A Unified Characterization of Belief-Revision Rules

Franz Dietrich, Christian List, and Richard Bradley[1]

October 2010, revised November 2014

Abstract

This paper characterizes several 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 a new conditional probability function. Despite their differences, these revision rules can be characterized in terms of the same two axioms: *responsiveness*, which requires that revised beliefs incorporate what has been learnt, and *conservativeness*, which requires that beliefs on which the learnt input is 'silent' do not change. So, the four revision rules apply the same principles, albeit to different learnt inputs. To illustrate that there is room for non-Bayesian belief revision in economic theory, we also sketch a simple decision-theoretic application.

Keywords: Subjective probability, Bayes's rule, Jeffrey's rule, axiomatic foundations, fine-grained versus coarse-grained beliefs, unawareness

1 Introduction

A belief-revision rule captures how an agent's subjective probabilities should change when the agent learns something new. The standard example is Bayes's rule. Here, the agent learns that some event has occurred, and the response is to raise the subjective probability of that event to 1, while retaining all probabilities conditional on it. More formally, let Ω be the underlying set of possible worlds (where Ω is non-empty and, for simplicity, finite or countably infinite).[2] Subsets of Ω are called *events*. Beliefs are represented by some probability measure on the set of all events. Bayes's rule says that, upon learning that some event $B \subseteq \Omega$ has occurred (with $p(B) \neq 0$), one should move from the prior probability measure p to the posterior probability measure p' given by

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

In economic theory, belief changes are almost always modelled in this way. The aim of this paper is to draw attention to some alternative, non-Bayesian belief-revision

[1] F. Dietrich, Paris School of Economics (CNRS) and University of East Anglia; C. List, LSE; R. Bradley, LSE. We have presented this paper at several seminars and conferences and are grateful to the audiences for helpful comments. These occasions include: D-TEA 2010 (HEC & Polytechnique, Paris, June 2010), LSE's Choice Group Seminar (LSE, September 2010), Pluralism in the Foundations of Statistics (University of Kent, September 2010), and Decisions, Games, and Logic 2012 (University of Munich, 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.

[2] Everything we say could be generalized to an arbitrary measurable set Ω.

rules, which are seldom discussed in economics. We show that Bayes's rule is just one special case of a larger family of rules, which can all be axiomatically characterized in a unified way. They differ only in what they assume about the nature of the input prompting the agent's belief change. Under Bayes's rule, the learnt input is always the occurrence of some event, but this is more restrictive than often recognized.

To illustrate, we begin with a story that Kotaro Suzumura has kindly shared with us. It is about how Suzumura was offered his first faculty position at a British university. He was at the time a British Council Visiting Scholar in Cambridge and had been encouraged by his supervisor, Frank Hahn, to apply for a lectureship at the LSE. After being interviewed for the position, he was telephoned by Terence Gorman to inform him of the outcome. As Suzumura tells the story, Gorman had a very thick accent which, combined with the fact that Suzumura was not a native English speaker and that the conversation was over the telephone, meant that he had great difficulty understanding Gorman. He got the impression that he was being offered the post, but was naturally reluctant to insist on greater clarity. At the end of the conversation, he was still far from sure that he had received an offer – so unsure that he felt compelled to telephone Amartya Sen, who was then at the LSE, to ask whether it was true. Fortunately, Sen was able to confirm that it was, and Suzumura could accept the offer. The rest, as they say, is history.

This story illustrates an instance of belief revision triggered by a noisy signal. Such cases present challenges to the Bayesian modeller. Before the telephone conversation, Suzumura presumably attached a fairly low probability to the event of him being appointed to the lectureship. After the conversation, he attached a somewhat higher probability to it, but one that still fell short of certainty. For this, a second conversation (with Sen) was needed. If the first change in Suzumura's probabilities is to be modelled as an application of Bayes's rule, then it will clearly not suffice to restrict attention to the 'naïve' set of possible worlds $\Omega = \{\text{appointed, not appointed}\}$. Relative to that 'naïve' set, a Bayesian belief change could never increase the probability of the event 'appointed' without raising it all the way to 1.

The modeller will need to enrich the set Ω to capture the possible sensory experiences responsible for Suzumura's shift in probabilities over the events 'appointed' and 'not appointed'. So, the enriched set of worlds will have to be something like

$$\Omega' = \{\text{appointed, not appointed}\} \times \mathcal{A},$$

where $\mathcal{A}$ is the set of possible analogue auditory signals received by Suzumura's eardrums. The signal he received from Gorman will then correspond to some subset of Ω', specifically one of the form $B = \{\text{appointed, not appointed}\} \times A$, where $A \subseteq \mathcal{A}$ is a particular 'auditory event'. Not much less than this representation will do. Even replacing $\mathcal{A}$ with a smaller set of possible verbal messages would not suffice, since Gorman's words were subject to a triple distortion, first by his thick accent, then by imperfections of the telephone line, and finally by the interpretation of someone new to an English-speaking environment.

Enriching the set Ω in this way, however, has definite modelling costs. First, the agent (Suzumura) almost certainly did not have prior subjective probabilities over the events from such a rich set. In light of the huge range of possible signals, the set Ω' is of dizzying size and complexity, when compared to the 'naïve' set Ω on which the agent's attention was originally focused. Second, it is doubtful that before the conversation he was even aware of the possibility of such complex auditory signals (probably he had never heard, or even heard of, an accent like Gorman's). So, a Bayesian model of Suzumura's story, and others like it, must inevitably involve a heavy dose of fiction. It ascribes to the agent greater *prior opinionation* (ability to assign prior probabilities) and greater *awareness* (conceptualization or consideration of events) than psychologically plausible. In a similar vein, Diaconis and Zabell (1982, p. 823) have called the assignment of prior subjective probabilities to 'many classes of sensory experiences [...] forced, unrealistic, or impossible' (see also Jeffrey 1957 and Shafer 1981). Of course, whether this is a problem or not will depend on the uses to which the model is put; we are not denying the usefulness of 'as-if' modelling in all cases. But good scientific practice should encourage us to investigate whether other belief-revision rules are better at capturing cases like the present one and how these other rules relate to Bayes's rule. This is what motivates this paper.

A prominent generalization of Bayes's rule is Jeffrey's rule (e.g., Jeffrey 1957, Shafer 1981, Diaconis and Zabell 1982, Grünwald and Halpern 2003). Here, the agent learns a new probability of some event, for instance a 20% probability of an accident or a 75% probability of the offer of a lectureship, as perhaps in Suzumura's case. More generally, the agent learns a new probability distribution of some random variable such as the level of rainfall or GDP. The response, then, is to assign the new distribution to that random variable, while retaining all probabilities conditional on it. More formally, let $\mathcal{B}$ be a partition of the set Ω into finitely many non-empty events, and suppose the agent learns a new probability π_B for each event B in $\mathcal{B}$. The family of learnt probabilities, $(\pi_B)_{B\in\mathcal{B}}$, is a *probability distribution* over $\mathcal{B}$ (i.e., consists of non-negative numbers with sum-total 1). Jeffrey's rule says that, upon learning $(\pi_B)_{B\in\mathcal{B}}$, one should move from the prior probability measure p to the posterior probability measure p' given by

$$p'(A) = \sum_{B\in\mathcal{B}} p(A|B)\pi_B \text{ for all events } A \subseteq \Omega.$$[3]

For instance, suppose the agent learns that it will rain with probability $\frac{1}{2}$, snow with probability $\frac{1}{3}$, and remain dry with probability $\frac{1}{6}$. Then the partition $\mathcal{B}$ (of a suitable set Ω) contains the events of rain (R), snow (S), and no precipitation (N), where $\pi_R = \frac{1}{2}$, $\pi_S = \frac{1}{3}$, and $\pi_N = \frac{1}{6}$. Bayes's rule is the special case where $\mathcal{B}$ partitions Ω into an event B and its complement $\overline{B}$, with $\pi_B = 1$ and $\pi_{\overline{B}} = 0$. (The *complement* of any event B is $\overline{B} = \Omega\backslash B$.)

We develop a general framework in which different belief-revision rules can be defined and compared. In this framework, what is being learnt by the agent can take a variety of forms; we call this the *learnt input*. It can be interpreted as the constraint that a particular experience – say the receipt of some signal – imposes on the agent's beliefs. Examples of learnt inputs are event occurrences for Bayes's rule and learnt probability distributions for Jeffrey's. We show that four salient belief-revision rules – Bayes's rule, Jeffrey's rule, and two others introduced below (Adams's rule and the

[3] For p' to be well-defined, we must have $\pi_B = 0$ whenever $p(B) = 0$. This ensures that if a term $p(A|B)$ is undefined in the displayed formula (because $p(B) = 0$), then this term does not matter (because it is multiplied by $\pi_B = 0$).

'dual-Jeffrey' rule) – can be uniquely characterized in terms of the same two axioms, simply applied to different domains of learnt inputs.

Our axioms are (i) a *responsiveness axiom*, which requires that revised beliefs be consistent with the learnt input, and (ii) a *conservativeness axiom*, which requires that those beliefs on which the input is 'silent' (in a sense to be made precise) do not change. The fact that several non-standard belief-revision rules can be justified in complete analogy to Bayes's rule should assuage some economists' worry that non-Bayesian rules automatically involve costly departures from compelling principles of belief revision. We hope that this will, in turn, inspire further work on economic applications and behavioural implications of non-Bayesian forms of belief revision. To suggest some steps in this direction, we conclude the paper with a discussion of how non-Bayesian belief revision may be introduced into decision and game theory, especially to capture 'unforeseen learning'; we also briefly revisit the issue of unawareness.[4]

Prior literature. Three of the four belief-revision rules we discuss – Bayes's, Jeffrey's, and Adams's rules – have been axiomatically characterized in previous work, though in different ways. The approach has usually been a 'distance-based' one. This consists in showing that a given revision rule generates posterior beliefs that incorporate the information learnt, while deviating as little as possible from prior beliefs, relative to some notion of 'distance' between beliefs.[5] Bayes's and Jeffrey's rules have been characterized relative to either the *variation distance* (defined by the maximal absolute difference in probability, over all events in the algebra), the *Hellinger distance*, or the *relative-entropy distance* (e.g., Csiszar 1967, 1977, van Fraassen 1981, Diaconis and Zabell 1982, Grünwald 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 (2011) have recently characterized Adams's rule by invoking 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 characterizations may be, they give a non-unified picture of belief revision. Different notions of distance are invoked to justify different revision rules, and their interpretation and relative advantages are controversial.

Another set of characterization results invokes the idea of 'rigidity' rather than distance-minimization (see Jeffrey 1957 for Bayes's and Jeffrey's rules, and Bradley 2005 for Adams's). For example, Bayesian belief revision is 'rigid' in the sense of preserving the conditional probability of any event, given the learnt event. Although the rigidity-based approach is closer in spirit to ours, it also lacks unification. However, within our framework, one may interpret the conservativeness axiom as a more unified version of earlier rigidity axioms, applicable to any belief-revision rule. For an overview of various forms of probabilistic belief and belief revision, we refer the reader to Halpern's handbook (2003). Since we here deal exclusively with beliefs that are represented by subjective probability measures, we set aside the literature on the revision of beliefs that do not take this form.

[4] For a brief discussion of dynamic-consistency arguments for Jeffrey's rule, similar to classic dynamic-consistency arguments for Bayes's rule, see Vineberg (2011).

[5] Distance is represented by a function $d : \mathcal{P} \times \mathcal{P} \to \mathbb{R}$, where $\mathcal{P}$ is the set of all probability measures over the events from Ω and $d(p, p')$ is interpreted as the distance between two such measures, $p, p' \in \mathcal{P}$.

2 A general framework

We can study attitude revision in general by specifying

(i) a set $\mathcal{P}$ of possible *attitudinal states* in which a given agent can be, and

(ii) a set $\mathcal{I}$ of possible *inputs* which can influence the agent's attitudinal state.

A *revision rule* maps pairs (p, I) of an initial attitudinal state p in $\mathcal{P}$ and an input I in $\mathcal{I}$ to a new attitudinal state $p' = p_I$ in $\mathcal{P}$. The pair (p, I) belongs to some domain $\mathcal{D} \subseteq \mathcal{P} \times \mathcal{I}$ containing those attitude-input pairs that are admissible under the given revision rule. A revision rule is thus a function from $\mathcal{D}$ to $\mathcal{P}$ (see also Dietrich 2012).

