

Belief revision generalized: A joint characterization of Bayes's and Jeffrey's rules

Franz Dietrich, Christian List, and Richard Bradley¹

October 2010, final version December 2015

Abstract

We present a general framework for representing belief-revision rules and use it to characterize Bayes's rule as a classical example and Jeffrey's rule as a non-classical one. In Jeffrey's rule, the input to a belief revision is not simply the information that some event has occurred, as in Bayes's rule, but a new assignment of probabilities to some events. Despite their differences, Bayes's and Jeffrey's rules can be characterized in terms of the same 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. To illustrate the use of non-Bayesian belief revision in economic theory, we sketch a simple decision-theoretic application.

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

JEL classification: C73, D01, D80, D81, D83, D90

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. 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

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

¹F. Dietrich, Paris School of Economics & CNRS; C. List, London School of Economics; R. Bradley, London School of Economics. We are grateful for comments from the editors and referees and from audiences at 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 is jointly authored, List and Bradley wish to note that the bulk of the mathematical credit should go to Dietrich. Dietrich was supported by a Ludwig Lachmann Fellowship at the LSE and the French Agence Nationale de la Recherche (ANR-12-INEG-0006-01). List was supported by a Leverhulme Major Research Fellowship (MRF-2012-100) and the Franco-Swedish Program in Philosophy and Economics (via a visit to the University of Uppsala). Bradley was supported by the Arts and Humanities Research Council (via a grant on 'Managing Severe Uncertainty', AH/J006033/1).

²We expect that our results can be generalized to a set Ω of arbitrary cardinality.

In economic theory, belief changes are almost always modelled in this way. The aim of this paper is to draw attention to a more general form of belief revision, which is seldom discussed in economics. We develop a general framework in which different belief-revision rules – Bayesian and non-Bayesian – can be characterized. In this framework, the key difference between different belief-revision rules lies in what they take to be 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, and we show that there is scope for useful generalization.

We begin with an example that will resonate with any international traveller. Kotaro, a junior academic from Japan, is applying for a faculty position in the UK. After his interview, he is telephoned by Terence, the chair of the department, to inform him of the outcome. Kotaro gets a vague impression that he is being offered the job, but struggles to understand Terence’s thick Irish accent. At the end of the call, he is still unsure whether he has received an offer. He becomes convinced only after a subsequent conversation with another member of the department.

This example illustrates an instance of belief revision triggered by a noisy signal. Before the telephone conversation with the chair of the department, Kotaro attaches a very low probability to the event of getting the job. After the conversation, he attaches a somewhat higher probability to it, but one that still falls short of certainty. For this, a second conversation (with another person) is needed. Such cases present challenges to the Bayesian modeller. If the first change in Kotaro’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 Kotaro’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 Kotaro’s eardrums. The signal he receives from Terence 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 Terence’s words are subject to a triple distortion: by his thick accent, by imperfections of the telephone line, and by the interpretation of a non-native speaker.

Enriching the set Ω in this way, however, has definite modelling costs. First, the agent (Kotaro) almost certainly does 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 is initially focused. Second, it is doubtful that before the conversation he is even aware of the possibility of such complex auditory signals (probably he has never heard, or even imagined, an accent like Terence’s). So, a Bayesian model of this 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.

One of the most prominent generalizations 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 a job offer, as perhaps in the example of Kotaro, the junior academic. 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. 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 \bar{B} , with $\pi_B = 1$ and $\pi_{\bar{B}} = 0$. (The *complement* of any event B is $\bar{B} = \Omega \setminus B$.)

The framework we develop allows us to define and characterize different belief-revision rules. 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. But, in principle, the learnt inputs could also be very different, such as new conditional probabilities of certain events. We show that, despite their differences, Bayes’s and Jeffrey’s rules can be 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

³For 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 $\pi_B = 0$).

not change. The fact that one of most prominent non-classical belief-revision rules, namely Jeffrey’s rule, 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. Our characterization can be extended to other non-classical belief-revision rules too, but for expositional simplicity, we set these aside in this paper.

We hope that our framework will 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 Jeffrey 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. Bayes’s and Jeffrey’s rules have been axiomatically characterized in previous work, but the existing approaches are less unified than ours. One approach is a ‘distance-based’ one. This consists in showing that a given belief-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. These results, however, do not generally carry over to other belief-revision rules without changing the distance metric,⁶ and the interpretation of different distance metrics is often controversial.

