will rain with probability $\frac{9}{10}$ given a ‘rainy’ forecast and with probability $\frac{3}{10}$ given a ‘dry’ forecast, then partition \mathcal{B} contains the events of rain (B) and no rain (\overline{B}), and partition \mathcal{C} contains the events of a ‘rainy’ forecast (C) and a ‘dry’ forecast (\overline{C}), where $\pi_B^C = \frac{9}{10}$, $\pi_{\overline{B}}^C = \frac{1}{10}$, $\pi_B^{\overline{C}} = \frac{3}{10}$, and $\pi_{\overline{B}}^{\overline{C}} = \frac{7}{10}$. To represent the scenario in which the agent learns only a single conditional probability, say only the new probability of rain given the ‘rainy’ forecast, one could define \mathcal{B} as containing the events of rain (B) and no rain (\overline{B}) and define \mathcal{C} as containing the event C of a ‘rainy’ forecast and trivial singleton events $\{\omega\}$ for all $\omega \in \overline{C}$, where we still have $\pi_B^C = \frac{9}{10}$ and $\pi_{\overline{B}}^C = \frac{1}{10}$ and where any $\pi_{B'}^{\{\omega\}}$ ($B' \in \mathcal{B}$) takes the trivial value of 1 if $\omega \in B'$ and 0 if $\omega \notin B'$. Adams revision generalizes dual-Jeffrey revision, which is obtained if \mathcal{B} is the finest partition $\{\{a\} : a \in \Omega\}$. It also ‘almost’ generalizes Jeffrey revision, since if \mathcal{C} is the coarsest partition $\{\Omega\}$ we obtain Jeffrey revision with family $(\pi_B)_{B \in \mathcal{B}} \equiv (\pi_B^\Omega)_{B \in \mathcal{B}}$, where this Jeffrey revision is not of the most general kind since each $\pi_B (= \pi_B^\Omega)$ is non-zero.

To give formal definitions of these four revision rules, we must first define the notion of a learning experience. Looking at Bayesian revision alone, one may be tempted

⁹ As a consequence, any new belief \hat{p} (with full support Ω) can be acquired in two steps: a Jeffrey revision step of learning the new probability $\pi_B = \hat{p}(B)$ of each event $B \in \mathcal{B}$, and a dual-Jeffrey revision step of learning the new conditional probability function $\pi^B = \hat{p}(\cdot|B)$ for each event $B \in \mathcal{B}$. In other words, revision towards \hat{p} is the composition of a Jeffrey revision and a dual-Jeffrey revision, in any order.

¹⁰ The revised belief p' is only defined under the condition that $p(B \cap C) \neq 0$ for all $B \in \mathcal{B}$ and $C \in \mathcal{C}$ such that $B \cap C \neq \emptyset$ and $p(C) \neq 0$. This condition ensures that in expression (4) the term $p(A|B \cap C)$ is defined whenever it matters, i.e., whenever the term $\pi_B^C p(C)$ with which it is multiplied is non-zero.

to define experiences as observed events $B \subseteq \Omega$. But the other three revision rules are based on mathematical objects distinct from events, namely families of probabilities (or probability functions) of the forms (π_B) , (π^C) , and (π_B^C) . Methodologically, one should not tie the notion of an experience (i.e., the definition of \mathcal{E}) to a particular kind of mathematical object that is tailor-made for a specific revision rule. Such a notion would not only exclude other revision rules from the framework, but also prevent one from giving a fully convincing axiomatic justification for the revision rule in question: key features of that rule would already have been built into the definitions themselves.

We thus need an abstract notion of a learning experience. We define an experience simply as a set of belief states $E \subseteq \mathcal{P}$, representing the constraint that the agent's revised belief state must belong to E . So, the set of logically possible experiences is $\mathcal{E} = 2^{\mathcal{P}}$ (note that this is deliberately general). An agent's belief change from p to p_E upon learning $E \in \mathcal{E}$ is *responsive to the experience* if $p_E \in E$. Our four revision rules involve the following experiences:

Definition 1 A (learning) experience $E \subseteq \mathcal{P}$ is

- **Bayesian** if $E = \{p' : p'(B) = 1\}$ for some (learned) event $B \neq \emptyset$;
- **Jeffrey** if $E = \{p' : p'(B) = \pi_B \ \forall B \in \mathcal{B}\}$ for some (learned) probability distribution $(\pi_B)_{B \in \mathcal{B}}$ on some partition \mathcal{B} ;
- **dual-Jeffrey** if $E = \{p' : p'(\cdot|C) = \pi^C \ \forall C \in \mathcal{C} \text{ such that } p'(C) \neq 0\}$ for some (learned) conditional probability distribution $(\pi^C)^{C \in \mathcal{C}}$ given some partition \mathcal{C} ;
- **Adams** if $E = \{p' : p'(B|C) = \pi_B^C \ \forall B \in \mathcal{B} \ \forall C \in \mathcal{C} \text{ such that } p'(C) \neq 0\}$ for some (learned) conditional probability distribution $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ on some partition \mathcal{B} given some partition \mathcal{C} .

Every Bayesian experience is a Jeffrey experience and every dual-Jeffrey and ‘almost’ every Jeffrey experience is an Adams experience (see the earlier remarks for details). Some experiences are of none of these kinds, such as the experience $E = \{p' : p'(A \cap B) > p'(A)p'(B)\}$ that two given events A and B are positively correlated, the experience $E = \{p' : p'(A) \geq 9/10\}$ that A is very probable, and so on. In general, the smaller the set E , the stronger (more constraining) the experience. The strongest logically consistent experiences are the singleton sets $E = \{p'\}$, which require adopting the new belief state p' regardless of the initial belief state. The logically weakest experience is the set $E = \mathcal{P}$, which allows the agent to retain his old belief state.

We can now formally define the four revision rules.

Definition 2 *Bayesian (respectively Jeffrey, dual-Jeffrey, Adams) revision is the revision rule $(p, E) \mapsto p' = p_E$ given by formula (1) (respectively (2), (3), (4)) and defined on the domain $\mathcal{D}_{\text{Bayes}}$ (respectively $\mathcal{D}_{\text{Jeffrey}}$, $\mathcal{D}_{\text{dual-Jeffrey}}$, $\mathcal{D}_{\text{Adams}}$) consisting of all belief-experience pairs $(p, E) \in \mathcal{P} \times 2^{\mathcal{P}}$ such that E is a Bayesian (respectively Jeffrey, dual-Jeffrey, Adams) experience compatible with p (i.e., for which expression (1) (respectively (2), (3), (4)) is defined¹¹).*

¹¹Recall that expression (1) is defined under the condition that $p(B) \neq 0$, expression (2) under the condition stated in footnote 8, expression (3) always, and expression (4) under the condition stated in footnote 10.

Jeffrey revision extends Bayesian revision, i.e., it coincides with Bayesian revision on the subdomain $\mathcal{D}_{\text{Bayes}} (\subseteq \mathcal{D}_{\text{Jeffrey}})$. Similarly, Adams revision extends dual-Jeffrey revision. The definition of each revision rule makes implicit use of the fact that the mathematical object entering the revision formula (and entering the definition of the rule's domain) – i.e., the event B or family (π_B) , (π^C) or (π_B^C) – is uniquely determined by the set E , or is at least determined to the extent needed for uniqueness of the revised belief state (and the domain).¹²

As mentioned, the literature has focused on characterizing Bayesian, Jeffrey, and Adams revision as minimizing the distance between the revised belief state and the initial one, subject to constraint given by the learning experience. Formally, a revision rule $(p, E) \mapsto p_E$ on a domain $\mathcal{D} \subseteq \mathcal{P} \times 2^{\mathcal{P}}$ is *distance-minimizing* with respect to distance function $d : 2^{\mathcal{P}} \times 2^{\mathcal{P}} \rightarrow \mathbb{R}$ if, for every $(p, E) \in \mathcal{D}$, $p' = p_E$ minimizes $d(p', p)$ subject to $p' \in E$. Different distance functions, however, have been used for characterizing each of the different revision rules. We aim at a different, more unified characterization.

3 A unified axiomatic characterization

We now introduce two plausible conditions on belief revision and show that these force the agent to revise his beliefs in a Bayesian way in response to any Bayesian experience, in a Jeffrey way in response to any Jeffrey experience, in a dual-Jeffrey way in response to any dual-Jeffrey experience, and in an Adams way in response to any Adams experience. This shows that the four seemingly different revision rules follow from the same two principles and differ only in the kinds of experience for which they are defined.

The first principle is that the revised belief state should be responsive to the learning experience, i.e., respect the constraint posed by it.

Responsiveness: $p_E \in E$ for all belief-experience pairs $(p, E) \in \mathcal{D}$.

Responsiveness guarantees, for instance, that in response to a Bayesian experience the learned event is assigned probability one.

The second principle is a natural conservativeness requirement: those ‘parts’ of the agent’s belief state on which the experience is ‘silent’ should not change in response to it. In short, the experience should have no effect where it has nothing to say. To define the principle formally, we must answer two questions: what do we mean by

¹²In the case of Bayesian revision, the event B is uniquely determined by the (Bayesian) experience E . Similarly, in the case of dual-Jeffrey revision the family (π^C) is uniquely determined. In the case of Jeffrey revision, the family $(\pi_B)_{B \in \mathcal{B}}$ is *essentially* uniquely determined, in the sense that the subpartition $\{B \in \mathcal{B} : \pi_B \neq 0\}$ and the corresponding subfamily $(\pi_B)_{B \in \mathcal{B} : \pi_B \neq 0}$ are unique. (The subpartition $\{B \in \mathcal{B} : \pi_B = 0\}$ is sometimes non-unique. Uniqueness can be achieved by imposing the convention that $|\{B \in \mathcal{B} : \pi_B = 0\}| \leq 1$.) Jeffrey revision is well-defined because the revision formula (2) (and the definition of the domain $\mathcal{D}_{\text{Jeffrey}}$) only depend on this subfamily. As for Adams revision, the family (π_B^C) is far from uniquely determined by the Adams experience E , but Adams revision is nonetheless well-defined because the revision formula (4) (and the domain definition) are invariant under the choice of family (π_B^C) representing E . This non-trivial fact is shown in the appendix, where we also give a characterization of the families (π_B^C) representing a given Adams experience.

‘parts of a belief state’, and when is an experience ‘silent’ on them? To answer these questions, note that, intuitively:

- a Bayesian experience of learning an event B is not silent on the probability of B , but silent on the conditional probabilities given B ;
- a Jeffrey experience is not silent on the probabilities of the events in the relevant partition \mathcal{B} , but silent on the conditional probabilities given these events;
- a dual-Jeffrey experience is silent on the unconditional probability of the events in the relevant partition;
- an Adams experience is silent on the unconditional probability of the events in the second partition as well as on conditional probabilities given events from the join of the two partitions.

So, the parts of the agent’s belief state on which these experiences are intuitively silent are conditional probabilities of certain events A given certain other events B (where possibly $B = \Omega$). These conditional probabilities are preserved by the corresponding revision rule, so that the rule is intuitively conservative.

We now define formally when an experience E is silent on the probability of an event A given another B . We need to define silence only for the case that $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(E)$, where $\text{Supp}(E)$ is the *support* of E , defined as $\cup_{p' \in E} \text{Supp}(p')$ ($= \{\omega : p'(\omega) \neq 0 \text{ for some } p' \in E\}$).

There are in fact two plausible notions of ‘silence’, and hence of ‘conservativeness’. We begin with the weaker notion of silence. An experience is *weakly silent* on the probability of A given B if it permits this conditional probability to take any value. Formally:

Definition 3 *Experience E is **weakly silent on the probability of A given B** (for $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(E)$) if, for every value α in $[0, 1]$, E contains some belief state p' (with $p'(B) \neq 0$) such that $p'(A|B) = \alpha$.*

For instance, the experience $E = \{p' : p'(B) = 1/2\}$ is weakly silent on the probability of A given B . So is the experience $E = \{p' : p'(A) \leq 1/2\}$. This weak notion of silence gives rise to the following strong notion of conservativeness:

Strong Conservativeness: For all belief-experience pairs $(p, E) \in \mathcal{D}$, if E is weakly silent on the probability of an event A given another B (where $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(E)$), this conditional probability is preserved, i.e., $p_E(A|B) = p(A|B)$ (provided $p_E(B), p(B) \neq 0$).

Although strong conservativeness may look like a plausible requirement, it leads to an impossibility result.

Proposition 1 *If $\#\Omega \geq 3$, no belief revision rule on a domain $\mathcal{D} \supseteq \mathcal{D}_{\text{Jeffrey}}$ is both responsive and strongly conservative.*

Note that, on the small domain $\mathcal{D}_{\text{Bayes}}$, the impossibility does not hold, because Bayesian revision is responsive as well as strongly conservative. On that domain, the present strong conservativeness condition is no stronger than the later, weaker one, which we now introduce to avoid the impossibility more generally (so, the impossibility occurs on domains on which the two conservativeness conditions come apart).