Since we focus on belief revision, attitudinal states are subjective probability measures. So, the set $\mathcal{P}$ of possible attitudinal states is the set of all probability measures over the events from Ω. Formally, a *probability measure* is a countably additive function $p : 2^{\Omega} \to [0, 1]$ with $p(\Omega) = 1$ (where Ω is the underlying finite or countably infinite set of worlds). We call any p in $\mathcal{P}$ a *belief state*. How can we define the set $\mathcal{I}$ of possible inputs? Looking at Bayes's rule alone, one might be tempted to define them as observed events $B \subseteq \Omega$. But Jeffrey's rule and the other rules introduced below permit different inputs, such as a family $(\pi_B)_{B\in\mathcal{B}}$ of learnt probabilities in Jeffrey's case.

Methodologically, one should not tie the notion of a 'learnt input' too closely to one particular revision rule, by defining it as a mathematical object that is tailor-made for that rule. This would exclude other revision rules from the outset and thereby prevent us from giving a fully compelling axiomatic characterization of the rule in question. Instead, we need an abstract notion of a 'learnt input'.

We define a *learnt input* as a set of belief states $I \subseteq \mathcal{P}$, interpreted as the set of those belief states that are consistent with the input. We can think of the input I as the constraint that a particular experience, such as the receipt of some signal, imposes on the agent's belief state. The set of logically possible inputs is $\mathcal{I} = 2^{\mathcal{P}}$. Note that this is deliberately general. An agent's belief change from p to p_I upon learning $I \in \mathcal{I}$ is *responsive to the input* if $p_I \in I$. We can now define the inputs involved in Bayesian revision and Jeffrey revision.

Definition 1 *A learnt input $I \in \mathcal{I}$ is*

- ***Bayesian*** *if $I = \{p' : p'(B) = 1\}$ for some event $B \neq \varnothing$; we then write $I = I_B$;*[6]
- ***Jeffrey*** *if $I = \{p' : p'(B) = \pi_B$ for all $B \in \mathcal{B}\}$ for some probability distribution $(\pi_B)_{B\in\mathcal{B}}$ on some partition $\mathcal{B}$; we then write $I = I_{(\pi_B)_{B\in\mathcal{B}}}$.*[7]

Here, and in what follows, we use the term *partition* to refer to a partition of Ω into finitely many non-empty events.

Clearly, every Bayesian input is also a Jeffrey input, while the converse is not true. Later, we introduce the notions of Adams inputs and dual-Jeffrey inputs. Some

[6] The representation $I = I_B$ is unique, because for any Bayesian input I, there exists a unique event B such that $I = I_B$.

[7] For any Jeffrey input I, the corresponding 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 $|\{B \in \mathcal{B} : \pi_B = 0\}| \leq 1$.

possible inputs $I \in \mathcal{I}$ are of none of these kinds, such as $I = \{p' : p'(A \cap B) > p'(A)p'(B)\}$, which captures the constraint that the events A and B are positively correlated, or $I = \{p' : p'(A) \geq 9/10\}$, which captures the constraint that A is very probable. In general, the smaller the set I, the stronger (more constraining) the input. The strongest consistent inputs are the singleton sets $I = \{p'\}$, which require adopting the new belief state p' regardless of the initial belief state. The weakest input is the set $I = \mathcal{P}$, which allows the agent to retain his old belief state.

We are now able to define Bayes's and Jeffrey's rules in this framework.

Definition 2

- *Let $\mathcal{D}_{Bayes}$ be the set of all pairs $(p, I) \in \mathcal{P} \times \mathcal{I}$ such that $I = I_B$ is a Bayesian input compatible with p (which means $p(B) \neq 0$).* ***Bayes's rule*** *is the revision rule on $\mathcal{D}_{Bayes}$ which maps each $(p, I_B) \in \mathcal{D}_{Bayes}$ to $p' \in \mathcal{P}$, where*

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

- *Let $\mathcal{D}_{Jeffrey}$ be the set of all pairs $(p, I) \in \mathcal{P} \times \mathcal{I}$ such that $I = I_{(\pi_B)_{B\in\mathcal{B}}}$ is a Jeffrey input compatible with p (which means $\pi_B = 0$ whenever $p(B) = 0$).* ***Jeffrey's rule*** *is the revision rule on $\mathcal{D}_{Jeffrey}$ which maps each $(p, I_{(\pi_B)_{B\in\mathcal{B}}}) \in \mathcal{D}_{Jeffrey}$ to $p' \in \mathcal{P}$, where*

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

The domains $\mathcal{D}_{\text{Bayes}}$ and $\mathcal{D}_{\text{Jeffrey}}$ are the maximal domains for which formulas (1) and (2) are well-defined.[8] Jeffrey's rule extends Bayes's, i.e., it coincides with Bayes's rule on the subdomain $\mathcal{D}_{\text{Bayes}}$ ($\subseteq \mathcal{D}_{\text{Jeffrey}}$). Later, we introduce two further revision rules: Adams's rule and the dual-Jeffrey rule.

3 An axiomatic characterization

We now introduce two plausible axioms that a belief-revision rule may be expected to satisfy and show that they imply that the agent must revise his beliefs in accordance with Bayes's rule in response to any Bayesian input and in accordance with Jeffrey's rule in response to any Jeffrey input. Later, we extend that characterization to the two other revision rules we have announced. All proofs are given in the Appendix.

3.1 Two axioms

Let $\mathcal{D} \subseteq \mathcal{P} \times \mathcal{I}$ be the domain of the belief-revision rule. For each belief-input pair $(p, I) \in \mathcal{D}$, we write $p_I \in \mathcal{P}$ to denote the revised belief state. Our first axiom says that the revised belief state should be responsive to the learnt input.

Responsiveness: $p_I \in I$ for all belief-input pairs $(p, I) \in \mathcal{D}$.

[8] The definition of each revision rule and its domain relies on the fact that each Bayesian input I is uniquely representable as $I = I_B$ and that each Jeffrey input is 'almost' uniquely representable as $I = I_{(\pi_B)_{B\in\mathcal{B}}}$, where any residual non-uniqueness makes no difference to the revised belief state or the criterion for including (p, I) in the domain. For details, see Lemmas 1 and 3 in the Appendix.

Responsiveness guarantees that the agent's revised belief state respects the constraint given by the input. For example, in response to a Bayesian input, the agent assigns probability one to the learnt event.

The second axiom expresses a natural conservativeness requirement: those 'parts' of the agent's belief state on which the learnt input is 'silent' should not change in response to it. In short, the learnt input should have no effect where it has nothing to say. To define that axiom formally, we must answer two questions: what do we mean by 'parts of a belief state', and when is a given input 'silent' on them? To answer these questions, note that, intuitively:

- a Bayesian input is not silent on the probability of the learnt event B, but is silent on all conditional probabilities, given B; and
- a Jeffrey input is not silent on the probabilities of the events in the relevant partition $\mathcal{B}$, but is silent on all conditional probabilities, given these events.

So, the 'parts' of the agent's belief state on which Bayesian inputs and Jeffrey inputs are 'silent' are conditional probabilities of some events, given others. The relevant conditional probabilities are preserved by Bayes's and Jeffrey's rules, so that these rules are intuitively conservative.

In the next subsection, we define formally what it means for a learnt input to be 'silent' on the probability of one event, given another. Once we have this definition, we can formulate our conservativeness axiom as follows.

Conservativeness (axiom scheme): For all belief-input pairs $(p, I) \in \mathcal{D}$, if I is 'silent' on the probability of a (relevant) event A given another B, this conditional probability is preserved, i.e., $p_I(A|B) = p(A|B)$ (if $p_I(B), p(B) \neq 0$).

3.2 When is a learnt input silent on the probability of one event, given another?

Our aim is to define when a learnt input $I \in \mathcal{I}$ is 'silent' on the probability of one event A, given another event B (where possibly $B = \Omega$). Our analysis will be fully general, i.e., not restricted to any particular class of inputs, such as Bayesian inputs or Jeffrey inputs. We first note that we need to define silence only for the case in which

$$\varnothing \subsetneq A \subsetneq B \subseteq Supp(I),$$

where $Supp(I)$ is the *support* of I, defined as $\{\omega \in \Omega : p'(\omega) \neq 0 \text{ for some } p' \in I\}$.[9]

There are two plausible notions of 'silence', which lead to two different variants of our conservativeness axiom. We begin with the weaker notion. A learnt input is *weakly silent* on the probability of A given B if it permits this conditional probability to take any value. Formally:

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

[9] Here, and elsewhere, we write $p'(\omega)$ as an abbreviation for $p'(\{\omega\})$ when we refer to the probability of a singleton event $\{\omega\}$.

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

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

Although this axiom may seem plausible, it leads to an impossibility result.

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

Note that, on the small domain $\mathcal{D}_{Bayes}$, there is no such impossibility, because Bayes's rule is responsive as well as strongly conservative. On that domain, the present strong conservativeness axiom is equivalent to our later, weaker one. The impossibility occurs on domains on which the two conservativeness axioms come apart.

We weaken strong conservativeness by strengthening the notion of silence. The key insight is that even if a learnt input I is weakly silent on the probability of A given B, it may still implicitly constrain the relationship between this conditional probability and others. Suppose, for instance, that $\Omega = \{0, 1\}^2$, where the first component of a world $(g, j) \in \Omega$ represents whether Richard has gone out (g=1) or not (g=0), and the second whether Richard is wearing a jacket (j=1) or not (j=0). Consider the event that Richard has gone out, $G = \{(g, j) : g = 1\}$, and the event that he is wearing a jacket, $J = \{(g, j) : j = 1\}$. Some inputs 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, especially those conditional on J. Consider, for instance, the Jeffrey input which says that G is 90% probable, formally $I = \{p' : p'(G) = 0.9\}$. It is compatible with any probability of J and is 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 I), then the probability of J can no longer exceed 0.9. If it did, the probability of G would exceed 0.9, which would contradict the learnt input I. In short, although I does not directly constrain the agent's subjective probability for J, it constrains it indirectly, i.e., after other parts of the belief state have been fixed.

A learnt input is *strongly silent* on the 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. In other words, they are all the beliefs of which the belief state p' consists, *over and above the probability of A given B*. More precisely, assuming again that A is included in B, they are given by the quadruple consisting of the unconditional probability $p'(B)$ and the conditional probabilities $p'(\cdot|A)$, $p'(\cdot|B\backslash A)$, and $p'(\cdot|\overline{B})$.[10] This quadruple and the conditional

[10] This informal discussion assumes that $p'(A), p'(B\backslash C), p'(\overline{B}) \neq 0$.

probability $p'(A|B)$ jointly determine the belief state p', because

$$p' = p'(\cdot|A) \underbrace{p'(A)}_{p'(A|B)p'(B)} + p'(\cdot|B\backslash A) \underbrace{p'(B\backslash A)}_{p'(B)-p'(A|B)p'(B)} + p'(\cdot|\overline{B}) \underbrace{p'(\overline{B})}_{1-p'(B)}.$$

If an input I 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 I (which requires them to match those of some belief state p^* in I). This idea is illustrated in Figure 1, where a learnt input I is

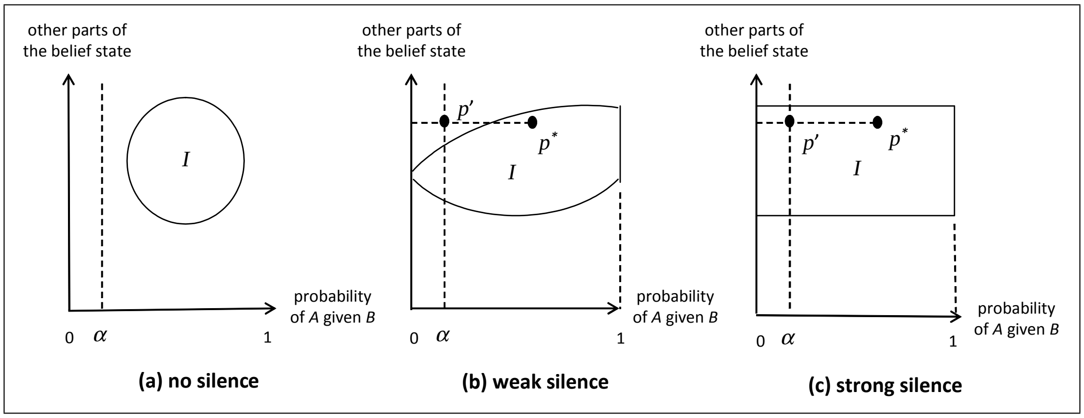

Figure 1: An input's weak or strong silence on some conditional probability

represented in the space whose horizontal coordinate represents the probability of A given B and whose vertical) coordinate represents the other parts of the agent's belief state (collapsed into a single dimension for illustration). In part (a), input I (represented by the circular region) is not silent at all on the probability of A given B, since many values of this probability, such as α, are ruled out by I. In part (b), input I (represented by the 'oval' region) is weakly but not strongly silent on the probability of A given B. This is because I is consistent with *any* value of that probability, but to combine it with a particular value, such as α, other parts of the belief state can no longer be freely chosen. In part (c), input I (represented by the rectangular region) is strongly silent on the probability of A given B. It is consistent with any value of that probability, even after other parts of the belief state have been fixed.