Another approach to characterizing belief-revision rules invokes the idea of ‘rigidity’ rather than distance-minimization (e.g., Jeffrey 1957 and Bradley 2005 for some extensions). For example, Bayesian belief revision is ‘rigid’ in the sense of preserving the conditional probability of any event, given the learnt event. Although this approach is closer in spirit to ours, it also lacks unification (since the notion of rigidity is not fully general). Still, one might interpret our present conservativeness axiom as a unified alternative to 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 textbook (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.

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

⁵Distance is represented by a function $d : \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}$, 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, $p, p' \in \mathcal{P}$.

⁶For example, Douven and Romeijn (2011) characterize another belief-revision rule, Adams’s rule (which we briefly discuss in Section 5), by invoking a different measure of distance, the *inverse relative-entropy distance*, which differs from ordinary relative-entropy distance in the inverted order of its arguments.

2 A general framework

We can study belief revision in general by specifying

- (i) a set \mathcal{P} of possible *belief states* in which a given agent can be, and
- (ii) a set \mathcal{I} of possible *inputs* which can influence the agent's belief state.

A *revision rule* maps pairs (p, I) of an initial belief state p in \mathcal{P} and an input I in \mathcal{I} to a new belief 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 belief–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 beliefs that take the form of subjective probability measures, the set \mathcal{P} of possible *belief states* is the set of all probability measures over the events from Ω . Formally, a *probability measure* is a countably additive function $p : 2^\Omega \rightarrow [0, 1]$ with $p(\Omega) = 1$ (where Ω is the underlying finite or countably infinite set of worlds). 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 $B \neq \emptyset$; we then write $I = I_B$;⁷
- **Jeffrey** if $I = \{p' : p'(B) = \pi_B \text{ for all } B \in \mathcal{B}\}$ for some probability distribution $(\pi_B)_{B \in \mathcal{B}}$ on some partition \mathcal{B} ; we then write $I = I_{(\pi_B)_{B \in \mathcal{B}}}$.⁸

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. Our framework also allows us to represent many other kinds of inputs, for instance $I = \{p' : p'(A \cap B) > p'(A)p'(B)\}$, which captures the constraint that the events

⁷The representation $I = I_B$ is unique, because for any Bayesian input I , there exists a unique event B such that $I = I_B$.

⁸For 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 $|\{B \in \mathcal{B} : \pi_B = 0\}| \leq 1$.

A and B are positively correlated, or $I = \{p' : p'(A) \geq 9/10\}$, which captures the constraint that A is highly 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 or her old belief state.

We are now able to define Bayes's and Jeffrey's rules in this framework. (Of course, the framework equally permits the definition of other belief-revision rules.)

Definition 2

- Let $\mathcal{D}_{\text{Bayes}}$ 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 $\mathcal{D}_{\text{Bayes}}$ which maps each $(p, I_B) \in \mathcal{D}_{\text{Bayes}}$ to $p' \in \mathcal{P}$, where

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

- Let $\mathcal{D}_{\text{Jeffrey}}$ 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 $\mathcal{D}_{\text{Jeffrey}}$ which maps each $(p, I_{(\pi_B)_{B \in \mathcal{B}}}) \in \mathcal{D}_{\text{Jeffrey}}$ to $p' \in \mathcal{P}$, where

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

The domains $\mathcal{D}_{\text{Bayes}}$ and $\mathcal{D}_{\text{Jeffrey}}$ are the maximal domains for which formulas (1) and (2) are well-defined.⁹ Jeffrey's rule extends Bayes's, i.e., it coincides with Bayes's rule on the subdomain $\mathcal{D}_{\text{Bayes}} (\subseteq \mathcal{D}_{\text{Jeffrey}})$.

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 or her 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. As already noted, the same axioms can also be used to characterize other belief-revision rules, but we relegate these extensions to follow-up work (as briefly discussed in Section 5). 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.

⁹The 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: $p_I \in I$ for all belief–input pairs $(p, I) \in \mathcal{D}$.

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 is 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

$$\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(I),$$

where $\text{Supp}(I)$ is the *support* of I , defined as $\{\omega \in \Omega : p'(\omega) \neq 0 \text{ for some } p' \in I\}$.¹⁰

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:

¹⁰Here, and elsewhere, we write $p'(\omega)$ as an abbreviation for $p'(\{\omega\})$ when we refer to the probability of a singleton event $\{\omega\}$.