We weaken strong conservativeness by strengthening the underlying notion of silence. The key insight is that even if an experience E is weakly silent on the probability of A given B , it may still implicitly ‘say’ something about how this conditional probability should relate to other probability assignments within the agent’s belief state. Suppose for instance $\Omega = \{0, 1\}^2$, where the first component of a world $(g, j) \in \Omega$ represents whether Peter has gone out ($g = 1$) or not ($g = 0$), and the second whether Peter is wearing a jacket ($j = 1$) or not ($j = 0$). Consider the events that Peter has gone out $G = \{(g, j) : g = 1\}$ and that he is wearing a jacket $J = \{(g, j) : j = 1\}$. Some experiences are weakly silent on the probability of J (given Ω) and yet require this probability to be related in certain ways to other probability assignments, notably probabilities conditional on J . Consider for instance the (Jeffrey) experience that G is 90% probable, $E = \{p' : p'(G) = 0.9\}$. It is compatible with any probability of J and thus weakly silent on the probability of J given Ω . But it requires this probability to be related in certain ways to the probability of G given J : if this conditional probability is 1 (which is compatible with E), then the probability of J can no longer exceed 0.9, because, if it did, the probability of G would exceed 0.9, in contradiction with the experience E . In short, although E does not directly constrain the belief on J , it constrains this belief *indirectly*, i.e., after other parts of the agent’s belief state have been fixed.

An experience is *strongly silent* on a probability of A given B if it permits this conditional probability to take any value even after other parts of the agent’s belief state have been fixed. Let us first explain this idea informally. What exactly are the ‘other parts of the agent’s belief state’? They are those probability assignments that are ‘orthogonal’ to the probability of A given B . Expressed differently, they are all the beliefs of which the belief state p' is made up, *over and above the probability of A given B* . More precisely, assuming again that A is included in B , they can be captured by the quadruple consisting of the unconditional probability $p'(B)$ and the conditional probabilities $p'(\cdot|A)$, $p'(\cdot|B \setminus A)$, and $p'(\cdot|\bar{B})$.¹³ This quadruple and the conditional probability $p'(A|B)$ jointly determine the belief state p' , because

$$p' = \underbrace{p'(\cdot|A) p'(A)}_{\alpha p'(B)} + \underbrace{p'(\cdot|B \setminus A) p'(B \setminus A)}_{p'(B) - \alpha p'(B)} + \underbrace{p'(\cdot|\bar{B}) p'(\bar{B})}_{1 - p'(B)}.$$

If an experience E is strongly silent on the conditional probability of A given B , then this probability can be chosen freely even after the other parts of the agent’s belief state have been fixed in accordance with E (which requires them to match those of some belief state p^* in E). This idea is illustrated in Figure 1, where an experience E is represented in the space whose first coordinate represents the probability of A given B and whose second (multi-dimensional) coordinate represents the other parts of the agent’s belief state.

To define strong silence formally, we say that two belief states p' and p^* *coincide outside the probability of A given B* if the other parts of these belief states coincide, i.e., if $p'(B) = p^*(B)$ and $p'(\cdot|C) = p^*(\cdot|C)$ for all $C \in \{A, B \setminus A, \bar{B}\}$ such that $p'(C), p^*(C) \neq 0$. Clearly, two belief states that coincide both (i) outside the probability of A given B and (ii) on the probability of A given B are identical.

¹³This informal discussion assumes that $p'(A), p'(B \setminus A), p'(\bar{B}) \neq 0$.




Figure 1: Experiences E which are (a) not even weakly silent, (b) weakly silent, or (c) strongly silent on the probability of A given B , respectively.

Definition 4 *Experience $E \subseteq \mathcal{P}$ is **strongly silent on the probability of A given B** (for $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(E)$) if, for all $\alpha \in [0, 1]$ and all $p^* \in E$, E contains some belief state p' (with $p'(B) \neq 0$) which*

- (a) *coincides with α on the probability of A given B , i.e., $p'(A|B) = \alpha$,*
- (b) *coincides with p^* outside the probability of A given B (if $p^*(A), p^*(B \setminus A) \neq 0$).*

In this definition, there is only one belief state p' satisfying (a) and (b), given by

$$p' := p^*(\cdot|A)\alpha p^*(B) + p^*(\cdot|B \setminus A)(1 - \alpha)p^*(B) + p^*(\cdot \cap \overline{B}), \quad (5)$$

so that the requirement that there exists some p' in E satisfying (a) and (b) reduces to the requirement that E contains the belief (5).¹⁴

For example, the experiences $E = \{p' : p' \text{ is uniform on } \overline{B}\}$ and $E = \{p' : p'(B) \geq 1/2\}$ are strongly silent on the probability of A given B , since this conditional probability can take any value independently of other parts of the agent's belief state (e.g., independently of the probability of B).

There is an alternative and equivalent way of defining weak and strong silence, which gives a different perspective on these notions. Informally, on this alternative approach, weak silence means that the experience implies nothing for the probability of A given B , whereas strong silence means that it implies only something outside the probability of A given B , i.e., for those parts of the agent's belief state that are orthogonal to the probability of A given B . To state the alternative definitions formally, we first define the 'implication' of an experience for the probability of A given B and for other parts of the agent's belief state (where $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(E)$):

- The **implication of E for the probability of A given B** is the experience, denoted $E_{A|B}$, which says everything that E says about the probability of A given B , but nothing else (see Figure 2). So, $E_{A|B}$ contains all belief states p' which are compatible with E on the probability of A given B . Formally, $E_{A|B}$ is the set of all belief states p' such that $p'(A|B) = p^*(A|B)$ for some p^* in E (more precisely, such that if $p'(B) \neq 0$ then $p'(A|B) = p^*(A|B)$ for some $p^* \in E$ satisfying $p^*(B) \neq 0$).

¹⁴To be precise, this is true whenever $p^*(A), p^*(B \setminus A) \neq 0$.




Figure 2: The experiences $E_{A|B}$ and $E_{\overline{A|B}}$ derived from an experience E

- The **implication of E outside the probability of A given B** is the experience, denoted $E_{\overline{A|B}}$, which says everything that E says outside the probability of A given B , but nothing else (see Figure 2). So, $E_{\overline{A|B}}$ contains all belief states which are compatible with E outside the probability of A given B . Formally, $E_{\overline{A|B}}$ is the set of all belief states p' which outside the probability of A given B coincide with some belief state in E (more precisely, with some belief state p^* in E satisfying the non-triviality condition

$$p^*(C) \neq 0 \text{ for all } C \in \{A, B \setminus A\} \text{ such that } p'(C) \neq 0$$

if at least one belief state in E satisfies this condition).

Clearly, $E \subseteq E_{A|B}$ and $E \subseteq E_{\overline{A|B}}$. The experiences $E_{A|B}$ and $E_{\overline{A|B}}$ capture two orthogonal components ('subexperiences') of the full experience E . Each component reflects what E has to say on a particular aspect. Weak and strong silence can now be characterized by the following salient properties, which constitute the announced alternative definitions:

Proposition 2 *For all experiences $E \subseteq \mathcal{P}$ and events A and B (where $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(E)$),*

- E is weakly silent on the probability of A given B if and only if $E_{A|B} = \mathcal{P}$ (i.e., E implies nothing for the probability of A given B),*
- E is strongly silent on the probability of A given B if and only if $E_{\overline{A|B}} = E$ (i.e., E implies only something outside the probability of A given B).*

An intuition for this result is obtained by combining Figures 1 and 2. By part (a), weak silence means that the subexperience $E_{A|B}$ about the probability of A given B is vacuous; graphically, it covers the entire area in the plot. By part (b), strong silence means that the experience E contains no more information than its subexperience $E_{\overline{A|B}}$ about the parts of the agent's belief state that are orthogonal to the probability

of A given B ; graphically, E covers a rectangular area reaching from the very left to the very right.

This strengthened notion of silence leads to a weaker notion of conservativeness, to be called just ‘conservativeness’. This condition is defined exactly like strong conservativeness except that ‘weak silence’ is replaced by ‘strong silence’:

Conservativeness: For all belief-experience pairs $(p, E) \in D$, if E is strongly silent on the probability of an event A given another B (for $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(E)$), this conditional probability is preserved, i.e., $p_E(A|B) = p(A|B)$ (if $p_E(B), p(B) \neq 0$).

This weaker condition does not lead to an impossibility result, but to a characterization of our four revision rules:

Theorem 1 *Bayesian, Jeffrey, dual-Jeffrey and Adams revision are the only responsive and conservative belief revision rules on their respective domains.*

Corollary 1 *Every responsive and conservative revision rule on an arbitrary domain $\mathcal{D} \subseteq \mathcal{P} \times 2^{\mathcal{P}}$ coincides with Bayesian (respectively Jeffrey, dual-Jeffrey, Adams) revision on the intersection of \mathcal{D} with $\mathcal{D}_{\text{Bayes}}$ (respectively $\mathcal{D}_{\text{Jeffrey}}$, $\mathcal{D}_{\text{dual-Jeffrey}}$, $\mathcal{D}_{\text{Adams}}$).*

It is easier to prove that *if* a revision rule on the domain of one of these four revisions rules is responsive and conservative, then it must be that classic revision rule, than to prove the converse implication that each of these four rules is in fact responsive and conservative on its domain. For instance, if a belief-experience pair (p, E) belongs to $\mathcal{D}_{\text{Bayes}}$, say $E = \{p' : p'(B) = 1\}$, then the new belief state p_E equals $p_E(\cdot|B)$ (as $p_E(B) = 1$ by responsiveness), which equals $p(\cdot|B)$ (by conservativeness, as E is strongly silent on probabilities given B). The reason why the converse implication is harder to prove is that, for each of the four kinds of experience, it is non-trivial to identify all the conditional probabilities on which this experience is strongly silent. There are more such conditional probabilities than one might expect. For example, a dual-Jeffrey experience is strongly silent *not only* on the unconditional probabilities of events in the relevant partition, but also on a number of other probabilities, as detailed in the Appendix. After having identified all the conditional probabilities on which an experience of each kind is strongly silent, one must verify that the corresponding revision rule does indeed preserve all these probabilities, as required by conservativeness.

4 Conclusion

We have shown that four salient belief revision rules follow from the same two basic principles: responsiveness to the learning experience and conservativeness. The only difference between the four rules lies in the kind of learning experience that is admitted by each of them. This characterization contrasts with known characterizations of Bayesian, Jeffrey, and Adams revision as distance-minimizing rules with respect to *different* distance functions between probability measures.

Our two principles can guide belief revision not just in the face of a learning experience of one of the four kinds we have discussed. They constitute a general

recipe for belief revision. An important question for future research is *how far* the principles can take us. Can they deal with completely different learning experiences, such as learning that the probability of rain exceeds the square root of the probability of a thunder storm? This question has two parts. First, for which learning experiences is responsive and conservative belief revision possible at all? Secondly, when is belief revision in accordance with these principles unique? Another challenge is to extend the conservativeness-based approach towards the revision of belief states distinct from probability measures, such as Dempster-Shafer belief functions, general non-additive probability measures, or sets of probability measures.

5 References

- Bradley, R. (2005) Radical Probabilism and Bayesian Conditioning, *Philosophy of Science* 72: 342-364
- Bradley, R. (2007) The Kinematics of Belief and Desire, *Synthese* 56(3): 513-535
- Csiszar, I. (1967) Information type measures of difference of probability distributions and indirect observations, *Studia Scientiarum Mathematicarum Hungarica* 2: 299-318
- Csiszar, I. (1977) Information Measures: A Critical Survey, *Transactions of the Seventh Prague Conference*: 73-86
- Dekel, E., Lipman, B., Rustichini, A. (1998) Standard state-space models preclude unawareness, *Econometrica* 66(1): 159-174
- Dempster, A. P. (1967) Upper and lower probabilities induced by a multi-valued mapping, *Annals of Mathematical Statistics* 38: 325-399
- Diaconis, P., Zabell, S. (1982) Updating subjective probability, *Journal of the American Statistical Association* 77: 822-830
- Dietrich, F. (2010) Bayesian group belief, *Social Choice and Welfare* 35(4): 595-626
- Dietrich, F. (2012) Modelling change in individual characteristics: an axiomatic approach, *Games and Economic Behavior*, in press
- Douven, I., Romeijn, J. W. (2012) A new resolution of the Judy Benjamin Problem, *Mind*, in press
- Fagin, R., Halpern, J. Y. (1991a) A new approach to updating beliefs, *Uncertainty in Artificial Intelligence* 6 (Bonissone *et al.* (eds.), Elsevier Science Publishers)
- Fagin, R., Halpern, J. Y. (1991b), Uncertainty, belief, and probability, *Computational Intelligence* 7: 160-173
- Genest, C., McConway, K. J., Schervish, M. J. (1986) Characterization of externally Bayesian pooling operators, *Annals of Statistics* 14, 487-501
- Genest, C., Zidek, J. V. (1986) Combining probability distributions: a critique and an annotated bibliography, *Statist. Sci.* 1: 114-148
- Gilboa, I., Schmeidler, D. (1989) Maximin expected utility with a non-unique prior, *Journal of Mathematical Economics* 18: 141-53
- Gilboa, I., Schmeidler, D. (2001) *A Theory of Case-Based Decisions*, Cambridge University Press
- Grove, A., Halpern, J. (1998) Updating Sets of Probabilities. In: D. Poole *et al.* (eds.) *Proceedings of the 14th Conference on Uncertainty in AI*, Morgan