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 $C \in \{A, B\backslash A, \overline{B}\}$ 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.

Definition 4 *Input $I \in \mathcal{I}$ is* ***strongly silent on the probability of A given B*** *(for $\varnothing \subsetneq A \subsetneq B \subseteq Supp(I)$) if, for all $\alpha \in [0,1]$ and all $p^* \in I$, the set I 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 $p^*(A), p^*(B\backslash A) \neq 0$).

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

$$p' := p^*(\cdot|A)\alpha p^*(B) + p^*(\cdot|B\backslash A)(1-\alpha)p^*(B) + p^*(\cdot \cap \overline{B}), \tag{3}$$

so that the requirement that there exists some p' in I satisfying (a) and (b) reduces to the requirement that I contains the belief state (3).[11]

For example, the inputs $I = \{p' : p' \text{ is uniform on } \overline{B}\}$ and $I = \{p' : p'(B) \geq 1/2\}$ 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). This strengthened notion of silence leads to a weaker conservativeness axiom, which we call just 'conservativeness'.

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

3.3 An alternative perspective on weak and strong silence

Before stating our characterization theorem, we note that there is an alternative and equivalent way of defining weak and strong silence, which gives a different perspective on these notions. Informally, weak silence can be taken to mean that the learnt input implies nothing for the probability of A given B. Strong silence can be taken to mean that all its implications are 'outside' the probability of A given B (i.e., the input constrains only parts of the agent's belief state that are orthogonal to the probability of A given B). To make this more precise, we first define the 'implication' of a learnt

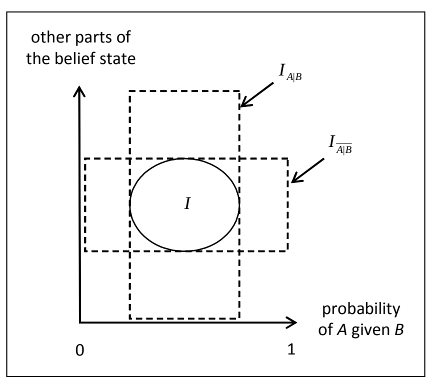

Figure 2: The implications $I_{A|B}$ and $I_{\overline{A|B}}$ derived from input I

input I for the probability of A given B and for other parts of the agent's belief state. Again, we assume that $\varnothing \subsetneq A \subsetneq B \subseteq Supp(I)$.

- The **implication of** I **for the probability of** A **given** B is the input, denoted $I_{A|B}$, which says everything that I says about the probability of A given B, and nothing else (see Figure 2). So, $I_{A|B}$ contains all belief states p'

[11] To be precise, this is true whenever $p^*(A), p^*(B\backslash A) \neq 0$.

which are compatible with I on the probability of A given B. Formally, $I_{A|B}$ is the set of all p' in $\mathcal{P}$ such that $p'(A|B) = p^*(A|B)$ for some p^* in I (modulo a non-triviality constraint).[12]

- The **implication of I outside the probability of A given B** is the input, denoted $I_{\overline{A|B}}$, which says everything that I says outside the probability of A given B, and nothing else (see Figure 2). So, $I_{\overline{A|B}}$ contains all belief states which are compatible with I outside the probability of A given B. Formally, $I_{\overline{A|B}}$ is the set of all p' in $\mathcal{P}$ which coincide with some p^* in I outside the probability of A given B (modulo a non-triviality constraint).[13]

Clearly, $I \subseteq I_{A|B}$ and $I \subseteq I_{\overline{A|B}}$. The inputs $I_{A|B}$ and $I_{\overline{A|B}}$ capture two orthogonal components ('sub-inputs') of the full input I. Each component encodes part of the information conveyed by I. Weak and strong silence can now be characterized (and thereby alternatively defined) as follows.

Proposition 2 *For all inputs $I \in \mathcal{I}$ and events A, B (where $\varnothing \subsetneq A \subsetneq B \subseteq Supp(I)$),*
(a) *I is weakly silent on the probability of A given B if and only if $I_{A|B} = \mathcal{P}$ (i.e., I implies* nothing *for the probability of A given B),*
(b) *I is strongly silent on the probability of A given B if and only if $I_{\overline{A|B}} = I$ (i.e., I implies* only something outside *the probability of A given B).*

We can illustrate this proposition by combining Figures 1 and 2. According to part (a), weak silence means that the sub-input $I_{A|B}$, which pertains to the probability of A given B, is vacuous. Graphically, it covers the entire area in the plot. According to part (b), strong silence means that the input I conveys no information beyond the sub-input $I_{\overline{A|B}}$, which pertains to those parts of the agent's belief state that are orthogonal to the probability of A given B. Graphically, the input I covers a rectangular area ranging from the far left to the far right.

3.4 The theorem

We have seen that the strong version of our conservativeness axiom, defined in terms of weak silence, leads to an impossibility result. By contrast, its weaker counterpart, defined in terms of strong silence, yields a characterization of Bayes's and Jeffrey's rules. As we will see in the next section, it also yields a structurally identical characterization of two other rules.

Theorem 1 *Bayes's and Jeffrey's rules are the only responsive and conservative belief-revision rules on the domains $\mathcal{D}_{Bayes}$ and $\mathcal{D}_{Jeffrey}$, respectively.*

Corollary 1 *Every responsive and conservative belief-revision rule on some domain $\mathcal{D} \subseteq \mathcal{P} \times \mathcal{I}$ coincides with Bayes's rule on $\mathcal{D} \cap \mathcal{D}_{Bayes}$ and with Jeffrey's rule on $\mathcal{D} \cap \mathcal{D}_{Jeffrey}$.*

[12] In full precision, $I_{A|B}$ is the set of all p' in $\mathcal{P}$ such that if $p'(B) \neq 0$ then $p'(A|B) = p^*(A|B)$ for some p^* in I satisfying $p^*(B) \neq 0$.

[13] In full precision, $I_{\overline{A|B}}$ is the set of all p' in $\mathcal{P}$ such that if there is a belief state p^* in I satisfying $[p^*(C) \neq 0$ for all $C \in \{A, B\backslash A\}$ such that $p'(C) \neq 0]$, then p' coincides with some such p^* outside the probability of A given B.

It is easier to prove that *if* a belief-revision rule on one of these domains is responsive and conservative, then it must be Bayes's or Jeffrey's rule, than to prove the converse implication, namely that each of these rules is responsive and conservative on its domain. To illustrate the easier implication, note, for instance, that if a belief-input pair (p, I) belongs to $\mathcal{D}_{\text{Bayes}}$, such as $I = \{p' : p'(B) = 1\}$, then the new belief state p_I equals $p_I(\cdot|B)$ (since $p_I(B) = 1$, by responsiveness), which equals $p(\cdot|B)$ (by conservativeness, as I is strongly silent on probabilities given B). The reason why the converse implication is harder to prove is that it is difficult to identify all the conditional probabilities on which a given input is strongly silent; there are more such conditional probabilities than one might expect. Once we have identified all those conditional probabilities, we must verify that the corresponding belief-revision rule does indeed preserve all of them, as required by conservativeness.

4 Two further belief-revision rules

We now extend our characterization to two further non-Bayesian belief-revision rules, showing that they, too, are the only responsive and conservative rules on their respective domains. They are less well known than Jeffrey's rule, but are of interest in their own right. The first, which we call the 'dual-Jeffrey' rule, stands out for its natural 'duality' to Jeffrey's rule. The second, Adams's rule, is inspired by Ernst Adams's work on the logic of conditionals (Adams 1975), although it was introduced formally by Bradley (2005). We first explain these rules informally, before representing them in our framework. We begin with the 'dual-Jeffrey' rule.

4.1 The dual-Jeffrey rule

As in our definition of Jeffrey's rule, let $\mathcal{C}$ be a partition of the set Ω into finitely many non-empty events. Now suppose that, for each event C in $\mathcal{C}$, the agent learns a new assignment of conditional probabilities to all other events, given C. More formally, the agent learns a new *conditional probability distribution given the partition* $\mathcal{C}$, i.e., a family $(\pi^C)^{C\in\mathcal{C}}$ of probability measures $\pi^C \in \mathcal{P}$, each of which has support $Supp(\pi^C) = C$. (The *support* of a probability measure p is $Supp(p) := \{\omega \in \Omega : p(\omega) \neq 0\}$.) Then π^C assigns to each event $A \subseteq \Omega$ a new conditional probability, given C. The 'dual-Jeffrey' rule says that, upon learning this new conditional probability distribution $(\pi^C)^{C\in\mathcal{C}}$, one should move from the prior probability measure p to the posterior probability measure p' given by

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

For instance, the agent might learn a new conditional probability distribution given the events of 'rain' (C) and 'no rain' ($\overline{C}$); so $\mathcal{C} = \{C, \overline{C}\}$. He might learn that, given rain, flooding is more likely than previously expected, whereas without rain, there is a greater risk of forest fire. Under the dual-Jeffrey rule, the agent would then revise these conditional probabilities accordingly, while leaving his unconditional probabilities for the events C and $\overline{C}$ unchanged.

Dual-Jeffrey revision also permits the possibility of learning some new conditional probabilities given a single event C, without learning any new conditional probabilities given its complement $\overline{C}$. In this case, the partition $\mathcal{C}$ consists of the event C and all singleton events $\{\omega\}$, where $\omega \in \overline{C}$.[14] The new conditional probabilities given C are then specified by π^C, while each $\pi^{\{\omega\}}$ is trivially given by $\pi^{\{\omega\}}(\omega) = 1$.

The duality between Jeffrey's rule and the dual-Jeffrey rule lies in the fact that, when $\mathcal{B} = \mathcal{C}$, the two rules affect complementary parts of the agent's belief state. While the former affects only unconditional probabilities of events in $\mathcal{B}$ and leaves any conditional probabilities given these events unaffected, the latter does the reverse.

4.2 Adams's rule

While the definitions of Jeffrey's rule and the dual-Jeffrey rule each involve a single partition of the set Ω into finitely many non-empty events, we now consider two such partitions, $\mathcal{B}$ and $\mathcal{C}$. Suppose that the agent learns a new conditional probability π_B^C of any event B in partition $\mathcal{B}$, given any event C in partition $\mathcal{C}$ (without learning any new unconditional probability of C or any new conditional probabilities given $B \cap C$). More formally, the agent learns a new *conditional probability distribution on $\mathcal{B}$ given $\mathcal{C}$*, i.e., a family of numbers indexed by both $\mathcal{B}$ and $\mathcal{C}$, denoted $(\pi_B^C)_{B\in\mathcal{B}}^{C\in\mathcal{C}}$, subject to two conditions. First, $\sum_{B\in\mathcal{B}} \pi_B^C = 1$ for all $C \in \mathcal{C}$; and second $\pi_B^C > (=)$ 0 whenever $B \cap C \neq (=) \varnothing$. Adams's rule (e.g., Bradley 2005, 2007, Douven and Romeijn 2011) says that, upon learning $(\pi_B^C)_{B\in\mathcal{B}}^{C\in\mathcal{C}}$, one should move from the prior probability measure p to the posterior probability measure p' given by

$$p'(A) = \sum_{B\in\mathcal{B}, C\in\mathcal{C}} p(A|B \cap C)\pi_B^C p(C) \text{ for all } A \subseteq \Omega.$$[15]

For instance, suppose 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 $\pi_{\overline{B}}^{\overline{C}} = \frac{7}{10}$. Another example is learning an equation of the form $X = f(Y) + \epsilon$, where X and Y are two 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.

Like the dual-Jeffrey rule, Adams's rule can also accommodate the case in which the agent learns only a single conditional probability, such as only the probability of rain given a 'rainy' forecast. In that case, $\mathcal{B}$ contains the events of rain (B) and no rain ($\overline{B}$), while $\mathcal{C}$ contains the event C of a 'rainy' forecast and all singleton events of the form $\{\omega\}$ for $\omega \in \overline{C}$.[16] We could then still have $\pi_B^C = \frac{9}{10}$ and $\pi_{\overline{B}}^C = \frac{1}{10}$ and define

[14] This assumes that $\Omega \backslash C$ is finite, since we require the partition $\mathcal{C}$ to be finite. Although a generalization to countable partitions is possible, we set this aside for expositional simplicity.