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

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 $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(I)$), 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 $\mathcal{D} \supseteq \mathcal{D}_{\text{Jeffrey}}$ is responsive and strongly conservative.*

Note that, on the small domain $\mathcal{D}_{\text{Bayes}}$, 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 $p'(\cdot|A)$, $p'(\cdot|B\setminus A)$, and $p'(\cdot|\overline{B})$.¹¹ This quadruple and the conditional probability $p'(A|B)$ jointly determine the belief state p' , because

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

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 $C \in \{A, B\setminus A, \overline{B}\}$ 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 $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(I)$) 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*

¹¹This informal discussion assumes that $p'(A), p'(B\setminus A), p'(\overline{B}) \neq 0$.

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

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

$$p' := p^*(\cdot|A)\alpha p^*(B) + p^*(\cdot|B \setminus A)(1 - \alpha)p^*(B) + p^*(\cdot \cap \bar{B}), \quad (3)$$

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).¹²

For example, the inputs $I = \{p' : p' \text{ is uniform on } \bar{B}\}$ 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 $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(I)$), 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

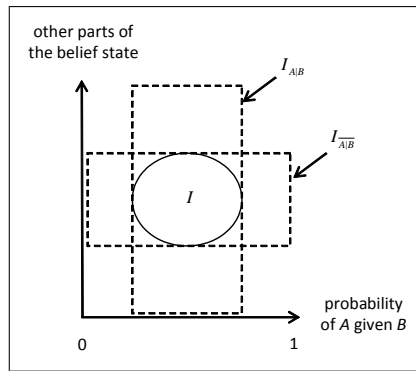


Figure 2: The implications $I_{A|B}$ and $I_{\bar{A}|\bar{B}}$ derived from input I

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

¹²To be precise, this is true whenever $p^*(A), p^*(B \setminus A) \neq 0$.

- 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' 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).¹³
- 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).¹⁴

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 $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(I)$),*

- 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),*
- 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.

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

¹³In 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 $p^*(B) \neq 0$.

¹⁴In 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 $[p^*(C) \neq 0 \text{ for all } C \in \{A, B \setminus A\} \text{ such that } p'(C) \neq 0]$, then p' coincides with some such p^* outside the probability of A given B .

It is also worth noting the following consequence of this result:

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 $\mathcal{D} \cap \mathcal{D}_{\text{Bayes}}$ and with Jeffrey’s rule on $\mathcal{D} \cap \mathcal{D}_{\text{Jeffrey}}$.*

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 one must take care 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, however, we can verify that the corresponding belief-revision rule does indeed preserve all of them, as required by conservativeness.

4 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 of Jeffrey revision 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 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 .




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

- 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 $(v(c) - 1)p(c|e) + (v(\bar{c}) - 1)p(\bar{c}|e) = 4\frac{4}{5} + (-6)\frac{1}{5} = 2$, 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 $p(c|\bar{e}) = \frac{p(c,\bar{e})}{p(\bar{e})} = \frac{0.1}{0.5} = \frac{1}{5}$. So, she rejects him as this yields an expected utility of -1 , whereas hiring him would have yielded an expected utility of $(v(c) - 1)p(c|\bar{e}) + (v(\bar{c}) - 1)p(\bar{c}|\bar{e}) = 4\frac{1}{5} + (-6)\frac{4}{5} = -4$.

So *ex ante* the expected utility of gathering information is the average $2p(e) + (-1)p(\bar{e}) = 2\frac{1}{2} + (-1)\frac{1}{2} = \frac{1}{2}$, 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

(a) Refined classical decision problem



(b) Decision problem with an unforeseen input




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

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 $p_I(c, e) = \frac{1}{10}$, $p_I(c, \bar{e}) = \frac{1}{40}$, $p_I(\bar{c}, e) = \frac{7}{40}$, and $p_I(\bar{c}, \bar{e}) = \frac{7}{10}$. 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 $p_I(c|e) = \frac{p_I(c, e)}{p_I(c, e) + p_I(\bar{c}, e)} = \frac{4}{11}$, 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 (\bar{w}), where the probability of w , denoted $t_{\alpha, \beta}$, is larger if Bob is competent than if he is not, i.e., $t_{c, \beta} > t_{\bar{c}, \beta}$. 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 , \bar{h} , and g), Ann is falsely taken to foresee the possibilities of learning w or learning \bar{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 .
- (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 $\Omega' = \{c, \bar{c}\} \times \{e, \bar{e}\} \times \{w, \bar{w}\}$.