- Kaufmann, Madison, WI, USA, 173-182
- Grunwald, P., Halpern, J. (2003) Updating probabilities, *Journal of AI Research* 19: 243-78
- Halpern, J. (2003) *Reasoning About Uncertainty*, MIT Press, Cambridge, MA, USA
- Heifetz, A., Meier, M. and B. C. Schipper (2006). Interactive unawareness, *Journal of Economic Theory*, 130, 78-94.
- Hylland, A., Zeckhauser, R. (1979) The impossibility of group decision making with separate aggregation of beliefs and values, *Econometrica* 47: 1321-36
- Jeffrey, R. (1957) Contributions to the theory of inductive probability, *PhD Thesis*, Princeton University
- McConway, K. (1981) Marginalization and linear opinion pools, *Journal of the American Statistical Association* 76: 410-414
- Modica, S., Rustichini, A. (1999) Unawareness and partitional information structures, *Games and Economic Behavior* 27: 265-298
- Sarin, R., Wakker, P. (1994) A General Result for Quantifying Beliefs, *Econometrica* 62, 683-685
- Schmeidler, D. (1989) Subjective probability and expected utility without additivity, *Econometrica* 57: 571-87
- Shafer, G. (1976) *A Mathematical Theory of Evidence*, Princeton University Press
- Shafer, G. (1981) Jeffrey's rule of conditioning, *Philosophy of Science* 48: 337-62
- van Fraassen, B. C. (1981) A Problem for Relative Information Minimizers in Probability Kinematics, *British Journal for the Philosophy of Science* 32: 375-379
- Wakker, P. (1989) Continuous Subjective Expected Utility with Nonadditive Probabilities, *Journal of Mathematical Economics* 18: 1-27
- Wakker, P. (2001) Testing and Characterizing Properties of Nonadditive Measures through Violations of the Sure-Thing Principle, *Econometrica* 69: 1039-59
- Wakker, P. (2010) *Prospect Theory: For Risk and Ambiguity*, Cambridge University Press

A Appendix: proofs

Notation in proofs: For all $a \in \Omega$, let $\delta_a \in \mathcal{P}$ be the Dirac measure in a , defined by $\delta_a(a) = 1$. For every non-empty event $A \subseteq \Omega$, let $\text{unifom}_A \in \mathcal{P}$ be the uniform probability measure on A , defined by $\text{unifom}_A(B) = \frac{|B \cap A|}{|A|}$ for all $B \subseteq \Omega$.

A.1 Well-definedness of each revision rule

As mentioned, our four revision rules (i.e., Bayesian, Jeffrey, dual-Jeffrey and Adams revision) have been well-defined because the mathematical object used in the definition of the new belief state (and of the rule's domain) – i.e., the learned event B respectively the learned family (π_B) , (π^C) or (π_B^C) – is uniquely determined by the relevant experience E , or is at least *sufficiently* determined so that the definition does not depend on any underdetermined features. This fact deserves a proof. For the first three revision rules, the proof is trivial and given by the following three lemmas (which the reader can easily show):

Lemma 1 *Every Bayesian experience is generated by exactly one event $B \subseteq E$.*

Lemma 2 *Every dual-Jeffrey experience is generated by exactly one family $(\pi^C)_{C \in \mathcal{C}}$.*

Lemma 3 *For every Jeffrey experience E ,*

- (a) *all families $(\pi_B)_{B \in \mathcal{B}}$ generating E have the same subfamily $(\pi_B)_{B \in \mathcal{B}: \pi_B \neq 0}$ (especially, the same set $\{B \in \mathcal{B} : \pi_B \neq 0\}$);*
- (b) *in particular, for every (initial) belief state $p \in \mathcal{P}$, the (revised) belief state (2) is either defined and the same for all families $(\pi_B)_{B \in \mathcal{B}}$ generating E , or undefined for all these families.¹⁵*

Well-definedness of Adams revision is harder to establish. We start by a lemma which characterizes the common features of all families $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ generating the same given Adams experience E . On a first reading of the lemma, one might assume that no $C \in \mathcal{C}$ is included in any $B \in \mathcal{B}$ (so that $\mathcal{C}_{\text{triv}} = \emptyset$). In this case, the lemma implies that all these families share the same partition \mathcal{C} and the same join of partition $\mathcal{B} \vee \mathcal{C} = \{B \cap C : B \in \mathcal{B}, C \in \mathcal{C}\} \setminus \{\emptyset\}$. The sets $C \in \mathcal{B}$ which are included in some $B \in \mathcal{B}$ are special because any value π_B^C ($B \in \mathcal{B}$) is then trivially one (if $C \subseteq B$) or zero (if $B \cap C = \emptyset$).

Lemma 4 *Let E be an Adams experience. All families $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ generating E have*

- (a) *the same set $\mathcal{C} \setminus \mathcal{C}_{\text{triv}}$, where $\mathcal{C}_{\text{triv}} := \{C \in \mathcal{C} : \exists B \in \mathcal{B} \text{ such that } C \subseteq B\}$,*
- (b) *the same set $(\mathcal{B} \vee \mathcal{C}) \setminus \mathcal{C}_{\text{triv}}$, where $\mathcal{C}_{\text{triv}}$ is defined as in part (a),*
- (c) *for each $a \in \Omega$ the same value $\pi_{B_a}^{C_a}$, where B_a (resp. C_a) denotes the member of \mathcal{B} (resp. \mathcal{C}) which contains a .*

Proof. Consider an Adams experience E . The proof consists of showing several claims about an arbitrary family $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ generating E . Claims 5, 7 and 8 complete the proofs of parts (a), (b) and (c), respectively. For each $a \in \Omega$ let B_a (resp. C_a , D_a) denote the set in \mathcal{B} (resp. \mathcal{C} , $\mathcal{B} \vee \mathcal{C}$) containing a . Note that $D_a = B_a \cap C_a$ for all $a \in \Omega$.

Our strategy is to show that the sets $\mathcal{C} \setminus \mathcal{C}_{\text{triv}}$ and $(\mathcal{B} \vee \mathcal{C}) \setminus \mathcal{C}_{\text{triv}}$ and the values $\pi_{B_a}^{C_a}$ ($a \in \Omega$) can be defined in terms of E alone rather than in terms of the family $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ generating E , which shows independence from the choice of family. We first prove that several other objects – such as in Claim 1 the number $|\{B \in \mathcal{B} \vee \mathcal{C} : B \subseteq C_a\}|$ and in Claim 2 the set $C_a \setminus D_a$ (where $a \in \Omega$) – can be defined in terms of E alone.

Claim 1: For each $a \in \Omega$, $|\{B \in \mathcal{B} \vee \mathcal{C} : B \subseteq C_a\}| = \min_{p' \in E: p'(a) \neq 0} |\text{Supp}(p')|$.

Let $a \in \Omega$. To show that $\min_{p' \in E: p'(a) \neq 0} |\text{Supp}(p')| \geq |\{B \in \mathcal{B} \vee \mathcal{C} : B \subseteq C_a\}|$, consider any $p' \in E$ such that $p'(a) \neq 0$. It suffices to consider any $B \in \mathcal{B} \vee \mathcal{C}$ such that $B \subseteq C_a$ and show that $p'(B) \neq 0$. If $a \in B$ the latter is evident since $p'(a) \neq 0$. Now let $a \notin B$. Since $B \in \mathcal{B} \vee \mathcal{C}$ and $B \subseteq C_a$ we have $B = B' \cap C_a$ for some $B' \in \mathcal{B}$. Noting that $p' \in E$ and $p'(C_a) \neq 0$, we have $p'(B'|C_a) = \pi_{B'}^{C_a} \neq 0$; so, $p'(B' \cap C_a) \neq 0$, i.e., $p'(B) \neq 0$.

To show the converse inequality,

$$\min_{p' \in E: p'(a) \neq 0} |\text{Supp}(p')| \leq |\{B \in \mathcal{B} \vee \mathcal{C} : B \subseteq C_a\}|,$$

¹⁵Footnote 8 specifies when (2) is defined.

note that one can find a $p' \in E$ with $p'(a) \neq 0$ such that

$$|Supp(p')| = |\{B \in \mathcal{B} \vee \mathcal{C} : B \subseteq C_a\}|,$$

namely by picking an element a_B from each set B in $\{B \in \mathcal{B} \vee \mathcal{C} : B \subseteq C_a\}$, where $a_{D_a} = a$, and defining p' as the unique probability function in \mathcal{P} such that $Supp(p') = \{a_B : B \in \mathcal{B} \vee \mathcal{C} : B \subseteq C_a\}$ and $p'(a_B|C_a) = \pi_{B'}^{C_a}$ for all $B \in \mathcal{B} \vee \mathcal{C}$ such that $B \subseteq C_a$ (where B' again stands for the set in \mathcal{B} such that $B = B' \cap C_a$). Q.e.d.

In the rest of this proof, for all $a \in \Omega$ we let E^a be the set of all $p' \in E$ such that $Supp(p')$ is minimal (w.r.t. set inclusion) subject to $p'(a) \neq 0$.

Claim 2: For all $a \in \Omega$, $C_a \setminus D_a = (\cup_{p' \in E^a} Supp(p')) \setminus \{a\}$.

Let $a \in \Omega$. The claim follows from the fact that, as the reader may verify, E^a is the set of all $p' \in \mathcal{P}$ such that for every $B \in \mathcal{B} \vee \mathcal{C}$ included in C_a there is an $a_B \in B$ such that (i) $a_{D_a} = a$, (ii) $Supp(p') = \{a_B : B \in \mathcal{B} \vee \mathcal{C}, B \subseteq C_a\}$ (hence, $p'(C_a) = 1$), and (iii) $p'(a_B) = \pi_{B'}^{C_a}$ (i.e., $p'(a_B|C_a) = \pi_{B'}^{C_a}$) for all $B \in \mathcal{B} \vee \mathcal{C}$ included in C_a , where B' again stands for the set in \mathcal{B} for which $B = B' \cap C_a$. Q.e.d.

Claim 3: For all $a \in \Omega$, the following are equivalent: (i) $D_a = C_a$, (ii) $C_a \subseteq B_a$, (iii) $\delta_a \in E$, and (iv) $E^a = \{\delta_a\}$.

For all $a \in \Omega$, (i) is equivalent to (ii) since $D_a = B_a \cap C_a$; (ii) is clearly equivalent to (iii); and (iii) is equivalent to (iv) by definition of E^a . Q.e.d.

In the following, for each $a \in \Omega$ such that $D_a \neq C_a$, i.e., such that $C_a \not\subseteq B_a$, let $c(a)$ be a fixed element of $C_a \setminus D_a$.

Claim 4: For all $a \in \Omega$ such that $\delta_a \notin E$ (i.e., such that $D_a \neq C_a$ by Claim 3), $C_a = \cup_{p' \in E^a \cup E^{c(a)}} Supp(p')$.

Consider $a \in \Omega$ such that $\delta_a \notin E$, i.e., by Claim 3 such that $D_a \neq C_a$. Note that $C_{c(a)} = C_a$ and that $D_{c(a)}$ and D_a are non-empty disjoint subsets of $C_a (= C_{c(a)})$. We may write C_a as

$$C_a = (C_a \setminus D_a) \cup (C_a \setminus D_{c(a)}).$$

So, by Claim 2 applied to a and to $c(a)$,

$$C_a = [(\cup_{p' \in E^a} Supp(p')) \setminus \{a\}] \cup [(\cup_{p' \in E^{c(a)}} Supp(p')) \setminus \{c(a)\}].$$

Since

$$c(a) \in (\cup_{p' \in E^a} Supp(p')) \setminus \{a\} \text{ and } a \in (\cup_{p' \in E^{c(a)}} Supp(p')) \setminus \{c(a)\},$$

it follows that

$$\begin{aligned} C_a &= (\cup_{p' \in E^a} Supp(p')) \cup (\cup_{p' \in E^{c(a)}} Supp(p')) \\ &= \cup_{p' \in E^a \cup E^{c(a)}} Supp(p'). \text{ Q.e.d.} \end{aligned}$$

Claim 5: We have

$$\mathcal{C} \setminus \mathcal{C}_{\text{triv}} = \left\{ \cup_{p' \in E^a \cup E^{c(a)}} Supp(p') : a \in \Omega, \delta_a \notin E \right\}$$

(which proves part (a) since $\mathcal{C} \setminus \mathcal{C}_{\text{triv}}$ depends on E alone rather than on the particular family (π_B^C)).