[15] 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 $B \cap C \neq \varnothing$ and $p(C) \neq 0$. This condition ensures that, in the present formula, 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.

[16] As in the dual-Jeffrey case, this requires $\Omega \backslash C$ to be finite.

each $\pi^{\{\omega\}}_{B'}$ ($B' \in \mathcal{B}$) to be 1 if $\omega \in B'$ and 0 if $\omega \notin B'$.

When Ω is finite, Adams's rule generalizes the dual-Jeffrey rule, which is obtained if $\mathcal{B}$ is the finest partition $\{\{a\} : a \in \Omega\}$.[17] It also 'almost' generalizes Jeffrey's rule, since if $\mathcal{C}$ is the coarsest partition $\{\Omega\}$ we obtain Jeffrey revision with learnt input $(\pi_B)_{B\in\mathcal{B}} \equiv (\pi^{\Omega}_B)_{B\in\mathcal{B}}$, where this input is not maximally general since each π_B $(= \pi^{\Omega}_B)$ is non-zero.

4.3 The two rules in our framework

We can now define and characterize these two rules in our general framework. We begin by defining the relevant inputs.

Definition 5 *A learnt input $I \in \mathcal{I}$ is*

- ***dual-Jeffrey*** *if $I = \{p' : p'(\cdot|C) = \pi^C$ for all $C \in \mathcal{C}$ with $p'(C) \neq 0\}$ for some conditional probability distribution $(\pi^C)^{C\in\mathcal{C}}$ given some partition $\mathcal{C}$; we then write $I = I_{(\pi^C)^{C\in\mathcal{C}}}$;*[18]
- ***Adams*** *if $I = \{p' : p'(B|C) = \pi^C_B$ for all $B \in \mathcal{B}$ and all $C \in \mathcal{C}$ with $p'(C) \neq 0\}$ for some conditional probability distribution $(\pi^C_B)^{C\in\mathcal{C}}_{B\in\mathcal{B}}$ on some partition $\mathcal{B}$ given some partition $\mathcal{C}$; we then write $I = I_{(\pi^C_B)^{C\in\mathcal{C}}_{B\in\mathcal{B}}}$.*[19]

As already noted, the class of Adams inputs is very general: it includes all dual-Jeffrey inputs (for finite Ω) and 'almost' all Jeffrey inputs (for details, see our earlier remarks). The two belief-revision rules can be defined as follows.

Definition 6

- *Let $\mathcal{D}_{dual\text{-}Jeffrey}$ be the set of all pairs $(p, I) \in \mathcal{P} \times \mathcal{I}$ such that $I = I_{(\pi^C)^{C\in\mathcal{C}}}$ is a dual-Jeffrey input (every dual-Jeffrey input is 'compatible' with p). The **dual-Jeffrey rule** is the revision rule on $\mathcal{D}_{dual\text{-}Jeffrey}$ which maps each $(p, I_{(\pi^C)^{C\in\mathcal{C}}}) \in \mathcal{D}_{dual\text{-}Jeffrey}$ to $p' \in \mathcal{P}$, where*

$$p'(A) = \sum_{C\in\mathcal{C}} \pi^C(A)p(C) \text{ for all events } A \subseteq \Omega. \quad (4)$$

- *Let $\mathcal{D}_{Adams}$ be the set of all pairs $(p, I) \in \mathcal{P} \times \mathcal{I}$ such that $I = I_{(\pi^C_B)^{C\in\mathcal{C}}_{B\in\mathcal{B}}}$ is an Adams input compatible with p (which means $p(B \cap C) \neq 0$ for all $B \in \mathcal{B}$ and $C \in \mathcal{C}$ such that $B \cap C \neq \varnothing$ and $p(C) \neq 0$). **Adams's rule** is the revision rule on $\mathcal{D}_{Adams}$ which maps each $(p, I_{(\pi^C_B)^{C\in\mathcal{C}}_{B\in\mathcal{B}}}) \in \mathcal{D}_{Adams}$ to $p' \in \mathcal{P}$, where*

$$p'(A) = \sum_{B\in\mathcal{B}, C\in\mathcal{C}} p(A|B \cap C)\pi^C_B p(C) \text{ for all } A \subseteq \Omega. \quad (5)$$

[17] In principle, we could define Adams's rule more generally, by allowing the relevant partitions of Ω to be countable rather than finite. For expositional simplicity, we set this generalization aside here.

[18] The representation $I = I_{(\pi^C)^{C\in\mathcal{C}}}$ is unique, because for any dual-Jeffrey input I, there exists a unique family $(\pi^C)^{C\in\mathcal{C}}$ such that $I = I_{(\pi^C)^{C\in\mathcal{C}}}$.

[19] For any Adams input I, there are multiple families of the form $(\pi^C_B)^{C\in\mathcal{C}}_{B\in\mathcal{B}}$ such that $I = I_{(\pi^C_B)^{C\in\mathcal{C}}_{B\in\mathcal{B}}}$ (though we show in Lemma 5 in the Appendix that one family stands out as canonical). We return to the implications of this non-uniqueness in footnote 20.

In analogy to our earlier definitions, the domains $\mathcal{D}_{\text{dual-Jeffrey}}$ and $\mathcal{D}_{\text{Adams}}$ are the maximal domains for which formulas (4) and (5) are well-defined.[20] Of course, the notion of silence and our two axioms can be applied to the present two domains as well. For example,

- a dual-Jeffrey input is not silent on the conditional probabilities, given each event in the relevant partition $\mathcal{C}$, but is silent on the unconditional probabilities of the events in $\mathcal{C}$;
- an Adams input is silent on the unconditional probabilities of the events in partition $\mathcal{C}$ and on conditional probabilities given events from the join of partitions $\mathcal{B}$ and $\mathcal{C}$.

We can finally extend our main theorem to the domains of dual-Jeffrey and Adams inputs.

Theorem 2 *The dual-Jeffrey rule and Adams's rule are the only responsive and conservative belief-revision rules on the domains $\mathcal{D}_{\text{dual-Jeffrey}}$ and $\mathcal{D}_{\text{Adams}}$, respectively.*

Corollary 2 *Every responsive and conservative belief-revision rule on some domain $\mathcal{D} \subseteq \mathcal{P} \times \mathcal{I}$ coincides with the dual-Jeffrey rule on $\mathcal{D} \cap \mathcal{D}_{\text{dual-Jeffrey}}$ and with Adams's rule on $\mathcal{D} \cap \mathcal{D}_{\text{Adams}}$.*

As before, it is easier to prove that *if* a belief-revision rule on $\mathcal{D}_{\text{dual-Jeffrey}}$ or $\mathcal{D}_{\text{Adams}}$ is responsive and conservative, then it must be the dual-Jeffrey rule or Adams's rule, than to prove the converse, namely that each of these rules is responsive and conservative on its domain. The main challenge, again, is to identify all the conditional probabilities on which a dual-Jeffrey input or an Adams input is strongly silent, in order to establish conservativeness. The present section completes our unified characterization of four distinct belief-revision rules.

5 A decision-theoretic application

To show that there is room for non-Bayesian belief-revision rules in economic theory, we now sketch an illustrative application to decision and game theory. Standard dynamic decision and game theory is inherently Bayesian. As is widely recognized, this sometimes entails unrealistic assumptions of forward-looking rationality, which limit the ability to model real-life learning, reasoning, and behaviour. We give an example which illustrates some of these difficulties and shows how a non-Bayesian model can avoid them. The example suggests a new class of dynamic decision problems or games, those 'with surprises' or 'with unforeseen learning inputs'.

Ann, an employer, must decide whether to hire Bob, a job candidate. There is no time for a job interview, since a quick decision is needed. Ann is uncertain about

[20] The definition of the dual-Jeffrey rule and its domain relies on the fact that each dual-Jeffrey input I is uniquely representable as $I_{(\pi^C)^{C \in \mathcal{C}}}$. Although the representation of any Adams input as $I = I_{(\pi^C_B)^{C \in \mathcal{C}}_{B \in \mathcal{B}}}$ is far from unique, Adams's rule is nonetheless well-defined. This is because the revision formula (5) and the domain definition are invariant under the choice of family $(\pi^C_B)^{C \in \mathcal{C}}_{B \in \mathcal{B}}$ representing I. This non-trivial fact is shown in the Appendix, where we also give a characterization of the families $(\pi^C_B)^{C \in \mathcal{C}}_{B \in \mathcal{B}}$ representing a given Adams input. See Lemmas 2 and 4 for details.

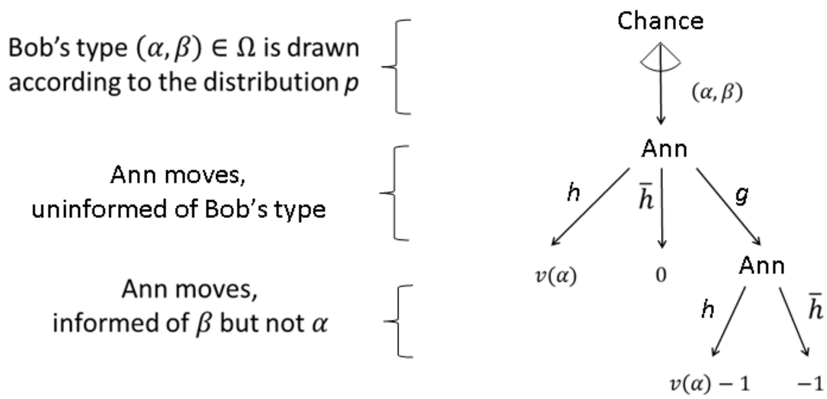

Figure 3: Ann's decision problem in its initial form

whether Bob is competent or not; both possibilities have prior probability $\frac{1}{2}$. It would help Ann to know whether Bob has previous work experience, since this is positively correlated with competence, but gathering this information takes time. Bob's type is thus a pair (α, β) whose first component indicates whether he is competent ($\alpha = c$) or not ($\alpha = \bar{c}$) and whose second component indicates whether he has work experience ($\beta = e$) or not ($\beta = \bar{e}$). To apply a belief-revision model, let the set of worlds be the set of possible types of Bob, i.e., $\Omega = \{c, \bar{c}\} \times \{e, \bar{e}\}$. Ann's initial beliefs about Bob's type are given by the belief state $p \in \mathcal{P}$ in which $p(c, e) = p(\bar{c}, \bar{e}) = 0.4$ and $p(c, \bar{e}) = p(\bar{c}, e) = 0.1$. Note the positive correlation between competence and work experience.

Ann initially seems to face the dynamic decision problem shown in Figure 3:

- First, a chance move determines Bob's type in Ω according to the probability measure p.
- Next, Ann, uninformed of the chance move, can hire Bob (h) or reject him ($\bar{h}$) or gather information about whether he has previous work experience (g).
- Finally, if Ann chooses g, she faces a subsequent choice between hiring Bob (h) or rejecting him ($\bar{h}$), but now she has information about β, i.e., about whether he has work experience.

Ann is an expected-utility maximizer, and her utility function is as follows: hiring Bob, who is of type (α, β), contributes an amount $v(\alpha)$ to her utility, where $v(\alpha) = 5$ if $\alpha = c$ and $v(\alpha) = -5$ if $\alpha = \bar{c}$; and gathering information about β reduces her utility by 1. Not hiring Bob yields a utility of 0. Ann has only one rational strategy: first she gathers information (g), and then she hires Bob if and only if she learns that Bob has work experience ($\beta = e$). To see why, note the following. Immediately hiring Bob yields an expected utility of $v(c)p(c) + v(\bar{c})p(\bar{c}) = 5\frac{1}{2} + (-5)\frac{1}{2} = 0$. Immediately rejecting Bob also yields an expected utility of 0. Gathering information leads Ann to a Bayesian belief revision:

- If she learns that he has work experience, she raises her probability that he is competent to $p(c|e) = \frac{p(c,e)}{p(e)} = \frac{0.4}{0.5} = \frac{4}{5}$. So, she hires Bob, since this yields an expected utility of $(v(c)-1)p(c|e)+(v(\overline{c})-1)p(\overline{c}|e) = 4\frac{4}{5}+(-6)\frac{1}{5} = 2$, while rejecting Bob would have yielded an expected utility of -1.

- If she learns that Bob has no work experience, she lowers her probability that he is competent to $p(c|\overline{e}) = \frac{p(c,\overline{e})}{p(\overline{e})} = \frac{0.1}{0.5} = \frac{1}{5}$. So, she rejects him as this yields an expected utility of -1, whereas hiring him would have yielded an expected utility of $(v(c)-1)p(c|\overline{e})+(v(\overline{c})-1)p(\overline{c}|\overline{e}) = 4\frac{1}{5}+(-6)\frac{4}{5} = -4$.