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 \bar{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.¹⁵ 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.

5 Concluding remarks

We have developed a unified framework for the study of belief revision and shown that Bayes’s and Jeffrey’s rules can be characterized in terms of the same two axioms: responsiveness to the learnt input and conservativeness. The only difference between those rules lies in the domains of learnt inputs to which they apply. We show in follow-up work that our analysis can be extended to other belief-revision rules, distinct from Bayes’s and Jeffrey’s rules. One such rule is Adams’s rule, inspired by Ernst Adams’s work on the logic of conditionals (Adams 1975) and formally introduced by Bradley (2005). It takes the learnt input to be a new assignment of conditional probabilities of some events, given some other events. Another rule is the ‘dual-Jeffrey’ rule, which stands out for its natural ‘duality’ to Jeffrey’s rule. Here, the learnt input is a new conditional probability distribution, given a certain partition. Like Bayes’s and Jeffrey’s rules, these two rules also uniquely satisfy our two axioms on their respective domains.

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 inputs (as well as other non-Bayesian inputs) in a Bayesian manner. However, as we have noted in our discussions of ‘Kotaro the junior academic’ and ‘Ann the employer’, such Bayesian remodelling comes at a cost:

Over-ascription of opinionation: A significant drawback of the Bayesian remodelling is that we must assume that the agent is able to assign prior probabilities to many complex events. In our initial example, Kotaro must assign prior 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

¹⁵ 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 ($\beta = e$).

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, once 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 Terence (the departmental chair with the thick accent), Kotaro even considered the possibility of receiving such 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) might 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.

6 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.
- 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): 159–173.

- Diaconis, P., Zabell, S. L. (1982) Updating subjective probability, *Journal of the American Statistical Association* 77(380): 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., Schipper, B. C. (2006). Interactive unawareness, *Journal of Economic Theory*, 130(1): 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(2): 265–298.
- Shafer, G. (1981) Jeffrey’s rule of conditioning, *Philosophy of Science* 48(3): 337–362.
- van Fraassen, B. C. (1981) A Problem for Relative Information Minimizers in Probability Kinematics, *British Journal for the Philosophy of Science* 32(4): 375–379.
- Vineberg, S. (2011) Dutch Book Arguments, *Stanford Encyclopedia of Philosophy* (Summer 2011 Edition).

A Appendix: Proofs

A.1 Well-definedness of each revision rule

Our two belief-revision rules are 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 family of probabilities (π_B) – is *uniquely* determined by the relevant input I (in the case of Bayes’s rule) or *sufficiently* determined (in the case of Jeffrey’s rule) so that the definition does not depend on non-unique features. The following two lemmas, which the reader can easily verify, make this precise:

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

Lemma 2 *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.¹⁶*

¹⁶Footnote 1 specifies when (2) is defined.

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 $A \cap B, B \setminus A, A \setminus B \neq \emptyset$. Consider an initial belief state p such that $p(A \cap B) = 1/4$ and $p(A \setminus B) = 3/4$, and define the Jeffrey 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, we have $p_I(A \cap B|B) = p(A \cap B|B)$ (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, we have $p_I(B) = 1/2$, a contradiction. ■

A.3 Proposition 2

We begin by offering a convenient reformulation of strong silence; we leave the proof to the reader.

Lemma 3 *For all inputs I and all events $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(I)$, I is strongly silent on the probability of A given B if and only if*

- *I contains a belief state p^* with $p^*(A), p^*(B \setminus A) \neq 0$, 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*

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

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

(a) First suppose $I_{A|B} = \mathcal{P}$. Consider any $\alpha \in [0, 1]$. As $\emptyset \subsetneq A \subsetneq B$, there exists a belief state p' such that $p'(B) \neq 0$ and $p'(A|B) = \alpha$. As $I_{A|B} = \mathcal{P}$, 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 $p'(A), p'(B \setminus A) \neq 0$, the equivalence holds because strong silence is violated (see Lemma 3) and moreover $I_{A|\overline{B}} \neq I$ because $I_{A|\overline{B}}$ but not I contains a belief state p' such that $p'(A), p'(B \setminus A) \neq 0$. Now consider the less trivial case in which I contains a \tilde{p} such that $\tilde{p}(A), \tilde{p}(B \setminus A) \neq 0$.