Note that $\mathcal{C} = \{C_a : a \in \Omega\}$ and $\mathcal{C}_{\text{triv}} = \{C_a : a \in \Omega, D_a = C_a\}$. So,

$$\mathcal{C} \setminus \mathcal{C}_{\text{triv}} = \{C_a : a \in \Omega, D_a \neq C_a\}.$$

This implies the claim by Claim 4. Q.e.d.

Claim 6: For all $a \in \Omega$ such that $\delta_a \notin E$ (i.e., such that $D_a \neq C_a$ by Claim 3),

$$D_a = \left[\bigcup_{p' \in E^{c(a)}} \text{Supp}(p') \right] \setminus \left[\bigcup_{p' \in E^a} \text{Supp}(p') \setminus \{a\} \right].$$

Consider any $a \in \Omega$ such that $\delta_a \notin E$. We have $D_a = C_a \setminus (C_a \setminus D_a)$. Hence, using the expressions for C_a and $C_a \setminus D_a$ found in Claims 4 and 2,

$$D_a = \left[\bigcup_{p' \in E^a \cup E^{c(a)}} \text{Supp}(p') \right] \setminus \left[\bigcup_{p' \in E^a} \text{Supp}(p') \setminus \{a\} \right].$$

It is clear that we can replace ' $E^a \cup E^{c(a)}$ ' by ' $E^{c(a)}$ ' without changing the resulting set D_a . Q.e.d.

Claim 7: We have

$$(\mathcal{B} \vee \mathcal{C}) \setminus \mathcal{C}_{\text{triv}} = \left\{ \left[\bigcup_{p' \in E^{c(a)}} \text{Supp}(p') \right] \setminus \left[\bigcup_{p' \in E^a} \text{Supp}(p') \setminus \{a\} \right] : a \in \Omega, \delta_a \notin E \right\}$$

(which proves part (b) since $(\mathcal{B} \vee \mathcal{C}) \setminus \mathcal{C}_{\text{triv}}$ depends on E alone rather than on the particular family (π_B^C)).

Since $\mathcal{B} \vee \mathcal{C} = \{D_a : a \in \Omega\}$ and $\mathcal{C}_{\text{triv}} = \{D_a : a \in \Omega, D_a = C_a\}$, we have

$$(\mathcal{B} \vee \mathcal{C}) \setminus \mathcal{C}_{\text{triv}} = \{D_a : a \in \Omega, D_a \neq C_a\}.$$

The claim now follows from Claim 6. Q.e.d.

Claim 8: Part (c) of the lemma holds.

Let $a \in \Omega$. Consider any other family $(\tilde{\pi}_{\tilde{C}}^{\tilde{C}})_{\tilde{B} \in \tilde{\mathcal{B}}}$ also generating E , define \tilde{B}_a (resp. \tilde{C}_a, \tilde{D}_a) as the set in $\tilde{\mathcal{B}}$ (resp. $\tilde{\mathcal{C}}, \tilde{\mathcal{B}} \vee \tilde{\mathcal{C}}$) containing a , and define $\tilde{\mathcal{C}}_{\text{triv}}$ as $\{C \in \tilde{\mathcal{C}} : C \subseteq B \text{ for some } B \in \tilde{\mathcal{B}}\}$. We have to show that $\pi_{B_a}^{C_a} = \tilde{\pi}_{\tilde{B}_a}^{\tilde{C}_a}$. By parts (a) and (b) (which we proved in Claims 5 and 7),

$$\mathcal{C} \setminus \mathcal{C}_{\text{triv}} = \tilde{\mathcal{C}} \setminus \tilde{\mathcal{C}}_{\text{triv}} \tag{6}$$

$$(\mathcal{B} \vee \mathcal{C}) \setminus \mathcal{C}_{\text{triv}} = (\tilde{\mathcal{B}} \vee \tilde{\mathcal{C}}) \setminus \tilde{\mathcal{C}}_{\text{triv}}. \tag{7}$$

By (6) we have $\bigcup_{C \in \mathcal{C} \setminus \mathcal{C}_{\text{triv}}} C = \bigcup_{C \in \tilde{\mathcal{C}} \setminus \tilde{\mathcal{C}}_{\text{triv}}} C$. So, taking complements in Ω on both sides,

$$\bigcup_{C \in \mathcal{C}_{\text{triv}}} C = \bigcup_{C \in \tilde{\mathcal{C}}_{\text{triv}}} C. \tag{8}$$

We distinguish between two cases.

Case 1: a belongs to a set in $\mathcal{C}_{\text{triv}}$, or equivalently by (8), a set in $\tilde{\mathcal{C}}_{\text{triv}}$. Since a belongs to a set in $\mathcal{C}_{\text{triv}}$, we have $C_a \subseteq B_a$, whence $\pi_{B_a}^{C_a} = 1$. Similarly, since a belongs to a set in $\tilde{\mathcal{C}}_{\text{triv}}$, we have $\tilde{C}_a \subseteq \tilde{B}_a$, whence $\tilde{\pi}_{\tilde{B}_a}^{\tilde{C}_a} = 1$. So, $\pi_{B_a}^{C_a} = \tilde{\pi}_{\tilde{B}_a}^{\tilde{C}_a} (= 1)$.

Case 2: a does not belong to a set in $\mathcal{C}_{\text{triv}}$, or equivalently, a set in $\tilde{\mathcal{C}}_{\text{triv}}$. We deduce firstly, using (6), that a belongs to a set in $\mathcal{C} \setminus \mathcal{C}_{\text{triv}} = \tilde{\mathcal{C}} \setminus \tilde{\mathcal{C}}_{\text{triv}}$, so that $C_a = \tilde{C}_a$;

and secondly, using (7), that a belongs to a set in $(\mathcal{B} \vee \mathcal{C}) \setminus \mathcal{C}_{\text{triv}} = (\tilde{\mathcal{B}} \vee \tilde{\mathcal{C}}) \setminus \tilde{\mathcal{C}}_{\text{triv}}$, so that $D_a = \tilde{D}_a$. Choose any p' in E such that $p'(C_a) \neq 0$ (of course there is such a p' in E). Then, as the families (π_B^C) and $(\tilde{\pi}_{\tilde{B}}^{\tilde{C}})$ both generate E , we have $p'(B_a|C_a) = \pi_{B_a}^{C_a}$ and $p'(\tilde{B}_a|\tilde{C}_a) = \tilde{\pi}_{\tilde{B}_a}^{\tilde{C}_a}$. So, it suffices to show that $p'(B_a|C_a) = p'(\tilde{B}_a|\tilde{C}_a)$, i.e., that $p'(B_a \cap C_a)/p'(C_a) = p'(\tilde{B}_a \cap \tilde{C}_a)/p'(\tilde{C}_a)$, or equivalently, that $p'(D_a)/p'(C_a) = p'(\tilde{D}_a)/p'(\tilde{C}_a)$. This holds because $D_a = \tilde{D}_a$ and $C_a = \tilde{C}_a$. ■

Among the families representing a given Adams experience E , one stands out as canonical, as the next lemma shows.

Lemma 5 *Let E be an Adams experience. Among all families $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ generating E , there is exactly one ('canonical') one such that*

- (a) \mathcal{B} refines \mathcal{C} (i.e., each C in \mathcal{C} is a union of one or more sets in \mathcal{B}),
- (b) $|\mathcal{B} \cap \mathcal{C}| \leq 1$.

Condition (a) on the family – more precisely, on the partitions \mathcal{B} and \mathcal{C} – is the key requirement; essentially, it requires a fine choice of \mathcal{B} . Starting from an arbitrary family $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ generating E , one can ensure condition (a) by refining \mathcal{B} , i.e., replacing each $B \in \mathcal{B}$ by all non-empty set(s) of the form $B \cap C$ where $C \in \mathcal{C}$. Condition (b) is no more than a convention to avoid trivial redundancies. Any set $B \in \mathcal{B} \cap \mathcal{C}$ leads to the trivial value $\pi_B^B = 1$. It suffices to have at most one such set, since if there are many sets in $\mathcal{B} \cap \mathcal{C}$ then they can be replaced by their union. We have just given an intuition for the lemma's existence claim. The uniqueness claim will be proved using Lemma 4.

Proof. Let E be an Adams experience.

1. In this part we prove existence of a family which generates E and has the two properties (a) and (b). Let $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ be any family generating E , i.e.,

$$E = \{p' : p'(B|C) = \pi_B^C \ \forall B \in \mathcal{B} \ \forall C \in \mathcal{C} \text{ such that } p'(C) \neq 0\}. \quad (9)$$

We now define a new family $(\hat{\pi}_B^C)_{B \in \hat{\mathcal{B}}}^{C \in \hat{\mathcal{C}}}$, of which we later show that it generates the same experience E and has the two required properties that $\hat{\mathcal{B}}$ refines $\hat{\mathcal{C}}$ and $|\hat{\mathcal{B}} \cap \hat{\mathcal{C}}| \leq 1$.

Consider the 'trivial' part of the partitions \mathcal{B} and \mathcal{C} , defined as $\mathcal{C}_{\text{triv}} := \{C \in \mathcal{C} : C \subseteq B \text{ for some } B \in \mathcal{B}\}$. The partition $\hat{\mathcal{C}}$ is defined as \mathcal{C} if $\mathcal{C}_{\text{triv}} = \emptyset$, while otherwise it is defined from \mathcal{C} by replacing the trivial part by a single set:

$$\hat{\mathcal{C}} := \begin{cases} \mathcal{C} & \text{if } \mathcal{C}_{\text{triv}} = \emptyset \\ (\mathcal{C} \setminus \mathcal{C}_{\text{triv}}) \cup \{\cup_{C' \in \mathcal{C}_{\text{triv}}} C'\} & \text{if } \mathcal{C}_{\text{triv}} \neq \emptyset. \end{cases}$$

The partition $\hat{\mathcal{B}}$ is defined as the join of \mathcal{B} and \mathcal{C} if $\mathcal{C}_{\text{triv}} = \emptyset$, and otherwise it is derived from this join by replacing the trivial part by a single set:

$$\hat{\mathcal{B}} := \begin{cases} \mathcal{B} \vee \mathcal{C} & \text{if } \mathcal{C}_{\text{triv}} = \emptyset \\ ((\mathcal{B} \vee \mathcal{C}) \setminus \mathcal{C}_{\text{triv}}) \cup \{\cup_{C' \in \mathcal{C}_{\text{triv}}} C'\} & \text{if } \mathcal{C}_{\text{triv}} \neq \emptyset. \end{cases}$$

Finally, for all $B \in \widehat{\mathcal{B}}$ and $C \in \widehat{\mathcal{C}}$, define

$$\widehat{\pi}_B^C := \begin{cases} \pi_{B'}^C & \text{if } B \subsetneq C \text{ (so that } C \in \mathcal{C} \setminus \mathcal{C}_{\text{triv}}), \text{ where } B' \text{ is the set in } \mathcal{B} \text{ s.t. } B \subseteq B' \\ 1 & \text{if } B = C \text{ (so that } B = C = \cup_{C' \in \mathcal{C}_{\text{triv}}} C') \\ 0 & \text{if } B \cap C = \emptyset. \end{cases}$$

Note that the three mentioned cases – i.e., $B \subsetneq C$, $B = C$ and $B \cap C = \emptyset$ – are the only possible ones since $\widehat{\mathcal{B}}$ refines $\widehat{\mathcal{C}}$.

We now show that the so-defined family $(\widehat{\pi}_B^C)_{B \in \widehat{\mathcal{B}}}^{C \in \widehat{\mathcal{C}}}$ has the required properties. Clearly, $\widehat{\mathcal{B}}$ refines $\widehat{\mathcal{C}}$, and $|\widehat{\mathcal{B}} \cap \widehat{\mathcal{C}}| \leq 1$ since $\widehat{\mathcal{B}} \cap \widehat{\mathcal{C}}$ is empty (if $\mathcal{C}_{\text{triv}} = \emptyset$) or $\{\cup_{C' \in \mathcal{T}} C'\}$ (if $\mathcal{C}_{\text{triv}} \neq \emptyset$). It remains to show that $(\widehat{\pi}_B^C)_{B \in \widehat{\mathcal{B}}}^{C \in \widehat{\mathcal{C}}}$ generates E , i.e., that the sets (9) and

$$\widehat{E} := \{p' : p'(B|C) = \widehat{\pi}_B^C \forall B \in \widehat{\mathcal{B}} \forall C \in \widehat{\mathcal{C}} \text{ such that } p'(C) \neq 0\}.$$

coincide.

First, let $p' \in E$. To show that $p' \in \widehat{E}$, consider any $B \in \widehat{\mathcal{B}}$ and $C \in \widehat{\mathcal{C}}$ such that $p'(C) \neq 0$; we have to prove that $p'(B|C) = \widehat{\pi}_B^C$. We distinguish three cases:

- If $B \subsetneq C$, then $p'(B|C) = \widehat{\pi}_B^C$ since $p'(B|C)$ and $\widehat{\pi}_B^C$ both equal $\pi_{B'}^C$, where B' denotes the set in \mathcal{B} such that $B \subseteq B'$, i.e., such that $B = B' \cap C$. To see why $p'(B|C) = \pi_{B'}^C$, note that $p'(B|C)$ equals $p'(B'|C)$, which in turn equals $\pi_{B'}^C$, as $p' \in E$.
- If $B = C$, then $p'(B|C) = \widehat{\pi}_B^C$ since $p'(B|C) = 1$ and $\widehat{\pi}_B^C = 1$.
- If $B \cap C = \emptyset$, then $p'(B|C) = \widehat{\pi}_B^C$ since $p'(B|C) = 0$ and $\widehat{\pi}_B^C = 0$.