So *ex ante* the expected utility of gathering information is the average $2p(e) + (-1)p(\overline{e}) = 2\frac{1}{2} + (-1)\frac{1}{2} = \frac{1}{2}$, which exceeds the zero expected utility of the two other choices.

So far, everything is classical. Now suppose Ann follows her rational strategy. She writes to Bob to ask whether he has work experience. At this point, however, something surprising happens. Bob's answer reveals right from the beginning that his written English is poor. Ann notices this even before figuring out what Bob says about his work experience. In response to this unforeseen learnt input, Ann lowers her probability that Bob is competent from $\frac{1}{2}$ to $\frac{1}{8}$. It is natural to model this as an instance of Jeffrey revision. Formally, Ann learns the Jeffrey input $I = \{p' \in \mathcal{P} : p'(c) = \frac{1}{8}\}$, and by Jeffrey's rule her revised belief state p_I is given by $p_I(c,e) = \frac{1}{10}$, $p_I(c,\overline{e}) = \frac{1}{40}$, $p_I(\overline{c},e) = \frac{7}{40}$, and $p_I(\overline{c},\overline{e}) = \frac{7}{10}$. As she reads the rest of Bob's letter, Ann eventually learns that he has previous work experience, which prompts a Bayesian belief revision, so that her final belief state is $p_I(\cdot|e)$ (or equivalently, $(p_I)_{I'}$ where I' is the Bayesian input $I' = \{p' \in \mathcal{P} : p'(e) = 1\}$). Since Ann's posterior probability for Bob's competence is only $p_I(c|e) = \frac{p_I(c,e)}{p_I(c,e)+p_I(\overline{c},e)} = \frac{4}{11}$, she decides not to hire him, despite his work experience.

Can classical decision theory explain this? Of course, the dynamic decision problem shown in Figure 3 is no longer adequate, as it wrongly predicts that Ann hires Bob after learning that he has work experience. The natural response, from a classical perspective, would be to refine the decision problem as shown in Figure 4(a). After Ann's information-gathering move g an additional chance move is introduced, which determines whether Bob's written English is normal (w) or poor ($\overline{w}$), where the probability of w, denoted $t_{\alpha,\beta}$, is larger if Bob is competent than if he is not, i.e., $t_{c,\beta} > t_{\overline{c},\beta}$. After observing this chance move, Ann makes her hiring decision.

Although this refined classical model predicts that Ann turns down Bob after receiving his poorly written letter, it is inadequate in many ways. It ignores the fact that Ann is initially unaware of – or does not consider – the possibility that Bob's written English is poor (suppose, for instance, that based on her initial information, she had no reason to doubt, or even to think about, Bob's literacy). It treats the event of a poorly written letter from Bob as a *foreseen* rather than an *unforeseen* contingency. As a result, Ann's reasoning at each of her decision nodes is modelled in an inadequate manner:

(i) In her first decision (between h, $\overline{h}$, and g), Ann is falsely taken to foresee the possibilities of learning w or learning $\overline{w}$, i.e., to reason along the tree displayed in Figure 4(a) rather than that in Figure 3. This artificially complicates her expected-utility maximization exercise, for instance by assuming awareness of the four parameters $t_{\alpha,\beta}$, for all values of α and β.

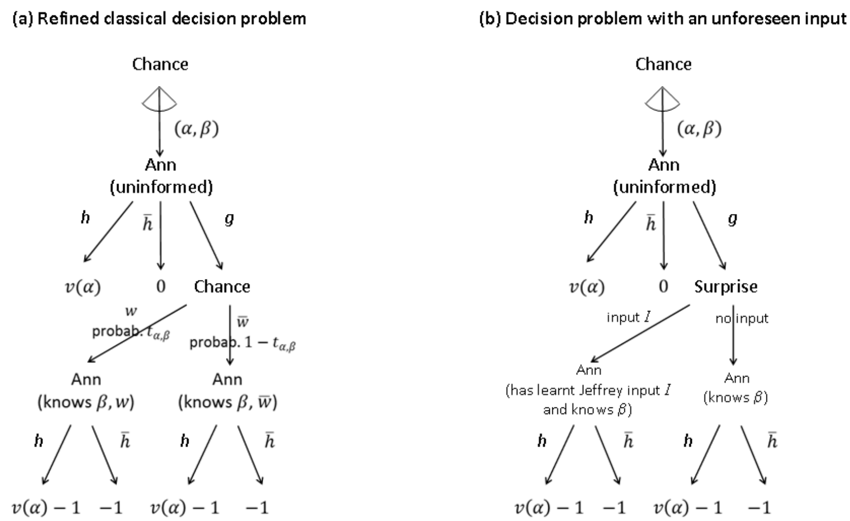

Figure 4: Two ways of refining Ann's decision problem

(ii) In her second decision, *in case Bob's written English is normal*, Ann is taken to have learnt not just the parameter β, but also the chance move w (normal written English), so that her posterior probability for Bob's competence is now $p(c|\beta, w)$ rather than $p(c|\beta)$. The additional conditionalization on w misrepresents Ann's beliefs, since the absence of linguistic errors in Bob's letter goes unnoticed: it is not an unforeseen event (she had taken Bob's normal literacy for granted). In fact, Ann continues to conceptualize her decision problem as the one shown in Figure 3 rather than the one in 4(a). The additional belief revision (upon learning w) departs from, and complicates, Ann's true reasoning.

(iii) In her second decision, *in case Bob's written English is poor*, Ann's reasoning is again misrepresented. Although it is true that the unforeseen news that Bob's written English is poor implies that Ann cannot uphold her original conceptualization of the decision problem (Figure 3), it does not follow that Ann re-conceptualizes her decision problem in line with Figure 4(a). Our informal description of Ann's reasoning takes her to perform a Jeffrey revision of her beliefs over $\Omega = \{c, \bar{c}\} \times \{e, \bar{e}\}$, whereas Figure 4(a) takes her to perform a Bayesian revision of beliefs over the refined set of worlds $\Omega' = \{c, \bar{c}\} \times \{e, \bar{e}\} \times \{w, \overline{w}\}$.

Arguably, the Bayesian model of Ann's behaviour is not only psychologically inadequate, but its predictive adequacy is also far from clear. Whether the model correctly predicts Ann's behaviour at the various decision nodes depends on the exact calibration of the parameters $t_{\alpha,\beta}$, for all values of α and β. Their most plausible (e.g., 'objective') values might not imply Ann's true decision behaviour, since that behaviour has a rather different psychological origin, which does not involve the parameters $t_{\alpha,\beta}$ at all.

We propose to model Ann's decision problem non-classically as a decision problem *with unforeseen inputs* or *surprises*. As illustrated in Figure 4(b), instead of introducing a chance move (selecting w or $\overline{w}$), we introduce a *surprise move*, which determines whether or not Ann receives a particular unforeseen input (here, the Jeffrey input $I = \{p' \in \Omega : p'(c) = \frac{1}{8}\}$). Then the problems in (i), (ii), and (iii) no longer arise:

- Problem (i) is avoided because Ann does not foresee or conceptualize the surprise move before its occurrence, so that she initially still reasons along the simple decision tree of Figure 3.
- Problem (ii) is avoided because in her second decision, without receiving the unforeseen input (the right branch at the surprise node), Ann only learns β and hence reasons like in her second decision in the simple decision problem of Figure 3.
- Problem (iii) is avoided because in her second decision, after receiving the unforeseen input (the left branch at the surprise node), Ann revises her beliefs in response to the Jeffrey input I.

In follow-up work, we formally define decision problems (or more generally games) *with unforeseen inputs* and introduce a corresponding equilibrium notion.[21] The details are beyond the scope of the present paper. Our aim in this section has simply been to illustrate that there is useful room for non-Bayesian belief-revision rules in a decision-theoretic model.

6 Concluding remarks

We have developed a unified framework for the study of belief revision and shown that Bayes's rule as well as three salient non-Bayesian alternatives can be characterized in terms of the same two axioms: responsiveness to the learnt input and conservativeness. The only difference between the four rules lies in the domain of learnt inputs to which they apply. Previous characterizations of Bayes's, Jeffrey's, and Adams's rules tended to be less unified. They typically characterized different rules either in terms of different axioms or as 'distance-minimizing' with respect to different notions of distance between probability measures.

Beyond offering a novel formal framework, the programmatic aim of this paper has been to put non-Bayesian belief revision onto the map for economic theorists. No doubt, skeptics will still wonder, 'why bother about non-Bayesian belief revision'. By suitably refining the set Ω of possible worlds, so the objection goes, we can always remodel Jeffrey, dual-Jeffrey, and Adams inputs in a Bayesian manner. However, as we have noted in our discussions of Suzumura's story and 'Ann the employer', such Bayesian remodelling comes at a cost:

- **Over-ascription of opinionation:** A key drawback of the Bayesian remodelling is that we must assume that the agent is able to assign prior probabilities to many complex events. In Suzumura's story, the agent must assign assign prior

[21] Ann's equilibrium strategy in her decision problem with unforeseen inputs has the intended form: she first gathers information (g), and then hires Bob if and only if she does not receive the unforeseen input I (Bob's poor written English) and learns that he has work experience ($\beta = e$).

probabilities to the various possible auditory signals that he might receive over the phone. Similarly, Ann the employer must assign prior probabilities to the various possible Jeffrey inputs she might receive. These may include not only learning that Bob's written English is poor, but also that he is a poetic writer, that he comes across as communicatively awkward in a way that she had not anticipated, and so on. To accommodate the possibility of belief changes in response to such inputs, we would have to ascribe to the agent beliefs over an ever 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 either have no beliefs about such events or have only imprecise ones. Even on a pure 'as-if' interpretation of the Bayesian model, taking an agent to have highly sophisticated beliefs is dubious, given the complexity of their behavioural implications, which may be hard to test empirically. By contrast, as soon as we restrict the complexity of the event-algebra, we may have to invoke non-Bayesian belief revision to capture the agent's belief dynamics adequately.

- **Over-ascription of awareness:** 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 or 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 as we have noted, it is far from clear whether, prior to the telephone conversation with Gorman, Suzumura even considered the possibility of receiving incomprehensible auditory signals, or whether Ann the employer would have considered the possibility that Bob's written English was so poor. In these examples, the agents plausibly lacked not only *knowledge* but also *awareness* of the 'surprise events'. Arguably, many real-life belief changes involve the observation or experience of something that was previously not just unknown, but even beyond awareness or imagination.

In sum, an economic modeller often faces a choice between

(i) ascribing to an agent *Bayesian revision* of beliefs *over a very complex, fine-grained algebra of events* and

(ii) ascribing *non-Bayesian revision* of beliefs *over a simpler, more coarse-grained algebra of events.*

Perhaps because of the elegance of Bayes's rule, many economists tend to assume that the first of these routes is more parsimonious than the second. But this overlooks the loss of parsimony at the level of the event-algebra. If all non-Bayesian belief-revision rules were *ad hoc* or otherwise unsatisfactory, the choice of route (i) would be understandable. But as we have shown, there are perfectly well-behaved non-Bayesian alternatives. This should make option (ii) at least a contender worth taking seriously.

7 References

Adams, E. (1975) *The Logic of Conditionals*. Dordrecht and Boston: Reidel.

Bradley, R. (2005) Radical Probabilism and Bayesian Conditioning, *Philosophy of Science* 72(2): 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 on Information Theory*: 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. (2012) Modelling change in individual characteristics: an axiomatic framework, *Games and Economic Behavior* 76(2): 471-494.

Douven, I., Romeijn, J. W. (2011) A new resolution of the Judy Benjamin Problem, *Mind* 120(479): 637-670.

Grünwald, P., Halpern, J. (2003) Updating probabilities, *Journal of AI Research* 19: 243-278.

Halpern, J. (2003) *Reasoning About Uncertainty*, MIT Press, Cambridge, MA.

Heifetz, A., Meier, M. and B. C. Schipper (2006). Interactive unawareness, *Journal of Economic Theory*, 130: 78-94.

Jeffrey, R. (1957) Contributions to the theory of inductive probability, *PhD Thesis*, Princeton University.

Modica, S., Rustichini, A. (1999) Unawareness and partitional information structures, *Games and Economic Behavior* 27: 265-298.

Shafer, G. (1981) Jeffrey's rule of conditioning, *Philosophy of Science* 48: 337-362.

van Fraassen, B. C. (1981) A Problem for Relative Information Minimizers in Probability Kinematics, *British Journal for the Philosophy of Science* 32: 375–379.