First suppose $I_{A|\overline{B}} = I$. To show strong silence, consider any $\alpha \in [0, 1]$ and any $p^* \in I$ with $p^*(A), p^*(B \setminus A) \neq 0$. By Lemma 3, 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 $C \in \{A, B \setminus A\}$ with $p'(C) \neq 0$.

We distinguish two cases. First suppose $p^*(A), p^*(B \setminus A) \neq 0$. Then $p'(B) = p^*(B) \neq 0$. By 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 $p^*(A), p^*(B \setminus A) \neq 0$, there can be only one belief state that coincides with p^* outside the probability of A given B and such that the probability of A given B takes a given value. Therefore, $p' = \tilde{p}$, and so $p' \in I$, as required.

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

A.4 The main theorem

As a key step towards proving our main theorem, we now answer the following question: on which conditional probabilities are the learnt inputs under Bayes's and Jeffrey's rules strongly silent?

Lemma 4 *Any Bayesian input I (of learning that some event B' has occurred) is strongly silent on any conditional probability of A given B such that $\emptyset \subsetneq A \subsetneq B \subseteq B'$.*

Lemma 5 *For all Jeffrey inputs I (of learning a new probability distribution on a partition \mathcal{B}) and all events $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(I)$, I is strongly silent on the conditional probability of A given B if and only if $B \subseteq B'$ for some $B' \in \mathcal{B}$.*

Since Lemma 4 is a special case of Lemma 5, we turn directly to the proof of the latter.

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 3. Conversely, suppose that $B \not\subseteq B'$ for all $B' \in \mathcal{B}$. For each $D \subseteq \Omega$, we write $\mathcal{B}_D := \{B' \in \mathcal{B} : B' \cap D \neq \emptyset\}$. Note that $\mathcal{B}_B = \mathcal{B}_A \cup \mathcal{B}_{B \setminus A}$, where $\#\mathcal{B}_A \geq 1$ (as $A \neq \emptyset$), $\#\mathcal{B}_{B \setminus A} \geq 1$ (as $B \setminus A \neq \emptyset$), and $\#\mathcal{B}_B \geq 2$ (as otherwise B would be included in a $B' \subseteq \mathcal{B}$). It follows that there are $B' \in \mathcal{B}_A$ and $B'' \in \mathcal{B}_{B \setminus A}$ with $B' \neq B''$. Note that I contains a p^* such that $p^*(B' \cap A) = \pi_{B'}$ and $p^*(B'' \cap (B \setminus A)) = \pi_{B''}$. Since each of B' and B'' has a non-empty intersection with B , and hence with $\text{Supp}(I) (\supseteq B)$, we have $\pi_{B'}, \pi_{B''} \neq 0$. Now $p^*(B'' \cap A) = p^*(B'' \cap \overline{B}) = 0$, since

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

By Lemma 3, 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 . ■

We can now complete the proof of our characterization theorem.

Proof of Theorem 1. It suffices to consider Jeffrey's rule, since Bayes's rule is extended by Jeffrey's. We first prove one direction of implication, and then prove the other direction.

Part 1: We consider a responsive and conservative revision rule on the domain $\mathcal{D}_{\text{Jeffrey}}$ and show that the rule is indeed Jeffrey's rule. 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

$$p_I = \sum_{B \in \mathcal{B}: p_I(B) \neq 0} p_I(\cdot|B)p_I(B), \quad (4)$$

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 5, I is strongly silent on the probability, given B , of any event strictly between \emptyset and B .)

Part 2: Conversely, we now show that Jeffrey's rule is responsive and conservative. Responsiveness is obvious. To establish conservativeness, consider any (p, I) in the domain $\mathcal{D}_{\text{Jeffrey}}$ and any events $\emptyset \subsetneq A \subsetneq B \subseteq \text{Supp}(I)$ 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 $p_I(A|B) = p(A|B)$.

As a Jeffrey input, I takes the form $\{p' : p'(B) = \pi_B \ \forall B \in \mathcal{B}\}$ for some learnt probability distribution $(\pi_B)_{B \in \mathcal{B}}$ on some partition \mathcal{B} . Since I is strongly silent on the probability of A given B , we must have $B \subseteq B'$ for some $B' \in \mathcal{B}$, by Lemma 5. 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 identity holds by the definition of Jeffrey revision. ■