Conversely, let $p' \in \widehat{E}$. To show that $p' \in E$, consider any $B \in \mathcal{B}$ and $C \in \mathcal{C}$ such that $p'(C) \neq 0$. We prove $p'(B|C) = \pi_B^C$ by again distinguishing three cases:

- If $C \setminus B, C \cap B \neq \emptyset$, then $p'(B|C) = \pi_B^C$ because $p'(B|C)$ and π_B^C both equal $\widehat{\pi}_{B'}^C$, where $B' := B \cap C \in \widehat{\mathcal{B}}$. To see why $p'(B|C) = \widehat{\pi}_{B'}^C$, note that $p'(B|C)$ equals $p'(B'|C)$, which in turn equals $\widehat{\pi}_{B'}^C$ as $p' \in \widehat{E}$.
- If $C \setminus B = \emptyset$ (i.e., $C \subseteq B$), then $p'(B|C) = \pi_B^C$ since $p'(B|C) = 1$ and $\pi_B^C = 1$.
- If $B \cap C = \emptyset$, then $p'(B|C) = \pi_B^C$ since $p'(B|C) = 0$ and $\pi_B^C = 0$.

2. In this part we prove the uniqueness claim. Let $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ and $(\widetilde{\pi}_B^C)_{B \in \widetilde{\mathcal{B}}}^{C \in \widetilde{\mathcal{C}}}$ be two such families. Define

$$\begin{aligned} \mathcal{C}_{\text{triv}} &\equiv \{C \in \mathcal{C} : C \subseteq B \text{ for some } B \in \mathcal{B}\} = \mathcal{B} \cap \mathcal{C} \\ \widetilde{\mathcal{C}}_{\text{triv}} &\equiv \{C \in \widetilde{\mathcal{C}} : C \subseteq B \text{ for some } B \in \widetilde{\mathcal{B}}\} = \widetilde{\mathcal{B}} \cap \widetilde{\mathcal{C}}, \end{aligned}$$

where the equalities on these two lines hold because \mathcal{B} refines \mathcal{C} and $\widetilde{\mathcal{B}}$ refines $\widetilde{\mathcal{C}}$. By Lemma 4,

$$\mathcal{C} \setminus \mathcal{C}_{\text{triv}} = \widetilde{\mathcal{C}} \setminus \widetilde{\mathcal{C}}_{\text{triv}}, \quad (10)$$

$$(\mathcal{B} \vee \mathcal{C}) \setminus \mathcal{C}_{\text{triv}} = (\widetilde{\mathcal{B}} \vee \widetilde{\mathcal{C}}) \setminus \widetilde{\mathcal{C}}_{\text{triv}}, \quad (11)$$

$$\pi_{B_a}^{C_a} = \widetilde{\pi}_{\widetilde{B}_a}^{\widetilde{C}_a} \text{ for all } a \in \Omega, \quad (12)$$

where for each $a \in \Omega$ the set B_a (resp. $C_a, \widetilde{B}_a, \widetilde{C}_a$) denotes the member of \mathcal{B} (resp. $\mathcal{C}, \widetilde{\mathcal{B}}, \widetilde{\mathcal{C}}$) which contains a . Since \mathcal{B} refines \mathcal{C} and $\widetilde{\mathcal{B}}$ refines $\widetilde{\mathcal{C}}$ we have $\mathcal{B} \vee \mathcal{C} = \mathcal{B}$ and $\widetilde{\mathcal{B}} \vee \widetilde{\mathcal{C}} = \widetilde{\mathcal{B}}$, so that equation (11) reduces to

$$\mathcal{B} \setminus \mathcal{C}_{\text{triv}} = \widetilde{\mathcal{B}} \setminus \widetilde{\mathcal{C}}_{\text{triv}}. \quad (13)$$

Further, from (10) and the fact that \mathcal{C} and $\tilde{\mathcal{C}}$ are partitions of Ω and that each of the sets $\mathcal{C}_{\text{triv}}$ ($= \mathcal{B} \cap \mathcal{C}$) and $\tilde{\mathcal{C}}_{\text{triv}}$ ($= \tilde{\mathcal{B}} \cap \tilde{\mathcal{C}}$) contains at most one member one can deduce that $\mathcal{C}_{\text{triv}} = \tilde{\mathcal{C}}_{\text{triv}}$, which together with equations (10) and (13) implies that

$$\mathcal{C} = \tilde{\mathcal{C}} \text{ and } \mathcal{B} = \tilde{\mathcal{B}}. \quad (14)$$

It remains to prove that $\pi_B^C = \tilde{\pi}_B^C$ for all $B \in \mathcal{B}$ ($= \tilde{\mathcal{B}}$) and $C \in \mathcal{C}$ ($= \tilde{\mathcal{C}}$). Consider any $B \in \mathcal{B}$ ($= \tilde{\mathcal{B}}$) and $C \in \mathcal{C}$ ($= \tilde{\mathcal{C}}$). If $B \cap C = \emptyset$ then $\pi_B^C = 0$ and $\tilde{\pi}_B^C = 0$, whence $\pi_B^C = \tilde{\pi}_B^C$, as required. Now assume $B \cap C \neq \emptyset$. Choose any $a \in B \cap C$. Since $a \in B \in \mathcal{B} = \tilde{\mathcal{B}}$ we have $B_a = \tilde{B}_a = B$, and similarly, since $a \in C \in \mathcal{C} = \tilde{\mathcal{C}}$ we have $C_a = \tilde{C}_a = C$. So, using (12), $\pi_B^C = \tilde{\pi}_B^C$. ■

We are now ready to prove that Adams revision has been well-defined.

Lemma 6 *For every Adams experience E and every (initial) belief state $p \in \mathcal{P}$, the (revised) belief state (4) is either defined and the same for all families $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ generating E , or undefined for all these families.¹⁶*

Proof. Let E be an Adams experience and $p \in \mathcal{P}$. We write Π for the set of families $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ generating E .

Claim 1: Expression (4) is defined for either every or no family in Π .

Consider two families $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ and $(\tilde{\pi}_B^{\tilde{C}})_{B \in \tilde{\mathcal{B}}}^{\tilde{C} \in \tilde{\mathcal{C}}}$ in Π . By footnote 10 we have to show that

$$[B \cap C \neq \emptyset \& p(C) \neq 0] \Rightarrow p(B \cap C) \neq 0 \text{ for all } B \in \mathcal{B}, C \in \mathcal{C} \quad (15)$$

if and only if

$$[\tilde{B} \cap \tilde{C} \neq \emptyset \& p(\tilde{C}) \neq 0] \Rightarrow p(\tilde{B} \cap \tilde{C}) \neq 0 \text{ for all } \tilde{B} \in \tilde{\mathcal{B}}, \tilde{C} \in \tilde{\mathcal{C}}. \quad (16)$$

We assume (15) and show (16); the converse implication holds analogously. To show (16), consider any $\tilde{B} \in \tilde{\mathcal{B}}$ and $\tilde{C} \in \tilde{\mathcal{C}}$ such that $\tilde{B} \cap \tilde{C} \neq \emptyset$ and $p(\tilde{C}) \neq 0$. We have to show that $p(\tilde{B} \cap \tilde{C}) \neq 0$. We suppose w.l.o.g. that $\tilde{C} \not\subseteq \tilde{B}$, since otherwise trivially $p(\tilde{B} \cap \tilde{C}) = p(\tilde{C}) \neq 0$. Again let $\mathcal{C}_{\text{triv}}$ ($\tilde{\mathcal{C}}_{\text{triv}}$) be the set of sets in \mathcal{C} ($\tilde{\mathcal{C}}$) included in a set in \mathcal{B} ($\tilde{\mathcal{B}}$). As $\tilde{C} \not\subseteq \tilde{B}$ and $\tilde{B} \cap \tilde{C} \neq \emptyset$, we have $\tilde{C} \notin \tilde{\mathcal{C}}_{\text{triv}}$. So, since by Lemma 4 $\mathcal{C} \setminus \mathcal{C}_{\text{triv}} = \tilde{\mathcal{C}} \setminus \tilde{\mathcal{C}}_{\text{triv}}$, we have $\tilde{C} \in \mathcal{C}$. Moreover, since $\tilde{\mathcal{C}}_{\text{triv}}$ does not contain \tilde{C} , it also does not contain any subset of \tilde{C} , so that $\tilde{B} \cap \tilde{C} \notin \tilde{\mathcal{C}}_{\text{triv}}$. Hence, $\tilde{B} \cap \tilde{C} \in (\tilde{\mathcal{B}} \vee \tilde{\mathcal{C}}) \setminus \tilde{\mathcal{C}}_{\text{triv}}$. As by Lemma 4 $(\mathcal{B} \vee \mathcal{C}) \setminus \mathcal{C}_{\text{triv}} = (\tilde{\mathcal{B}} \vee \tilde{\mathcal{C}}) \setminus \tilde{\mathcal{C}}_{\text{triv}}$, it follows that $\tilde{B} \cap \tilde{C} \in \mathcal{B} \vee \mathcal{C}$. Thus there exist (unique) $B \in \mathcal{B}$ and $C \in \mathcal{C}$ such that $\tilde{B} \cap \tilde{C} = B \cap C$. Since $\tilde{C} \in \mathcal{C}$ we have $C = \tilde{C}$. Using that $p(C) = p(\tilde{C}) \neq 0$ and that $B \cap C = \tilde{B} \cap \tilde{C} \neq \emptyset$, we have $p(B \cap C) \neq 0$ by (15), i.e., $p(\tilde{B} \cap \tilde{C}) \neq 0$. Q.e.d.

Claim 2: The revised belief state (4) is the same for all families $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ in Π for which it is defined.

Let $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ and $(\tilde{\pi}_B^{\tilde{C}})_{B \in \tilde{\mathcal{B}}}^{\tilde{C} \in \tilde{\mathcal{C}}}$ be two families in Π for which the revised belief state is defined. We write p' and \tilde{p}' for the corresponding new belief states, respectively.

¹⁶Footnote 10 specifies when (4) is defined.

To show that $p' = \tilde{p}'$, we consider a fixed $a \in \Omega$ and show that $p'(a) = \tilde{p}'(a)$. Note that

$$p'(a) = p(a|B_a \cap C_a) \pi_{B_a}^{C_a} p(C_a), \quad (17)$$

$$\tilde{p}'(a) = p(a|\hat{B}_a \cap \hat{C}_a) \hat{\pi}_{\hat{B}_a}^{\hat{C}_a} p(\hat{C}_a), \quad (18)$$

where B_a (resp. C_a , \hat{B}_a , \hat{C}_a) denotes the element of \mathcal{B} (resp. \mathcal{C} , $\hat{\mathcal{B}}$, $\hat{\mathcal{C}}$) which contains a . By Lemma 4, we have $\mathcal{C} \setminus \mathcal{C}_{\text{triv}} = \hat{\mathcal{C}} \setminus \hat{\mathcal{C}}_{\text{triv}}$, where $\mathcal{C}_{\text{triv}} := \{C \in \mathcal{C} : C \subseteq B \text{ for some } B \in \mathcal{B}\}$ and $\hat{\mathcal{C}}_{\text{triv}} := \{\hat{C} \in \hat{\mathcal{C}} : \hat{C} \subseteq \hat{B} \text{ for some } \hat{B} \in \hat{\mathcal{B}}\}$. So, $\cup_{C \in \mathcal{C} \setminus \mathcal{C}_{\text{triv}}} C = \cup_{\hat{C} \in \hat{\mathcal{C}} \setminus \hat{\mathcal{C}}_{\text{triv}}} \hat{C}$, and hence, taking complements on both sides,

$$\cup_{C \in \mathcal{C}_{\text{triv}}} C = \cup_{\hat{C} \in \hat{\mathcal{C}}_{\text{triv}}} \hat{C}. \quad (19)$$

We consider two cases.

Case 1: a does *not* belong to a set in $\mathcal{C}_{\text{triv}}$, or equivalently by (19) a set in $\hat{\mathcal{C}}_{\text{triv}}$. By parts (a), (b) and (c) of Lemma 4 we therefore have $C_a = \hat{C}_a$, $B_a \cap C_a = \hat{B}_a \cap \hat{C}_a$ and $\pi_{B_a}^{C_a} = \hat{\pi}_{\hat{B}_a}^{\hat{C}_a}$, respectively. So, equations (17) and (18) imply that $p'(a) = \tilde{p}'(a)$.