Vineberg, S. (2011) Dutch Book Arguments, *Stanford Encyclopedia of Philosophy* (Summer 2011 Edition).

A Appendix: Proofs

Notationally, for all $a \in \Omega$, let $\delta_a \in \mathcal{P}$ be the *Dirac measure in* a, defined by $\delta_a(a) = 1$.

A.1 Well-definedness of each revision rule

As mentioned, each of our four belief-revision rules is well-defined because the mathematical object used in the definition of the new belief state and the rule's domain (i.e., the learnt event B or the learnt families (π_B), $(\pi^C)^{C \in \mathcal{C}}$, or $(\pi_B^C)_{B \in \mathcal{B}}^{C \in \mathcal{C}}$) is either *uniquely* determined by the relevant input I or at least *sufficiently* determined so that the definition does not depend on non-unique features. This fact deserves a proof. For Bayes's, Jeffrey's, and the dual-Jeffrey rules, the proof is trivial and given by the following three lemmas (which the reader can easily verify):

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

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

Lemma 3 *For every Jeffrey input I,*
(a) all families $(\pi_B)_{B\in\mathcal{B}}$ generating I 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 (i) defined and identical for all families $(\pi_B)_{B\in\mathcal{B}}$ generating I, or (ii) undefined for all these families.[22]

The well-definedness of Adams's rule is harder to establish. We begin with a lemma which characterizes the common features of all families $(\pi_B^C)_{B\in\mathcal{B}}^{C\in\mathcal{C}}$ generating a given Adams input I. This determines the extent to which the generating family $(\pi_B^C)_{B\in\mathcal{B}}^{C\in\mathcal{C}}$ is unique. We show that the partition $\mathcal{C}$ is *essentially unique*: its members are fixed, with the exception of 'trivial' members. A member C of $\mathcal{C}$ is 'trivial' if it is included in some member of $\mathcal{B}$, which implies that any π_B^C $(B \in \mathcal{B})$ must be 1 (if $C \subseteq B$) or 0 (if $B \cap C = \varnothing$). We also show that, while the partition $\mathcal{B}$ is far from unique, the join of partitions

$$\mathcal{B} \vee \mathcal{C} = \{B \cap C : B \in \mathcal{B}, C \in \mathcal{C}\} \backslash \{\varnothing\}$$

is *essentially unique*. On a first reading of the lemma, one may assume that $\mathcal{C}$ contains no trivial members (so that $\mathcal{C}_{\text{triv}} = \varnothing$). In this case, the lemma states that $\mathcal{C}$ and $\mathcal{B} \vee \mathcal{C}$ are fully unique.

Lemma 4 *Let I be an Adams input. All families $(\pi_B^C)_{B\in\mathcal{B}}^{C\in\mathcal{C}}$ generating I have*
(a) the same set $\mathcal{C}\backslash\mathcal{C}_{triv}$, where $\mathcal{C}_{triv} := \{C \in \mathcal{C} : C$ is a subset of some $B \in \mathcal{B}\}$,
(b) the same set $(\mathcal{B} \vee \mathcal{C})\backslash\mathcal{C}_{triv}$, where $\mathcal{C}_{triv}$ is defined as in part (a),
(c) for each $a \in \Omega$, the same value $\pi_{B_a}^{Ca}$, where B_a (resp. C_a) denotes the member of $\mathcal{B}$ (resp. $\mathcal{C}$) containing a.

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

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

[22] Footnote 1 specifies when (2) is defined.

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

Let $a \in \Omega$. To show that $\min_{p' \in I : p'(a) \neq 0} |Supp(p')| \geq |\{B \in \mathcal{B} \vee \mathcal{C} : B \subseteq C_a\}|$, consider any $p' \in I$ 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 to 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 $B' \in \mathcal{B}$. Note that $p' \in I$ and $p'(C_a) \neq 0$. So, we have $p'(B'|C_a) = \pi_{B'}^{C_a} \neq 0$; and therefore $p'(B' \cap C_a) \neq 0$, i.e., $p'(B) \neq 0$.

To show the converse inequality

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

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

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

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 $B \in \mathcal{B} \vee \mathcal{C}$ with $B \subseteq C_a$ (where B' again stands for the set in $\mathcal{B}$ such that $B = B' \cap C_a$). □

In the rest of this proof, for all $a \in \Omega$, we let I^a be the set of all $p' \in I$ such that $Supp(p')$ is minimal (with respect to set inclusion), subject to $p'(a) \neq 0$.

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

Let $a \in \Omega$. The claim follows from the fact that, as the reader may verify, I^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 exists $a_B \in B$ 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 $B = B' \cap C_a$. □

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

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 the definition of I^a. □

In the following, for each $a \in \Omega$ with $D_a \neq C_a$ (i.e., with $C_a \nsubseteq B_a$), let $c(a)$ be a fixed element of $C_a \backslash D_a$.

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

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

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

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

$$C_a = \left[(\cup_{p' \in I^a} Supp(p')) \backslash \{a\} \right] \cup \left[(\cup_{p' \in I^{c(a)}} Supp(p')) \backslash \{c(a)\} \right].$$

Since

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

it follows that

$$\begin{aligned} C_a &= (\cup_{p' \in I^a} Supp(p')) \cup (\cup_{p' \in I^{c(a)}} Supp(p')) \\ &= \cup_{p' \in I^a \cup I^{c(a)}} Supp(p'). \ \square \end{aligned}$$

Claim 5: We have

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

(which proves part (a), since $\mathcal{C} \backslash \mathcal{C}_{\text{triv}}$ depends on I 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,

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

This establishes the present claim, by Claim 4. $\square$

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

$$D_a = \left[\cup_{p' \in I^{c(a)}} Supp(p') \right] \Big\backslash \left[\cup_{p' \in I^a} Supp(p') \backslash \{a\} \right].$$

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

$$D_a = \left[\cup_{p' \in I^a \cup I^{c(a)}} Supp(p') \right] \Big\backslash \left[\cup_{p' \in I^a} Supp(p') \backslash \{a\} \right].$$

It is clear that we can replace '$I^a \cup I^{c(a)}$' with '$I^{c(a)}$' without changing the resulting set D_a. $\square$

Claim 7: We have

$$\begin{aligned} (\mathcal{B} \vee \mathcal{C}) \backslash \mathcal{C}_{\text{triv}} &= \Big\{ \left[\cup_{p' \in I^{c(a)}} Supp(p') \right] \Big\backslash \\ &\qquad \left[\cup_{p' \in I^a} Supp(p') \backslash \{a\} \right] : a \in \Omega, \delta_a \notin I \Big\} \end{aligned}$$

(which proves part (b), since $(\mathcal{B} \vee \mathcal{C}) \backslash \mathcal{C}_{\text{triv}}$ depends on I 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

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

The present claim now follows from Claim 6. $\square$

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

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

$$\mathcal{C} \backslash \mathcal{C}_{\text{triv}} = \widetilde{\mathcal{C}} \backslash \widetilde{\mathcal{C}}_{\text{triv}}; \tag{6}$$

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

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

$$\cup_{C\in\mathcal{C}_{\text{triv}}}C=\cup_{C\in\widetilde{\mathcal{C}}_{\text{triv}}}C. \tag{8}$$

We distinguish between two cases.

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

Case 2: a does not belong to a set in $\mathcal{C}_{\text{triv}}$, or equivalently, to a set in $\widetilde{\mathcal{C}}_{\text{triv}}$. We deduce firstly, using (6), that a belongs to a set in $\mathcal{C}\backslash\mathcal{C}_{\text{triv}}=\widetilde{\mathcal{C}}\backslash\widetilde{\mathcal{C}}_{\text{triv}}$, so that $C_a=\widetilde{C}_a$; and secondly, using (7), that a belongs to a set in $(\mathcal{B}\vee\mathcal{C})\backslash\mathcal{C}_{\text{triv}}=(\widetilde{\mathcal{B}}\vee\widetilde{\mathcal{C}})\backslash\widetilde{\mathcal{C}}_{\text{triv}}$, so that $D_a=\widetilde{D}_a$. Choose any p' in I such that $p'(C_a)\neq 0$ (of course there is such a p' in I). Then, as the families (π_B^C) and $(\widetilde{\pi}_{\widetilde{B}}^{\widetilde{C}})$ both generate I, we have $p'(B_a|C_a)=\pi_{B_a}^{C_a}$ and $p'(\widetilde{B}_a|\widetilde{C}_a)=\widetilde{\pi}_{\widetilde{B}_a}^{\widetilde{C}_a}$. So, it suffices to show that $p'(B_a|C_a)=p'(\widetilde{B}_a|\widetilde{C}_a)$, i.e., that $p'(B_a\cap C_a)/p'(C_a)=p'(\widetilde{B}_a\cap\widetilde{C}_a)/p'(\widetilde{C}_a)$, or equivalently, that $p'(D_a)/p'(C_a)=p'(\widetilde{D}_a)/p'(\widetilde{C}_a)$. This holds because $D_a=\widetilde{D}_a$ and $C_a=\widetilde{C}_a$. ■

The next lemma shows that, among the families representing a given Adams input I, one stands out as canonical.

Lemma 5 *Let I be an Adams input. Among all families $(\pi_B^C)_{B\in\mathcal{B}}^{C\in\mathcal{C}}$ generating I, there is exactly one 'canonical' family 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) $\mathcal{B}$ and $\mathcal{C}$ have at most one event in common.[23]

Condition (a) on the family – more precisely, on the partitions $\mathcal{B}$ and $\mathcal{C}$ – is the key requirement. Essentially, it requires a fine-grained choice of $\mathcal{B}$. Starting with an arbitrary family $(\pi_B^C)_{B\in\mathcal{B}}^{C\in\mathcal{C}}$ generating I, one can ensure the satisfaction of condition (a) by refining $\mathcal{B}$, i.e., by replacing each $B\in\mathcal{B}$ with all non-empty set(s) of the form $B\cap C$, where $C\in\mathcal{C}$. Condition (b) is simply 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 $\mathcal{B}\cap\mathcal{C}$, 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 I be an Adams input.

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

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

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

[23] Given condition (a), we can restate condition (b) equivalently as follows: $\pi_B^C=1$ for at most one pair of events $B\in\mathcal{B}$ and $C\in\mathcal{C}$.

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 $\widehat{\mathcal{C}}$ is defined as $\mathcal{C}$ if $\mathcal{C}_{\text{triv}} = \varnothing$, while otherwise it is derived from $\mathcal{C}$ by replacing the trivial part with a single set:

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

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

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

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

$$\widehat{\pi}_B^C := \begin{cases} \pi_{B'}^C & \text{if } B \subsetneq C \text{ (so that } C \in \mathcal{C} \backslash \mathcal{C}_{\text{triv}}\text{), where } B' \text{ is the set in } \mathcal{B} \text{ with } B \subseteq B' \\ 1 & \text{if } B = C \text{ (so that } B = C = \cup_{C' \in \mathcal{C}_{\text{triv}}} C'\text{)} \\ 0 & \text{if } B \cap C = \varnothing. \end{cases}$$

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

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

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

coincide.

First, let $p' \in I$. To show that $p' \in \widehat{I}$, 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 $p' \in I$.
- If $B = C$, then $p'(B|C) = \widehat{\pi}_B^C$, since $p'(B|C) = 1$ and $\widehat{\pi}_B^C = 1$.
- If $B \cap C = \varnothing$, then $p'(B|C) = \widehat{\pi}_B^C$, since $p'(B|C) = 0$ and $\widehat{\pi}_B^C = 0$.

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

- If $C \backslash B, C \cap B \neq \varnothing$, 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 $p' \in \widehat{I}$.
- If $C \backslash B = \varnothing$ (i.e., $C \subseteq B$), then $p'(B|C) = \pi_B^C$, since $p'(B|C) = 1$ and $\pi_B^C = 1$.
- If $B \cap C = \varnothing$, then $p'(B|C) = \pi_B^C$, since $p'(B|C) = 0$ and $\widehat{\pi}_B^C = 0$. □

Part 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,

$$\mathcal{C}\backslash\mathcal{C}_{\text{triv}} = \widetilde{\mathcal{C}}\backslash\widetilde{\mathcal{C}}_{\text{triv}}, \tag{10}$$

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

$$\pi_{B_a}^{Ca} = \widetilde{\pi}_{\widetilde{B}_a}^{\widetilde{C}_a} \text{ for all } a\in\Omega, \tag{12}$$

where, for each $a \in \Omega$, the set B_a (and C_a, $\widetilde{B}_a$, $\widetilde{C}_a$, respectively) denotes the member of $\mathcal{B}$ (and of $\mathcal{C}$, $\widetilde{\mathcal{B}}$, $\widetilde{\mathcal{C}}$, respectively) 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

$$\mathcal{B}\backslash\mathcal{C}_{\text{triv}} = \widetilde{\mathcal{B}}\backslash\widetilde{\mathcal{C}}_{\text{triv}}. \tag{13}$$

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

$$\mathcal{C}=\widetilde{\mathcal{C}} \text{ and } \mathcal{B}=\widetilde{\mathcal{B}}. \tag{14}$$

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

We are now ready to prove that Adams's rule is well-defined.