Case 2: a belongs to a set in $\mathcal{C}_{\text{triv}}$, or equivalently a set in $\hat{\mathcal{C}}_{\text{triv}}$. Then $C_a \subseteq B_a$ and $\hat{C}_a \subseteq \hat{B}_a$, whence $\pi_{B_a}^{C_a} = 1$ and $\hat{\pi}_{\hat{B}_a}^{\hat{C}_a} = 1$. So, equations (17) and (18) reduce to

$$\begin{aligned} p'(a) &= p(a|C_a)p(C_a) = p(a), \\ \tilde{p}'(a) &= p(a|\hat{C}_a)p(\hat{C}_a) = p(a). \end{aligned}$$

Hence, $p'(a) = \tilde{p}'(a)$. ■

A.2 Proposition 1

Proof of Proposition 1. Suppose that $\#\Omega \geq 3$. For a contraction, consider a responsive and conservative revision rule on a domain $\mathcal{D} \supseteq \mathcal{D}_{\text{Jeffrey}}$. As $\#\Omega \geq 3$ there are events $A, B \subseteq \Omega$ such that $A \cap B, B \setminus A, A \setminus B \neq \emptyset$. Consider an initial belief state p such that $p(A \cap B) = 1/4$ and $p(A \setminus B) = 3/4$, and define the Jeffrey experience $E := \{p' : p'(B) = 1/2\}$. Note that $(p, E) \in \mathcal{D}$. What is the new belief state p_E ?

First note that E is weakly silent on the probability of $A \cap B$ given B . So, by Strong Conservativeness $p_E(A \cap B|B) = p(A \cap B|B)$ (using that $p(B) \neq 0$ and that $p_E(B) \neq 0$ by Responsiveness), i.e., (*) $p_E(A|B) = 1$.

Similarly, (**) $p_E(A|\bar{B}) = 1$. (This is trivial if $A \cap \bar{B} = \bar{B}$, and can otherwise be shown like (*), using this time that E is weakly silent on the probability of $A \cap \bar{B}$ given \bar{B} .) By (*) and (**), $p_E(A) = 1$.

Further, E is weakly silent on the probability of $A \cap B$ given A , so that by Strong Conservativeness $p_E(A \cap B|A) = p(A \cap B|A)$ (using that $p_E(A), p(A) \neq 0$). Given the fact that $p_E(A) = 1$ and the definition of p , it follows that $p_E(B) = 1/4$. But by Responsiveness $p_E(B) = 1/2$, a contradiction. ■

A.3 Proposition 2

We start by giving a convenient reformulation of strong silence (we leave the proof to the reader).

Lemma 7 *For all experiences E and all events $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(E)$, E is strongly silent on the probability of A given B if and only if E contains a p^* with $p^*(A), p^*(B \setminus A) \neq 0$ and for every such $p^* \in E$ and every $\alpha \in [0, 1]$ E contains the belief state p' which coincides with α on and with p^* outside the probability of A given B , i.e., the belief state*

$$p' := p^*(\cdot|A)\alpha p^*(B) + p^*(\cdot|B \setminus A)(1 - \alpha)p^*(B) + p^*(\cdot \cap \overline{B}).$$

Proof of Proposition 2. Consider $E \subseteq \mathcal{P}$ and $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(E)$.

(a) First suppose $E_{A|B} = \mathcal{P}$. Consider any $\alpha \in [0, 1]$. As $\emptyset \subsetneq A \subsetneq B$ there exists a belief state p' such that $p'(B) \neq 0$ and $p'(A|B) = \alpha$. As $E_{A|B} = \mathcal{P}$, we have that $p' \in E_{A|B}$, so that E contains a p^* (with $p^*(B) \neq 0$) such that $p^*(A|B) = p'(A|B)$, i.e., such that $p^*(A|B) = \alpha$, as required to establish weak silence.

Now assume E is weakly silent on the probability of A given B . Trivially $E_{A|B} \subseteq \mathcal{P}$; we show that $\mathcal{P} \subseteq E_{A|B}$. Let $p' \in \mathcal{P}$. If $p'(B) = 0$ then clearly $p' \in E_{A|B}$. Otherwise, by weak silence as applied to $\alpha := p'(A|B)$, E contains a p^* such that $p^*(B) \neq 0$ and $p^*(A|B) = p'(A|B)$, so that again $p' \in E_{A|B}$.

(b) First, in the (degenerate) case that E contains no p' such that $p'(A), p'(B \setminus A) \neq 0$, the equivalence holds because strong silence is violated (see Lemma 7) and moreover $E_{\overline{A|B}} \neq E$ because $E_{\overline{A|B}}$ but not E contains a belief state p' such that $p'(A), p'(B \setminus A) \neq 0$. Now assume the less trivial case that E contains a \tilde{p} such that $\tilde{p}(A), \tilde{p}(B \setminus A) \neq 0$.

First suppose $E_{\overline{A|B}} = E$. To show strong silence, consider any $\alpha \in [0, 1]$ and any $p^* \in E$ with $p^*(A), p^*(B \setminus A) \neq 0$. By Lemma 7 it suffices to show that the belief state p' which coincides with p^* outside the probability of A given B and satisfies $p'(A|B) = \alpha$ belongs to E . Clearly, p' belongs to $E_{\overline{A|B}}$. Hence, as $E = E_{\overline{A|B}}$, p' belongs to E .

Conversely, assume E is strongly silent on the probability of A given B . Trivially, $E \subseteq E_{\overline{A|B}}$. To show the converse inclusion, suppose $p' \in E_{\overline{A|B}}$. Then there is a $p^* \in E$ such that p' and p^* coincide outside the probability of A given B and such that $p^*(C) \neq 0$ for all $C \in \{A, B \setminus A\}$ with $p'(C) \neq 0$.

We distinguish two cases. First suppose $p^*(A), p^*(B \setminus A) \neq 0$. Then $p'(B) = p^*(B) \neq 0$. By E 's strong silence on the probability of A given B , E contains a belief state \tilde{p} (with $\tilde{p}(B) \neq 0$) which satisfies $\tilde{p}(A|B) = p'(A|B)$ and coincides with p^* outside the probability of A given B . Note that, since $p^*(A), p^*(B \setminus A) \neq 0$, there can be only one belief state that coincides with p^* outside the probability of A given B and such that the probability of A given B takes a given value. Therefore, $p' = \tilde{p}$, and so $p' \in E$, as had to be shown.

Next assume the special case that $p^*(C) = 0$ for at least one $C \in \{A, B \setminus A\}$. As $p^*(C) = 0 \Rightarrow p'(C) = 0$ for each $C \in \{A, B \setminus A\}$ and as $p'(A) + p'(B \setminus A) = p'(B) = p^*(B) = p^*(A) + p^*(B \setminus A)$, it follows that $p'(C) = p^*(C)$ for each $C \in \{A, B \setminus A, \overline{B}\}$. This and the fact that $p'(\cdot|C) = p^*(\cdot|C)$ for all $C \in \{A, B \setminus A, \overline{B}\}$ for which $p'(C) (= p^*(C))$ is non-zero imply that $p' = p^*$. So again $p' \in E$. ■

A.4 Characterization of where each kind of experience is strongly silent

As a step in establishing Theorem 1, this section determine where Bayesian, Jeffrey, dual-Jeffrey and Adams experiences are strongly silent. We do not treat Bayesian experiences explicitly and instead move directly to Jeffrey experiences, since the latter generalize the former.

Lemma 8 *For all Jeffrey experiences E (of learning a new probability distribution on a partition \mathcal{B}) and all events $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(E)$, E is strongly silent on the probability of A given B if and only if $B \subseteq B'$ for some $B' \in \mathcal{B}$.*

Proof. Let E , \mathcal{B} , A and B be as specified, and let $(\pi_B)_{B \in \mathcal{B}}$ be the learned probability distribution on \mathcal{B} . First, if $B \subseteq B'$ for some $B' \in \mathcal{B}$ then E is strongly silent on the probability of A given B , as one easily checks using Lemma 7. Conversely, suppose that $B \not\subseteq B'$ for all $B' \in \mathcal{B}$. For each $D \subseteq \Omega$ we write $\mathcal{B}_D := \{B' \in \mathcal{B} : B' \cap D \neq \emptyset\}$. Note that $\mathcal{B}_B = \mathcal{B}_A \cup \mathcal{B}_{B \setminus A}$, where $\#\mathcal{B}_A \geq 1$ (as $A \neq \emptyset$), $\#\mathcal{B}_{B \setminus A} \geq 1$ (as $B \setminus A \neq \emptyset$) and $\#\mathcal{B}_B \geq 2$ (as otherwise B would be included in a $B' \subseteq \mathcal{B}$). It follows that there are $B' \in \mathcal{B}_A$ and $B'' \in \mathcal{B}_{B \setminus A}$ with $B' \neq B''$. Note that E contains a p^* such that $p^*(B' \cap A) = \pi_{B'}$ and $p^*(B'' \cap (B \setminus A)) = \pi_{B''}$. Since each of B' and B'' has a non-empty intersection with B , and hence with $\text{Supp}(E) (\supseteq B)$, we have $\pi_{B'}, \pi_{B''} \neq 0$. Now $p^*(B'' \cap A) = p^*(B'' \cap \overline{B}) = 0$, since

$$p^*((B'' \cap A) \cup (B'' \cap \overline{B})) = p^*(B'') - p^*(B'' \cap (B \setminus A)) = \pi_{B''} - \pi_{B''} = 0.$$

By Lemma 7, if E were strongly silent on the probability of A given B , E would also contain the belief state p' which coincides with p^* outside the probability of A given B and satisfies $p'(A|B) = 1$; i.e., E would contain the belief state $p' := p^*(\cdot|A)p^*(B) + p^*(\cdot \cap \overline{B})$. But this is not the case because

$$p'(B'') = p^*(B''|A)p^*(B) + p^*(B'' \cap \overline{B}) = 0 \neq \pi_{B''},$$

where the second equality uses the shown fact that $p^*(B'' \cap A) = p^*(B'' \cap \overline{B}) = 0$. Hence, E is not strongly silent on the probability of A given B . ■

Next, we determine where dual-Jeffrey experiences are strongly silent.

Lemma 9 *For all dual-Jeffrey experiences E (of learning a new conditional probability distribution given a partition \mathcal{C}) and all events $\emptyset \subsetneq A \subsetneq B \subseteq \Omega (= \text{Supp}(E))$, E is strongly silent on the probability of A given B if and only if $A = \cup_{C \in \mathcal{C}_A} C$ and $B = \cup_{C \in \mathcal{C}_B} C$ for some sets $\emptyset \subsetneq \mathcal{C}_A \subsetneq \mathcal{C}_B \subseteq \mathcal{C}$.*

Proof. Let E , \mathcal{C} , A and B be as specified, and let $(\pi^C)^{C \in \mathcal{C}}$ be the learned conditional probability distribution given \mathcal{C} . First, if $A = \cup_{C \in \mathcal{C}_A} C$ and $B = \cup_{C \in \mathcal{C}_B} C$ for some sets $\emptyset \subsetneq \mathcal{C}_A \subsetneq \mathcal{C}_B \subseteq \mathcal{C}$ then E is strongly silent on the probability of A given B , as one can check using Lemma 7. Conversely, suppose that one cannot express A , B as such unions. Consider the belief state $p^* := \frac{1}{\#\mathcal{C}} \sum_{C \in \mathcal{C}} \pi^C$. Clearly, $p^* \in E$. If E were strongly silent on the probability of A given B then E would also contain

the belief state p' which coincides with p^* outside the probability of A given B and satisfies $p'(A|B) = 1$, i.e., the belief state

$$p' := p^*(\cdot|A)p^*(B) + p^*(\cdot \cap \overline{B}).$$

But E fails to contain p' , for the following reason. We distinguish between two cases.

Case 1: There is no set $\mathcal{C}_A \subseteq \mathcal{C}$ such that $A = \cup_{C \in \mathcal{C}_A} C$. Then there exists a $C \in \mathcal{C}$ such that $C \cap A, C \setminus A \neq \emptyset$. By the definition of p' (and the fact that $C \cap A, C \setminus A \neq \emptyset$), $p'(C \cap A) > p^*(C \cap A)$ and $0 < p'(C \setminus A) \leq p^*(C \setminus A)$. This implies that $p'(C), p^*(C) \neq 0$ and $p'(A|C) > p^*(A|C)$. So, as $p^*(\cdot|C) = \pi^C$ (by $p^* \in E$), $p'(\cdot|C) \neq \pi^C$, and therefore $p' \notin E$.