Lemma 6 *For every Adams input I and every (initial) belief state $p\in\mathcal{P}$, the (revised) belief state (5) is either (i) defined and identical for all families $(\pi_B^C)_{B\in\mathcal{B}}^{C\in\mathcal{C}}$ generating I, or (ii) undefined for all these families.*[24]

Proof. Let I be an Adams input 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 I.

Claim 1: Expression (5) is defined for either (i) *every* or (ii) *no* family in Π.

Consider two families $(\pi_B^C)_{B\in\mathcal{B}}^{C\in\mathcal{C}}$ and $(\widetilde{\pi}_{\widetilde{B}}^{\widetilde{C}})_{\widetilde{B}\in\widetilde{\mathcal{B}}}^{\widetilde{C}\in\widetilde{\mathcal{C}}}$ in Π. By footnote 4.2, we have to show that

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

if and only if

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

We assume (15) and prove (16). The converse implication holds analogously. To prove (16), consider any $\widetilde{B}\in\widetilde{\mathcal{B}}$ and $\widetilde{C}\in\widetilde{\mathcal{C}}$ such that $\widetilde{B}\cap\widetilde{C}\neq\varnothing$ and $p(\widetilde{C})\neq 0$. We have to show that $p(\widetilde{B}\cap\widetilde{C})\neq 0$. We assume, without loss of generality, that $\widetilde{C}\not\subseteq\widetilde{B}$,

[24] Footnote 4.2 specifies when (5) is defined.

since otherwise trivially $p(\widetilde{B}\cap\widetilde{C})=p(\widetilde{C})\neq 0$. Again, let $\mathcal{C}_{\text{triv}}$ ($\widetilde{\mathcal{C}}_{\text{triv}}$) be the set of sets in $\mathcal{C}$ ($\widetilde{\mathcal{C}}$) included in a set in $\mathcal{B}$ ($\widetilde{\mathcal{B}}$). As $\widetilde{C}\not\subseteq\widetilde{B}$ and $\widetilde{B}\cap\widetilde{C}\neq\varnothing$, we have $\widetilde{C}\notin\widetilde{\mathcal{C}}_{\text{triv}}$. So, since $\mathcal{C}\backslash\mathcal{C}_{\text{triv}}=\widetilde{\mathcal{C}}\backslash\widetilde{\mathcal{C}}_{\text{triv}}$ by Lemma 4, we have $\widetilde{C}\in\mathcal{C}$. Moreover, since $\widetilde{\mathcal{C}}_{\text{triv}}$ does not contain $\widetilde{C}$, it also does not contain any subset of $\widetilde{C}$, so that $\widetilde{B}\cap\widetilde{C}\notin\widetilde{\mathcal{C}}_{\text{triv}}$. Hence, $\widetilde{B}\cap\widetilde{C}\in(\widetilde{\mathcal{B}}\vee\widetilde{\mathcal{C}})\backslash\widetilde{\mathcal{C}}_{\text{triv}}$. As we have $(\mathcal{B}\vee\mathcal{C})\backslash\mathcal{C}_{\text{triv}}=(\widetilde{\mathcal{B}}\vee\widetilde{\mathcal{C}})\backslash\widetilde{\mathcal{C}}_{\text{triv}}$ (by Lemma 4), it follows that $\widetilde{B}\cap\widetilde{C}\in\mathcal{B}\vee\mathcal{C}$. Thus there exist (unique) $B\in\mathcal{B}$ and $C\in\mathcal{C}$ such that $\widetilde{B}\cap\widetilde{C}=B\cap C$. Since $\widetilde{C}\in\mathcal{C}$, we have $C=\widetilde{C}$. Using the fact that $p(C)=p(\widetilde{C})\neq 0$ and that $B\cap C=\widetilde{B}\cap\widetilde{C}\neq\varnothing$, we have $p(B\cap C)\neq 0$ by (15), i.e., $p(\widetilde{B}\cap\widetilde{C})\neq 0$. □

Claim 2: The revised belief state (5) 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 $(\widehat{\pi}_{\widehat{B}}^{\widehat{C}})_{\widehat{B}\in\widehat{\mathcal{B}}}^{\widehat{C}\in\widehat{\mathcal{C}}}$ be two families in Π for which the revised belief state is defined. We write p' and $\widehat{p}'$ for the corresponding new belief states, respectively. To show that $p'=\widehat{p}'$, we consider a fixed $a\in\Omega$ and show that $p'(a)=\widehat{p}'(a)$. Note that

$$
\begin{aligned}
p'(a) &= p(a|B_a\cap C_a)\pi_{B_a}^{C_a}p(C_a), &(17)\\
\widehat{p}'(a) &= p(a|\widehat{B}_a\cap\widehat{C}_a)\widehat{\pi}_{\widehat{B}_a}^{\widehat{C}_a}p(\widehat{C}_a), &(18)
\end{aligned}
$$

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

$$\cup_{C\in\mathcal{C}_{\text{triv}}}C=\cup_{\widehat{C}\in\widehat{\mathcal{C}}_{\text{triv}}}\widehat{C}. \tag{19}$$

We consider two cases.

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

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

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

Hence, $p'(a)=\widehat{p}'(a)$. ■

A.2 Proposition 1

Proof of Proposition 1. Suppose that $\#\Omega\geq 3$. Suppose, for a contradiction, that there exists a responsive and conservative revision rule on a domain $\mathcal{D}\supseteq\mathcal{D}_{\text{Jeffrey}}$. Since $\#\Omega\geq 3$, we can find events $A,B\subseteq\Omega$ such that $A\cap B,B\backslash A,A\backslash B\neq\varnothing$. Consider an initial belief state p such that $p(A\cap B)=1/4$ and $p(A\backslash B)=3/4$, and define the Jeffrey input $I:=\{p':p'(B)=1/2\}$. Note that $(p,I)\in\mathcal{D}$. What is the new belief state p_I?

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

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

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

A.3 Proposition 2

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

Lemma 7 *For all inputs I and all events $\varnothing \subsetneq A \subsetneq B \subseteq Supp(I)$, I is strongly silent on the probability of A given B if and only if*

- *I contains a belief state p^* with $p^*(A), p^*(B \backslash A) \neq 0$, and*
- *for every such $p^* \in I$ and every $\alpha \in [0,1]$, I contains the belief state p' which coincides with α on the probability of A given B and with p^* outside that conditional probability, formally*

$$p' \in I \text{ where } p' = p^*(\cdot|A)\alpha p^*(B) + p^*(\cdot|B\backslash A)(1-\alpha)p^*(B) + p^*(\cdot \cap \overline{B}).$$

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

(a) First suppose $I_{A|B} = \mathcal{P}$. Consider any $\alpha \in [0,1]$. As $\varnothing \subsetneq A \subsetneq B$, there exists a belief state p' such that $p'(B) \neq 0$ and $p'(A|B) = \alpha$. As $I_{A|B} = \mathcal{P}$, we have $p' \in I_{A|B}$, so that I 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 for weak silence.

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

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

First suppose $I_{\overline{A|B}} = I$. To show strong silence, consider any $\alpha \in [0,1]$ and any $p^* \in I$ with $p^*(A), p^*(B\backslash A) \neq 0$. 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 I. Clearly, p' belongs to $I_{\overline{A|B}}$. Hence, as $I = I_{\overline{A|B}}$, p' belongs to I.

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

We distinguish two cases. First suppose $p^*(A), p^*(B\backslash A) \neq 0$. Then $p'(B) = p^*(B) \neq 0$. By I's strong silence on the probability of A given B, I 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 $p^*(A), p^*(B\backslash A) \neq 0$, 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 $p' \in I$, as had to be shown.

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

A.4 On which conditional probabilities is each kind of learnt input strongly silent?

As a key step towards proving our theorems, this section determines on which conditional probabilities learnt inputs of the four kinds (Bayesian, Jeffrey, dual-Jeffrey, and Adams) are strongly silent. Instead of discussing Bayesian inputs explicitly, we turn directly to Jeffrey inputs, since they generalize Bayesian inputs.

Lemma 8 *For all Jeffrey inputs I (of learning a new probability distribution on a partition $\mathcal{B}$) and all events $\varnothing \subsetneq A \subsetneq B \subseteq Supp(I)$, I 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 I, $\mathcal{B}$, A, and B be as specified, and let $(\pi_B)_{B\in\mathcal{B}}$ be the learnt probability distribution on $\mathcal{B}$. First, if $B \subseteq B'$ for some $B' \in \mathcal{B}$, then I is strongly silent on the probability of A given B, as one can easily check, using Lemma 7. Conversely, suppose that $B \not\subseteq B'$ for all $B' \in \mathcal{B}$. For each $D \subseteq \Omega$, we write $\mathcal{B}_D := \{B' \in \mathcal{B} : B' \cap D \neq \varnothing\}$. Note that $\mathcal{B}_B = \mathcal{B}_A \cup \mathcal{B}_{B\backslash A}$, where $\#\mathcal{B}_A \geq 1$ (as $A \neq \varnothing$), $\#\mathcal{B}_{B\backslash A} \geq 1$ (as $B\backslash A \neq \varnothing$), 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 $B'' \in \mathcal{B}_{B\backslash A}$ with $B' \neq B''$. Note that I contains a p^* such that $p^*(B' \cap A) = \pi_{B'}$ and $p^*(B'' \cap (B\backslash A)) = \pi_{B''}$. Since each of B' and B'' has a non-empty intersection with B, and hence with $Supp(I)$ $(\supseteq B)$, we have $\pi_{B'}, \pi_{B''} \neq 0$. Now $p^*(B'' \cap A) = p^*(B'' \cap \overline{B}) = 0$, since

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

By Lemma 7, if I were strongly silent on the probability of A given B, I 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., I 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 fact that $p^*(B'' \cap A) = p^*(B'' \cap \overline{B}) = 0$, which we have shown. Hence, I is not strongly silent on the probability of A given B. ■

Next, we determine on which conditional probabilities dual-Jeffrey inputs are strongly silent.

Lemma 9 *For all dual-Jeffrey inputs I (of learning a new conditional probability distribution given a partition $\mathcal{C}$) and all events $\varnothing \subsetneq A \subsetneq B \subseteq \Omega$ (= Supp(I)), I 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 $\varnothing \subsetneq \mathcal{C}_A \subsetneq \mathcal{C}_B \subseteq \mathcal{C}$.*

Proof. Let $I, \mathcal{C}$, A, and B be as specified, and let $(\pi^C)^{C \in \mathcal{C}}$ be the learnt 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 $\varnothing \subsetneq \mathcal{C}_A \subsetneq \mathcal{C}_B \subseteq \mathcal{C}$, then I 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, $p^* \in I$. If I were strongly silent on the probability of A given B, then I 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 I fails to contain p', for the following reason. We distinguish 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 $C \cap A, C \backslash A \neq \varnothing$. By the definition of p' (and the fact that $C \cap A, C \backslash A \neq \varnothing$), $p'(C \cap A) > p^*(C \cap A)$ and $0 < p'(C \backslash A) \leq p^*(C \backslash A)$. This implies that $p'(C), p^*(C) \neq 0$ and $p'(A|C) > p^*(A|C)$. So, as $p^*(\cdot|C) = \pi^C$ (by $p^* \in I$), $p'(\cdot|C) \neq \pi^C$, and therefore $p' \notin I$.

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 exists $C \in \mathcal{C}$ such that $C \cap B, C \backslash B \neq \varnothing$. As A is included in B and a union of sets in $\mathcal{C}$, $C \cap A = \varnothing$. Note that $p^*(C \cap B), p^*(C \backslash B) \neq 0$ (as $C \cap B, C \backslash B \neq \varnothing$ and by definition of p^*); further, that $p'(C \cap B) = p'(C \cap (B \backslash A)) = 0$ (where the first equality holds because $C \cap A = \varnothing$ and the second by definition of p'); and finally, that $p'(C) = p'(C \cap B) + p'(C \backslash B) = 0 + p^*(C \cap B) \neq 0$. 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 $p^* \in I$), $p'(\cdot|C) \neq \pi^C$, and so $p' \notin I$. ■

We now turn to Adams inputs. Before we show on which conditional probabilities they are strongly silent, we prove two useful lemmas.

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

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

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

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 $p'' \in I$.
- 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 $p' \in I$.
- 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 exists $\beta > 0$ 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 I is an Adams input, $(\pi_B^C)_{B\in\mathcal{B}}^{C\in\mathcal{C}}$ is its canonical family, $p_B \in \mathcal{P}$ with $Supp(p_B) \subseteq B$ 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 I contains*

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

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

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

The next lemma determines on which conditional probabilities Adams inputs are strongly silent, combining insights from Lemmas 8 and 9 about Jeffrey and dual-Jeffrey inputs. In fact, the next lemma implies Lemma 9 if Ω is finite – not surprisingly, since, for finite Ω, Adams inputs generalize dual-Jeffrey inputs.

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

Lemma 12 *Consider an Adams input I and let $(\pi_B^C)_{B\in\mathcal{B}}^{C\in\mathcal{C}}$ be the canonical family generating it (as defined in Lemma 5). For all events $\varnothing \subsetneq A \subsetneq B \subseteq \Omega$ $(= Supp(I))$, I 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 $\mathcal{C}_A \subseteq \mathcal{C}_B \subseteq \mathcal{C} \backslash (\mathcal{B}\cap\mathcal{C})$ and some $D_A \subseteq D_B \subseteq \cup_{D\in\mathcal{B}\cap\mathcal{C}} D$.[25]

Proof. Let I, $(\pi_B^C)_{B\in\mathcal{B}}^{C\in\mathcal{C}}$, 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 $\Omega^* := \Omega \backslash (\cup_{D\in\mathcal{D}} D)$.

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.

[25] Since $(\pi_B^C)_{B\in\mathcal{B}}^{C\in\mathcal{C}}$ is canonical, the set $\mathcal{B}\cap\mathcal{C}$ is either empty or a singleton set $\{D^*\}$. So, the union $\cup_{D\in\mathcal{B}\cap\mathcal{C}} D$ is either $\varnothing$ 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 $D_A = D_B = \varnothing$, and in the second case it reduces to $D_A \subseteq D_B \subseteq D^*$.

Now suppose I is strongly silent on the probability of A given B. Suppose, for a contradiction, that A and B are neither of the form (a) nor of the form (b). We derive a contradiction in each of the following cases. The proof will sometimes use a uniform distribution on a non-empty event $B' \subseteq \Omega$, denoted $uniform_{B'}$ and defined by $uniform_{B'}(A) = \frac{|A \cap B'|}{|B'|}$ for all $A \subseteq \Omega$. To be precise, $uniform_{B'}$ is only defined if $|B'| < \infty$, and thus the proof given here is literally valid only if Ω is finite. To extend the proof to general (countable) Ω, it suffices to replace each $uniform_{B'}$ by a distribution in $\mathcal{P}$ with support B'.

Case 1: There does not exist any $C \in \mathcal{C}\backslash\mathcal{D}$ such that $C \cap A, C\backslash A \neq \varnothing$. In other words, $A = (\cup_{C \in \mathcal{C}_A} C) \cup D_A$ for some $\mathcal{C}_A \subseteq \mathcal{C}\backslash\mathcal{D}$ 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 $\mathcal{C}_B \subseteq \mathcal{C}\backslash\mathcal{D}$ and $D_B \subseteq \cup_{D \in \mathcal{D}} D$. In other words, there exists $C \in \mathcal{C}\backslash\mathcal{D}$ such that $C \cap B, C\backslash B \neq \varnothing$. Since $B \cap C, C\backslash B \neq \varnothing$ and since the set $\mathcal{B}_C$ (which partitions C) has at least two members, there are distinct $\widehat{B}, \widetilde{B} \in \mathcal{B}_C$ such that $\widehat{B} \cap B, \widetilde{B}\backslash B \neq \varnothing$.

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

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

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

Hence, since I is strongly silent on the probability of A given B, and since $p^*(A) \neq 0$ (as $A \cap C^* \neq \varnothing$) and $p^*(B\backslash A) \neq 0$ (as $a_{\widehat{B}} \in B\backslash A$), Lemma 7 implies that I 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

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

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

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

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

$$\begin{aligned} p^* \;:&= \frac{1}{2}\left(\pi^{C^*}_{B^*} uniform_{B^* \cap (B\backslash A)} + \sum_{B' \in \mathcal{B}_{C^*}\backslash\{B^*\}} \pi^{C^*}_{B'} uniform_{B'}\right) \\ &+ \frac{1}{2} \sum_{B' \in \mathcal{B}_C} \pi^C_{B'} uniform_{B'}. \end{aligned}$$

So, because I is strongly silent on the probability of A given B (and because $p^*(A), p^*(B \backslash A) \neq 0$), by Lemma 7 I also contains the belief state p' that 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 $\widetilde{B} \cap A \neq \varnothing$, 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

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

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

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

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

$$p'(\widehat{B}|C) = \frac{p'(\widehat{B})}{\sum_{\widetilde{B} \in \mathcal{B}_C} p'(\widetilde{B})} < \frac{p^*(\widehat{B})}{\sum_{\widetilde{B} \in \mathcal{B}_C} p^*(\widetilde{B})} = p^*(\widehat{B}|C) = \pi^C_{\widehat{B}}.$$

So, $p'(\widehat{B}|C) \neq \pi^C_{\widehat{B}}$, a contradiction since $p' \in I$.

Subcase 2.2: $(B \backslash A) \cap C \neq \varnothing$ and no set in $\mathcal{B}_C$ includes $B \cap C$. Since $(B \backslash A) \cap C$ 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

$$\begin{aligned} \varnothing &\neq B_1 \cap ((B \backslash A) \cap C)\ (= B_1 \cap (B \backslash A)), \\ \varnothing &\neq B_2 \cap (A \cap C)\ (= B_2 \cap A). \end{aligned}$$

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

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

So, since I is strongly silent on the probability of A given B (and since $p^*(A) \neq 0$ as $a_{B_2} \in A$ and since $p^*(B \backslash A) \neq 0$ as $a_{B_1} \in B \backslash A$), by Lemma 7 I 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

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

We have

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

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

Subcase 2.3: $(B \backslash A) \cap C \neq \varnothing$ 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., $B \backslash C \neq \varnothing$. Hence there are $\widehat{C} \in \mathcal{C} \backslash \{C\}$ and $\widehat{B} \in \mathcal{B}_{\widehat{C}}$ such that $\widehat{B} \cap B \neq \varnothing$; hence, $\widehat{C} \cap B \neq \varnothing$. (Possibly $\widehat{B} = \widehat{C}$.)

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

$$\begin{aligned} p^* \quad : &= \frac{1}{2}\left(\pi^C_{B^*} uniform_{B^* \cap A} + \sum_{B' \in \mathcal{B}_C \backslash \{B^*\}} \pi^C_{B'} uniform_{B'}\right) \\ &+ \frac{1}{2} \sum_{B' \in \mathcal{B}_{\widehat{C}}} \pi^{\widehat{C}}_{B'} uniform_{B'} \end{aligned}$$

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

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

Notice that

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

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

Subsubcase 2.3.2: $A \cap \widehat{C} = \varnothing$. We re-define p^* by replacing '$unform_{B^* \cap A}$' with '$uniform_{B^* \cap (B \backslash A)}$' in the previous definition of p^*. Again, $p^* \in I$, by Lemma 11. So, since I is strongly silent on the probability of A given B (and since $p^*(A), p^*(B \backslash A) \neq 0$), I 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

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

where again the latter is positive. So, $p'(B^*|C) = 0 \neq \pi^C_{B^*}$, a contradiction since $p' \in I$. ■

A.5 Theorems 1 and 2

Using our previous lemmas, we finally prove our central characterization of the four revision rules (Theorems 1 and 2).

Proof of Theorems 1 and 2. It suffices to consider Jeffrey's, the dual-Jeffrey, and Adams's rules, since Bayes's rule is extended by Jeffrey's. We first prove one direction of implication of both theorems, and then we prove the other direction.

Part 1: First, we consider a responsive and conservative revision rule on one of the domains $\mathcal{D}_{\text{Jeffrey}}$, $\mathcal{D}_{\text{dual-Jeffrey}}$, and $\mathcal{D}_{\text{Adams}}$. We show that the rule is Jeffrey's, the dual-Jeffrey, or Adams's rule, respectively. We distinguish between the three domains.

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

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

where $p_I(B)$ reduces to π_B by Responsiveness, and $p_I(\cdot|B)$ reduces to $p(\cdot|B)$ by Conservativeness. (Note that, by Lemma 8, I is strongly silent on the probability given B of any event strictly between $\varnothing$ and B.)

Dual-Jeffrey: Suppose $(p, I) \in \mathcal{D}_{\text{dual-Jeffrey}}$, say $I = \{p' : p'(\cdot|C) = \pi^C \ \forall C \in \mathcal{C}$ such that $p'(C) \neq 0\}$. Then p_I is given by the dual-Jeffrey rule, because we may expand p_I as

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

where $p_I(\cdot|C)$ reduces to π^C by Responsiveness, and $p_I(C)$ reduces to $p(C)$ by Conservativeness. (Note that, by Lemma 9, I is strongly silent on the probability of C given Ω if $C \neq \Omega$.)

Adams: Suppose $(p, I) \in \mathcal{D}_{\text{Adams}}$, say $I = \{p' : p'(B|C) = \pi_B^C \ \forall B \in \mathcal{B} \ \forall C \in \mathcal{C}$ such that $p'(C) \neq 0\}$. Then p_I is given by Adams's rule, because we may expand p_I as

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

where $p_I(B|C)$ reduces to π_B^C by Responsiveness, $p_I(C)$ reduces to $p(C)$ by Conservativeness (since, by Lemma 12, I is strongly silent on the probability of C given Ω if $C \neq \Omega$), and $p_I(\cdot|B \cap C)$ reduces to $p(\cdot|B \cap C)$ by Conservativeness (since, by Lemma 12, I is strongly silent on the probability given $B \cap C$ of any event strictly between $\varnothing$ and $B \cap C$). □

Part 2: Conversely, we now show that Jeffrey's, the dual-Jeffrey, and Adams's rules are responsive and conservative. Responsiveness is obvious. To establish Conservativeness, consider any (p, I) in the rule's domain ($\mathcal{D}_{\text{Jeffrey}}$ or $\mathcal{D}_{\text{dual-Jeffrey}}$ or $\mathcal{D}_{\text{Adams}}$) and any events $\varnothing \subsetneq A \subsetneq B \subseteq Supp(I)$ such that I is strongly silent on the probability of A given B and $p_I(B), p(B) \neq 0$. We have to show that $p_I(A|B) = p(A|B)$. We again distinguish between the three rules.

Jeffrey: Suppose the rule in question is Jeffrey's rule. Then the input takes the form $I = \{p' : p'(B) = \pi_B \ \forall B \in \mathcal{B}\}$ for some learnt probability distribution $(\pi_B)_{B \in \mathcal{B}}$ on some partition $\mathcal{B}$. As I 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 $p_I(A|B) = p(A|B)$, because

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

where the second equality holds by the definition of Jeffrey revision.

Dual-Jeffrey: Consider the dual-Jeffrey rule. Then I is a dual-Jeffrey input, of the form $I = \{p' : p'(\cdot|C) = \pi^C \ \forall C \in \mathcal{C}$ such that $p'(C) \neq 0\}$ for some (unique) conditional probability distribution $(\pi^C)^{C\in\mathcal{C}}$ given some partition $\mathcal{C}$. By I's strong silence on the probability of A given B and Lemma 9, $A = \cup_{C\in\mathcal{C}_A} C$ and $B = \cup_{C\in\mathcal{C}_B} C$ for some sets $\varnothing \subsetneq \mathcal{C}_A \subsetneq \mathcal{C}_B \subseteq \mathcal{C}$. We have $p_I(A|B) = p(A|B)$, because

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

where, as one can easily verify, each $p_I(C)$ equals $p(C)$.

Adams: Consider Adams's rule. Then I is an Adams input, of the form $I = \{p' : p'(B|C) = \pi_B^C \ \forall B \in \mathcal{B} \ \forall C \in \mathcal{C}$ 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 I, i.e., that $\mathcal{B}$ refines $\mathcal{C}$ and $\mathcal{B} \cap \mathcal{C}$ is empty or singleton. By I'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 $\mathcal{C}_A \subseteq \mathcal{C}_B \subseteq \mathcal{C} \backslash (\mathcal{B} \cap \mathcal{C})$ 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 $D_A = D_B = \varnothing$ or $D_A \subseteq D_B \subseteq D$, respectively.)

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

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

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

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

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