Case 2: There is a set $\mathcal{C}_A \subseteq \mathcal{C}$ such that $A = \cup_{C \in \mathcal{C}_A} C$. Then there is no $\mathcal{C}_B \subseteq \mathcal{C}$ such that $B = \cup_{C \in \mathcal{C}_B} C$; and so, there is a $C \in \mathcal{C}$ such that $C \cap B, C \setminus B \neq \emptyset$. As A is included in B and a union of sets in \mathcal{C} , $C \cap A = \emptyset$. Note that $p^*(C \cap B), p^*(C \setminus B) \neq 0$ (as $C \cap B, C \setminus B \neq \emptyset$ and by definition of p^*); further, that $p'(C \cap B) = p'(C \cap (B \setminus A)) = 0$ (where the first equality holds because $C \cap A = \emptyset$ and the second by definition of p'); and finally, that $p'(C) = p'(C \cap B) + p'(C \setminus B) = 0 + p^*(C \cap B) \neq 0$. Since $p'(C), p^*(C) \neq 0$, the conditional belief states $p'(\cdot|C)$ and $p^*(\cdot|C)$ are defined; they differ since $p'(C \cap B) = 0$ but $p^*(C \cap B) \neq 0$. Hence, as $p^*(\cdot|C) = \pi^C$ (by $p^* \in E$), $p'(\cdot|C) \neq \pi^C$, and so $p' \notin E$. ■

We now turn to Adams experiences. Before we can show where they are strongly silent, we derive two useful lemmas.

Lemma 10 *Every Adams experience E is convex, i.e., if $p', p'' \in E$ and $\alpha \in [0, 1]$, then $\alpha p' + (1 - \alpha)p'' \in E$.*

Proof. Let E, p', p'' and α be as specified, and fix any family $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ generating E . To show that $q := \alpha p' + (1 - \alpha)p'' \in E$, we consider any $B \in \mathcal{B}$ and $C \in \mathcal{C}$ such that $q(C) \neq 0$, and have to prove that $q(B|C) = \pi_B^C$. Note that

$$q(B|C) = \frac{q(B \cap C)}{q(C)} = \frac{\alpha p'(B \cap C) + (1 - \alpha)p''(B \cap C)}{\alpha p'(C) + (1 - \alpha)p''(C)}. \quad (20)$$

There are three cases:

- First let $p'(C) = 0$. Then also $p'(B \cap C) = 0$; and so by (20) $q(B|C) = \frac{p''(B \cap C)}{p''(C)} = p''(B|C)$, which equals π_B^C as $p'' \in E$.
- Now let $p''(C) = 0$. Then also $p''(B \cap C) = 0$; hence by (20) $q(B|C) = \frac{p'(B \cap C)}{p'(C)} = p'(B|C)$, which equals π_B^C as $p' \in E$.
- Finally, let $p'(C), p''(C) \neq 0$. Then $p'(B|C) = p''(B|C)$ ($= \pi_B^C$), i.e., $\frac{p'(B \cap C)}{p'(C)} = \frac{p''(B \cap C)}{p''(C)}$. So there is a $\beta > 0$ such that

$$p''(B \cap C) = \beta p'(B \cap C) \text{ and } p''(C) = \beta p'(C),$$

so that by (20) $q(B|C) = \frac{p'(B \cap C)}{p'(C)}$, which equals π_B^C . ■

Lemma 11 *If E is an Adams experience, $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ is its canonical family, $p_B \in \mathcal{P}$ with $\text{Supp}(p_B) \subseteq B$ for all $B \in \mathcal{B}$, and $\beta_C \geq 0$ for all $C \in \mathcal{C}$ where $\sum_{C \in \mathcal{C}} \beta_C = 1$, then E contains*

$$p \equiv \sum_{C \in \mathcal{C}, B \in \mathcal{B}} \beta_C \pi_B^C p_B \left(= \sum_{C \in \mathcal{C}} \beta_C \sum_{B \in \mathcal{B}} \pi_B^C p_B = \sum_{C \in \mathcal{C}} \beta_C \sum_{B \in \mathcal{B}: B \subseteq C} \pi_B^C p_B \right).$$

Proof. The lemma follows from the convexity of Adams experiences (see Lemma 10) since each p_B belongs to E and the coefficients $\beta_C \pi_B^C$ satisfy

$$\sum_{C \in \mathcal{C}, B \in \mathcal{B}} \beta_C \pi_B^C = \sum_{C \in \mathcal{C}} \beta_C \sum_{B \in \mathcal{B}} \pi_B^C = \sum_{C \in \mathcal{C}} \beta_C \times 1 = 1. \blacksquare$$

The next lemma determines where Adams experiences are strongly silent, combining insights from Lemmas 8 and 9 about Jeffrey and dual-Jeffrey experiences. In fact, the next lemma implies Lemma 9 – not surprisingly, since Adams experiences generalize dual-Jeffrey experiences. (We have nonetheless stated and proved Lemma 9 separately, as a useful warm-up for the next complex lemma.)

On a first reading of the next lemma, one might assume that $\mathcal{B} \cap \mathcal{C} = \emptyset$, so that $D_A = D_B = \cup_{D \in \mathcal{B} \cap \mathcal{C}} D = \emptyset$.

Lemma 12 *Consider an Adams experiences E and let $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ be the canonical family generating E (defined in Lemma 5). For all events $\emptyset \subsetneq A \subsetneq B \subseteq \Omega$ ($= \text{Supp}(E)$), E is strongly silent on the probability of A given B if and only if*

- (a) *either $B \subseteq B'$ for some $B' \in \mathcal{B}$,*
- (b) *or $A = (\cup_{C \in \mathcal{C}_A} C) \cup D_A$ and $B = (\cup_{C \in \mathcal{C}_B} C) \cup D_B$ for some $\mathcal{C}_A \subseteq \mathcal{C}_B \subseteq \mathcal{C} \setminus (\mathcal{B} \cap \mathcal{C})$ and some $D_A \subseteq D_B \subseteq \cup_{D \in \mathcal{B} \cap \mathcal{C}} D$.¹⁷*

Proof. Let E , $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$, A and B be as specified. For each $C \in \mathcal{C}$ let $\mathcal{B}_C := \{B \in \mathcal{B} : B \subseteq C\}$. Also, let $\mathcal{D} := \mathcal{B} \cap \mathcal{C}$ (note that $|\mathcal{D}| \leq 1$) and let $\Omega^* := \Omega \setminus (\cup_{D \in \mathcal{D}} D)$.

First, if A and B take the form (a) or (b), then A is strongly silent on the probability of A given B , as one can verify using Lemma 7.

Now suppose E is strongly silent on the probability of A given B . For a contradiction, suppose A and B are neither of the form (a) nor of the form (b). We derive a contradiction in each of the following cases.

Case 1: There does not exist any $C \in \mathcal{C} \setminus \mathcal{D}$ such that $C \cap A, C \setminus A \neq \emptyset$. In other words, $A = (\cup_{C \in \mathcal{C}_A} C) \cup D_A$ for some $\mathcal{C}_A \subseteq \mathcal{C} \setminus \mathcal{D}$ and some $D_A \subseteq \cup_{D \in \mathcal{D}} D$. Since condition (b) does not hold, B cannot take the form $(\cup_{C \in \mathcal{C}_B} C) \cup D_B$ with $\mathcal{C}_B \subseteq \mathcal{C} \setminus \mathcal{D}$ and $D_B \subseteq \cup_{D \in \mathcal{D}} D$. In other words, there exists a $C \in \mathcal{C} \setminus \mathcal{D}$ such that $C \cap B, C \setminus B \neq \emptyset$. Since $B \cap C, C \setminus B \neq \emptyset$ and since the set \mathcal{B}_C (which partitions C) has at least two members, there are distinct $\hat{B}, \tilde{B} \in \mathcal{B}_C$ such that $\hat{B} \cap B, \tilde{B} \setminus B \neq \emptyset$.

Note that since $B \not\supseteq C$ and $A \subseteq B$, we have $A \not\supseteq C$, and so $A \cap C = \emptyset$. Hence, as $A \neq \emptyset$ there is a $C^* \in \mathcal{C} \setminus \{C\}$ such that $A \cap C^* \neq \emptyset$. (Possibly $C^* \in \mathcal{D}$, in which case $A \cap C^*$ can differ from C^* .)

¹⁷Since by the canonicity of $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ the set $\mathcal{B} \cap \mathcal{C}$ is either empty or a singleton set $\{D^*\}$, the union $\cup_{D \in \mathcal{B} \cap \mathcal{C}} D$ is either \emptyset or D^* . In the first case the requirement ' $D_A \subseteq D_B \subseteq \cup_{D \in \mathcal{B} \cap \mathcal{C}} D$ ' reduces to $D_A = D_B = \emptyset$, and in the second case it reduces to $D_A \subseteq D_B \subseteq D^*$.

Now for each $B' \in \mathcal{B}_C$ we choose an $a_{B'} \in B'$, where $a_{\widehat{B}} \in \widehat{B} \cap B (\neq \emptyset)$ and $a_{\widetilde{B}} \in \widetilde{B} \setminus B (\neq \emptyset)$. By Lemma 11 (applied with $\beta_C = \beta_{C^*} = \frac{1}{2}$ and $\beta_{C'} = 0$ for all $C' \in \mathcal{C} \setminus \{C, C^*\}$), E contains

$$p^* := \frac{1}{2} \sum_{B' \in \mathcal{B}_C} \pi_{B'}^C \delta_{a_{B'}} + \frac{1}{2} \sum_{B' \in \mathcal{B}_{C^*}} \pi_{B'}^{C^*} \text{uniform}_{B'}.$$

Hence, since E is strongly silent on the probability of A given B , and since $p^*(A) \neq 0$ (as $A \cap C^* \neq \emptyset$) and $p^*(B \setminus A) \neq 0$ (as $a_{\widehat{B}} \in B \setminus A$), by Lemma 7 E also contains the belief state p' which satisfies $p'(A|B) = 1$ and coincides with p^* outside the probability of A given B , i.e., the belief state

$$p' := p^*(\cdot|A)p^*(B) + p^*(\cdot \cap \overline{B}).$$

Now

$$\begin{aligned} p'(\widehat{B}) &= p^*(\widehat{B}|A)p^*(B) + p^*(\widehat{B} \cap \overline{B}) = 0 \times p^*(B) + 0 = 0 \\ p'(\widetilde{B}) &= p^*(\widetilde{B}|A)p^*(B) + p^*(\widetilde{B} \cap \overline{B}) = 0 \times p^*(B) + p^*(a_{\widehat{B}}) \neq 0. \end{aligned}$$

Note that $p'(C) \geq p'(\widetilde{B}) > 0$ and $p'(\widehat{B}|C) = 0 \neq \pi_{\widehat{B}}^C$, a contradiction since $p' \in E$.

Case 2: There exists a $C \in \mathcal{C} \setminus \mathcal{D}$ such that $C \cap A, C \setminus A \neq \emptyset$.

Subcase 2.1: $(B \setminus A) \cap C = \emptyset$ (i.e., $A \cap C = B \cap C$). So, as $A \subsetneq B$, there is a $C^* \in \mathcal{C}$ such that $(B \setminus A) \cap C^* \neq \emptyset$. (Possibly $C \in \mathcal{D}$.) Hence, there is a $B^* \in \mathcal{B}_{C^*}$ such that $B^* \cap (B \setminus A) \neq \emptyset$. By Lemma 11 (applied with $\beta_{C^*} = \beta_C = \frac{1}{2}$ and $\beta_{C'} = 0$ for all $C' \in \mathcal{C} \setminus \{C^*, C\}$), E contains

$$\begin{aligned} p^* &:= \frac{1}{2} \left(\pi_{B^*}^{C^*} \text{uniform}_{B^* \cap (B \setminus A)} + \sum_{B' \in \mathcal{B}_{C^*} \setminus \{B^*\}} \pi_{B'}^{C^*} \text{uniform}_{B'} \right) \\ &\quad + \frac{1}{2} \sum_{B' \in \mathcal{B}_C} \pi_{B'}^C \text{uniform}_{B'}. \end{aligned}$$

So, because E is strongly silent on the probability of A given B (and because $p^*(A), p^*(B \setminus A) \neq 0$), by Lemma 7 E also contains the belief state p' which satisfies $p'(A|B) = 1$ and coincides with p^* outside the probability of A given B , i.e., the belief state

$$p' := p^*(\cdot|A)p^*(B) + p^*(\cdot \cap \overline{B}).$$

For all $\widetilde{B} \in \mathcal{B}_C$ such that $\widetilde{B} \cap A \neq \emptyset$ we have $\widetilde{B} \cap A = \widetilde{B} \cap B$ and $(0 <) p^*(A) < p^*(B)$, so that

$$p^*(\widetilde{B}|A) = \frac{p^*(\widetilde{B} \cap A)}{p^*(A)} > \frac{p^*(\widetilde{B} \cap B)}{p^*(B)} = p^*(\widetilde{B}|B),$$

and hence,

$$\begin{aligned} p'(\widetilde{B}) &= p^*(\widetilde{B}|A)p^*(B) + p^*(\widetilde{B} \cap \overline{B}) \\ &> p^*(\widetilde{B}|B)p^*(B) + p^*(\widetilde{B} \setminus B) = p^*(\widetilde{B} \cap B) + p^*(\widetilde{B} \setminus B) = p^*(\widetilde{B}). \end{aligned}$$

Further, for all $\tilde{B} \in \mathcal{B}_C$ such that $\tilde{B} \cap A$ ($\tilde{B} \cap B$) is empty we have

$$p'(\tilde{B}) = p^*(\tilde{B}|A)p^*(B) + p^*(\tilde{B} \cap \bar{B}) = 0 \times p^*(B) + p^*(\tilde{B}) = p^*(\tilde{B}).$$

As we have shown, $p'(\tilde{B}) \geq p^*(\tilde{B})$ for all $\tilde{B} \in \mathcal{B}_C$, where some inequalities hold strictly and some hold as equalities. For every $\hat{B} \in \mathcal{B}_C$ such that $p'(\hat{B}) = p^*(\hat{B})$ we have

$$p'(\hat{B}|C) = \frac{p'(\hat{B})}{\sum_{\tilde{B} \in \mathcal{B}_C} p'(\tilde{B})} < \frac{p^*(\hat{B})}{\sum_{\tilde{B} \in \mathcal{B}_C} p^*(\tilde{B})} = p^*(\hat{B}|C) = \pi_{\hat{B}}^C.$$

So, $p'(\hat{B}|C) \neq \pi_{\hat{B}}^C$, a contradiction since $p' \in E$.

Subcase 2.2: $(B \setminus A) \cap C \neq \emptyset$ and no set in \mathcal{B}_C includes $B \cap C$. Since $(B \setminus A) \cap C$ and $A \cap C$ are both non-empty, and since the union of these two sets, $B \cap C$, is not included in any set in \mathcal{B}_C (hence, intersects with at least two sets in \mathcal{B}_C), there exist distinct $B_1, B_2 \in \mathcal{B}_C$ such that

$$\begin{aligned} \emptyset &\neq B_1 \cap ((B \setminus A) \cap C) (= B_1 \cap (B \setminus A)) \\ \emptyset &\neq B_2 \cap (A \cap C) (= B_2 \cap A). \end{aligned}$$

Now for each $B' \in \mathcal{B}_C$ we fix an $a_{B'} \in B'$ such that $a_{B_1} \in B_1 \cap (B \setminus A)$ ($\neq \emptyset$) and $a_{B_2} \in B_2 \cap A$ ($\neq \emptyset$). By Lemma 11, E contains the measure

$$p^* := \sum_{B' \in \mathcal{B}_C} \pi_{B'}^C \delta_{a_{B'}}.$$

So, since E is strongly silent on the probability of A given B (and since $p^*(A) \neq 0$ as $a_{B_2} \in A$ and since $p^*(B \setminus A) \neq 0$ as $a_{B_1} \in B \setminus A$), by Lemma 7 E also contains the belief state p' which satisfies $p'(A|B) = 1$ and coincides with p^* outside the probability of A given B , i.e., the belief state

$$p' := p^*(\cdot|A)p^*(B) + p^*(\cdot \cap \bar{B}).$$

We have

$$\begin{aligned} p'(B_2) &= p^*(B_2|A)p^*(B) + p^*(B_2 \cap \bar{B}) = 0 \times p^*(B) + 0 = 0, \\ p'(C) &= p^*(C|A)p^*(B) + p^*(C \cap \bar{B}) = 1 \times p^*(B) + p^*(C \setminus B) = p^*(C) = 1. \end{aligned}$$

So, $p'(B_2|C) = 0 \neq \pi_{B_2}^C$, a contradiction since $p' \in E$.

Subcase 2.3: $(B \setminus A) \cap C \neq \emptyset$ and some $B^* \in \mathcal{B}_C$ includes $B \cap C$. Since condition (a) does not hold, $B \not\subseteq B^*$. So, $B \neq B \cap C$, i.e., $B \setminus C \neq \emptyset$. Hence there are $\hat{C} \in \mathcal{C} \setminus \{C\}$ and $\hat{B} \in \mathcal{B}_{\hat{C}}$ such that $\hat{B} \cap B \neq \emptyset$; hence, $\hat{C} \cap B \neq \emptyset$. (Possibly $\hat{B} = \hat{C}$.)

Subsubcase 2.3.1: $A \cap \hat{C} \neq \emptyset$. By Lemma 11, the belief state

$$p^* := \frac{1}{2} \left(\pi_{B^*}^C \text{uniform}_{B^* \cap A} + \sum_{B' \in \mathcal{B}_C \setminus \{B^*\}} \pi_{B'}^C \text{uniform}_{B'} \right) + \frac{1}{2} \sum_{B' \in \mathcal{B}_{\hat{C}}} \pi_{B'}^{\hat{C}} \text{uniform}_{B'}$$

belongs to E . Since p^* belongs to E which is strongly silent on the probability of A given B (and since $p^*(A), p^*(B \setminus A) \neq 0$), E also contains the belief state p' for which $p'(A|B) = 0$ and which coincides with p^* outside the probability of A given B ,

$$p' := p^*(\cdot|B \setminus A)p^*(B) + p^*(\cdot \cap \bar{B}).$$

Notice that

$$\begin{aligned} p'(B^*) &= p^*(B^*|B \setminus A)p^*(B) + p^*(B^* \cap \overline{B}) = 0 \times p^*(B) + 0 = 0, \\ p'(C) &= p^*(C|B \setminus A)p^*(B) + p^*(C \cap \overline{B}) = 0 \times p^*(B) + p^*(C \setminus B^*) = p^*(C \setminus B^*), \end{aligned}$$

where the latter is positive since $B^* \neq C$. So, $p'(B^*|C) = 0 \neq \pi_{B^*}^C$, a contradiction as $p' \in E$.

Subsubcase 2.3.2: $A \cap \widehat{C} = \emptyset$. We re-define p^* by replacing ‘ $uniform_{B^* \cap A}$ ’ by ‘ $uniform_{B^* \cap (B \setminus A)}$ ’ in the previous definition of p^* . Again, $p^* \in E$ by Lemma 11, so that as E is strongly silent on the probability of A given B (and as $p^*(A), p^*(B \setminus A) \neq 0$) E also contains the belief state p' for which $p'(A|B) = 1$ and which coincides with p^* outside the probability of A given B ,

$$p' = p^*(\cdot|A)p^*(B) + p^*(\cdot \cap \overline{B}).$$

Notice that

$$\begin{aligned} p'(B^*) &= p^*(B^*|A)p^*(B) + p^*(B^* \cap \overline{B}) = 0 \times p^*(B) + 0 = 0, \\ p'(C) &= p^*(C|A)p^*(B) + p^*(C \cap \overline{B}) = 0 \times p^*(B) + p^*(C \cap \overline{B}) = p^*(C \cap \overline{B}) > 0. \end{aligned}$$

So, $p'(B^*|C) = 0 \neq \pi_{B^*}^C$, a contradiction since $p' \in E$. ■

A.5 Theorem 1

Based on previous lemmas, we now prove the central result.

Proof of Theorem 1. It suffices to consider Jeffrey and Adams revision, since Bayesian and dual-Jeffrey revision are extended by Jeffrey and Adams revision, respectively. We prove each direction of implication in turn.

1. First we consider a responsive and conservative revision rule on one of the domains $\mathcal{D}_{\text{Jeffrey}}$ and $\mathcal{D}_{\text{Adams}}$. We show that the rule is Jeffrey respectively Adams revision, by distinguishing between the two domains.

Jeffrey: Suppose $(p, E) \in \mathcal{D}_{\text{Jeffrey}}$, say $E = \{p' : p(B) = \pi_B \ \forall B \in \mathcal{B}\}$. Then p_E is given by Jeffrey revision because we may expand $p_E(a)$ as

$$p_E = \sum_{B \in \mathcal{B}: p_E(B) \neq 0} p_E(\cdot|B)p_E(B), \quad (21)$$

in which $p_E(B)$ reduces to π_B by Responsiveness, and $p_E(\cdot|B)$ reduces to $p(\cdot|B)$ by Conservativeness (as by Lemma 8 E is strongly silent on the probability given B of any event strictly between \emptyset and B).

Adams: Now suppose $(p, E) \in \mathcal{D}_{\text{Adams}}$, say $E = \{p' : p'(B|C) = \pi_B^C \ \forall B \in \mathcal{B} \ \forall C \in \mathcal{C} \text{ such that } p'(C) \neq 0\}$. Then p_E is given by Adams revision because we may expand p_E as

$$p_E = \sum_{B \in \mathcal{B}, C \in \mathcal{C}: p_E(B \cap C) \neq 0} p_E(\cdot|B \cap C)p_E(B|C)p_E(C),$$

in which $p_E(B|C)$ reduces to π_B^C by Responsiveness, $p_E(C)$ reduces to $p(C)$ by Conservativeness (as by Lemma 12 E is strongly silent on the probability of C given Ω if $C \neq \Omega$), and $p_E(\cdot|B \cap C)$ reduces to $p(\cdot|B \cap C)$ by Conservativeness (as by Lemma 12 E is strongly silent on the probability given $B \cap C$ of any event strictly between \emptyset and $B \cap C$).

2. Conversely, we now show that Jeffrey and Adams revision are responsive and conservative. Responsiveness is obvious. To show Conservativeness, consider any (p, E) in the rule's domain ($\mathcal{D}_{\text{Jeffrey}}$ or $\mathcal{D}_{\text{Adams}}$) and any events $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(E)$ such that E is strongly silent on the probability of A given B and $p_E(B), p(B) \neq 0$. We have to show that $p_E(A|B) = p(A|B)$, by distinguishing between the two rules.

Jeffrey: First suppose the rule in question is Jeffrey revision. Then the experience takes the form $E = \{p' : p'(B) = \pi_B \ \forall B \in \mathcal{B}\}$ for some learned probability distribution $(\pi_B)_{B \in \mathcal{B}}$ on some partition \mathcal{B} . As E is strongly silent on the probability of A given B , by Lemma 8 $B \subseteq B'$ for some $B' \in \mathcal{B}$. It follows that $p_E(A|B) = p(A|B)$, because

$$p_E(A|B) = \frac{p_E(A)}{p_E(B)} = \frac{p(A|B')\pi_{B'}}{p(B|B')\pi_{B'}} = \frac{p(A)/p(B')}{p(B)/p(B')} = p(A|B),$$

where the second equality holds by definition of Jeffrey revision.

Adams: Now consider Adams revision. Then E is an Adams experience, of the form $E = \{p' : p'(B|C) = \pi_B^C \ \forall B \in \mathcal{B} \ \forall C \in \mathcal{C} \text{ such that } p'(C) \neq 0\}$ where $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ is a conditional probability distribution on some partition \mathcal{B} given another \mathcal{C} . By Lemma 5 we may assume that the family $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$ is the canonical one for E , i.e., that \mathcal{B} refines \mathcal{C} and $\mathcal{B} \cap \mathcal{C}$ is empty or singleton. By E 's strong silence on the probability of A given B and Lemma 12, there are only two cases:

- (a) $B \subseteq B'$ for some $B' \in \mathcal{B}$, or
- (b) $A = (\cup_{C \in \mathcal{C}_A} C) \cup D_A$ and $B = (\cup_{C \in \mathcal{C}_B} C) \cup D_B$ for some $\mathcal{C}_A \subseteq \mathcal{C}_B \subseteq \mathcal{C} \setminus (\mathcal{B} \cap \mathcal{C})$ and some $D_A \subseteq D_B \subseteq \cup_{D \in \mathcal{B} \cap \mathcal{C}} D$. (So, as $\mathcal{B} \cap \mathcal{C}$ is empty or a singleton set $\{D\}$, we have $D_A = D_B = \emptyset$ or $D_A \subseteq D_B \subseteq D$, respectively.)

In case (a) we have $p_E(A|B) = p(A|B)$ because, writing C' for the member of \mathcal{C} which includes B' , we have

$$p_E(A|B) = \frac{p_E(A)}{p_E(B)} = \frac{p(A|B')\pi_{B'}^{C'}p(C')}{p(B|B')\pi_{B'}^{C'}p(C')} = \frac{p(A|B')}{p(B|B')} = p(A|B).$$

In case (b) we again have $p_E(A|B) = p(A|B)$, this time because

$$\begin{aligned} p(A|B) &= \frac{p(A)}{p(B)} = \frac{\sum_{C \in \mathcal{C}_A} p(C) + p(D_A)}{\sum_{C \in \mathcal{C}_B} p(C) + p(D_B)} \\ p_E(A|B) &= \frac{p_E(A)}{p_E(B)} = \frac{\sum_{C \in \mathcal{C}_A} p_E(C) + p_E(D_A)}{\sum_{C \in \mathcal{C}_B} p_E(C) + p_E(D_B)}, \end{aligned}$$

where, as one easily checks, each $p_E(C)$ equals $p(C)$, and $p_E(D_A) = p(D_A)$, and $p_E(D_B) = p(D_B)$. To see for instance why $p_E(D_A) = p(D_A)$, recall that either $D_A = \emptyset$ or $D_A \subseteq D$ ($\in \mathcal{B} \cap \mathcal{C}$). If $D_A = \emptyset$ then clearly $p_E(D_A) = p(D_A)$. If $D_A \subseteq D$ then $p_E(D_A) = p(D_A|D)\pi_D^D p(D) = p(D_A)$ (where, as usual, ' $p(D_A|D)\pi_D^D p(D)$ ' is defined as 0 if $p(D) = 0$). ■