

On the Revision of Probabilistic Beliefs using Uncertain Evidence

Hei Chan and Adnan Darwiche
Computer Science Department
University of California, Los Angeles
Los Angeles, CA 90095
{hei,darwiche}@cs.ucla.edu

Abstract

We revisit the problem of revising probabilistic beliefs using uncertain evidence, and report results on four major issues relating to this problem: How to specify uncertain evidence? How to revise a distribution? Should, and do, iterated belief revisions commute? And how to provide guarantees on the amount of belief change induced by a revision? Our discussion is focused on two main methods for probabilistic revision: Jeffrey's rule of probability kinematics and Pearl's method of virtual evidence, where we analyze and unify these methods from the perspective of the questions posed above.

1 Introduction

We consider in this paper the problem of revising beliefs given uncertain evidence, where beliefs are represented using a probability distribution. There are two main methods for revising probabilistic beliefs in this case. The first method is known as *Jeffrey's rule* and is based on the principle of *probability kinematics*, which can be viewed as a principle for minimizing belief change [Jeffrey, 1965]. The second method is called *virtual evidence* and is proposed by Pearl in the context of belief networks—even though it can be easily generalized to arbitrary probability distributions—and is based on reducing uncertain evidence into certain evidence on some virtual event [Pearl, 1988]. We analyze both of these methods in this paper with respect to the following four questions:

1. How should one specify uncertain evidence?
2. How should one revise a probability distribution?
3. Should, and do, iterated belief revisions commute?
4. What guarantees can be offered on the amount of belief change induced by a particular revision?

Our main findings can be summarized as follows. First, we show that Jeffrey's rule and Pearl's method both revise beliefs using the principle of probability kinematics. Jeffrey's rule explicitly commits to this principle, while Pearl's method is based on a different principle. Yet, we show that Pearl's method implies the principle of probability kinematics, leading to the same revision method as that of Jeffrey's. The difference between Jeffrey's rule and Pearl's method is in the

way uncertain evidence is specified. Jeffrey requires uncertain evidence to be specified in terms of the *effect* it has on beliefs once accepted, which is a function of both evidence strength and beliefs held before the evidence is obtained. Pearl, on the other hand, requires uncertain evidence to be specified in terms of its *strength* only. Despite this difference, we show that one can easily translate between the two methods of specifying evidence and provide the method for carrying out this translation.

The multiplicity of methods for specifying evidence also raises an important question: how should informal statements about evidence be captured formally using available methods? For example, what should the following statement translate to: "Seeing these clouds, I believe there is an 80% chance that it will rain?" We will discuss this problem of interpreting informal evidential statements later, where we emphasize its subtlety.

As to the question of iterated belief revision: It is well known that Jeffrey's rule does not commute; hence, the order in which evidence is incorporated matters [Diaconis & Zabell, 1982]. This has long been perceived as a problem, until clarified recently by the work of Wagner who observed that Jeffrey's method of specifying evidence is dependent on what is believed before the evidence is obtained and, hence, should not be commutative to start with [Wagner, 2002]. Wagner proposed a method for specifying evidence, based on the notion of *Bayes factor*, and argued that this method specifies only the strength of evidence, and is independent of the beliefs held when attaining evidence. Wagner argued that when evidence is specified in that particular way, iterated revisions should commute. He even showed that combining this method for specifying evidence with the principle of probability kinematics leads to a revision rule that commutes. We actually show that Pearl's method of virtual evidence is specifying evidence according to Bayes factor, exactly as proposed by Wagner and, hence, corresponds exactly to the proposal he calls for. Therefore, the results we discuss in this paper unify the two main methods for probabilistic belief revision proposed by Jeffrey and Pearl, and show that differences between them amount to a difference in the protocol for specifying uncertain evidence.

Our last set of results relate to the problem of providing guarantees on the amount of belief change induced by a revision. We have recently proposed a distance measure for

bounding belief changes, and showed how one can use it to provide such guarantees [Chan & Darwiche, 2002]. We show in this paper how this distance measure can be computed when one distribution is obtained from another using the principle of probability kinematics. We then show how the guarantees provided by this measure can be realized when applying either Jeffrey's rule or Pearl's method, since they both are performing revision based on the principle of probability kinematics.

2 Probability Kinematics and Jeffrey's Rule

Suppose that we have two probability distributions Pr and Pr' which disagree on the probabilities they assign to a set of mutually exclusive and exhaustive events $\gamma_1, \dots, \gamma_n$, yet:

$$Pr(\alpha | \gamma_i) = Pr'(\alpha | \gamma_i), \quad (1)$$

for $i = 1, \dots, n$, and for every event α in the probability space. We say here that Pr' is obtained from Pr by *probability kinematics* on $\gamma_1, \dots, \gamma_n$. This concept was proposed by Jeffrey [1965] to capture the notion that even though Pr and Pr' disagree on the probabilities of events γ_i , they agree on their relevance to every event α .

Consider now the problem of revising a probability distribution Pr given uncertain evidence relating to a set of mutually exclusive and exhaustive events $\gamma_1, \dots, \gamma_n$. One method of specifying uncertain evidence is through the *effect* that it would have on beliefs once accepted. That is, we can say that the evidence is such that the probability of γ_i becomes q_i once the evidence is accepted. If we adopt this method of evidence specification, we conclude that there is *only one* distribution Pr' such that:

- $Pr'(\gamma_i) = q_i$ for $i = 1, \dots, n$.
- Pr' is obtained from Pr by probability kinematics on $\gamma_1, \dots, \gamma_n$.

Moreover, this specific distribution is given by:

$$Pr'(\omega) \stackrel{\text{def}}{=} Pr(\omega) \frac{q_i}{Pr(\gamma_i)}, \quad \text{if } \omega \models \gamma_i, \quad (2)$$

where ω is an atomic event, also known as a world, and $\models$ is the logical entailment relationship. This is exactly the distribution that Jeffrey suggests and, hence, this method of revision is known as *Jeffrey's rule*. We stress here that we are drawing a distinction between the principle of probability kinematics and Jeffrey's rule, which are often considered synonymous. Specifically, Jeffrey's rule arises from a combination of two proposals: (1) the principle of probability kinematics, and (2) the specification of uncertain evidence using a posterior distribution. It is possible for one to combine the principle of probability kinematics with other methods for specifying evidence as we discuss later.

It is not hard to show that the above distribution Pr' is indeed obtained from Pr by probability kinematics on $\gamma_1, \dots, \gamma_n$, as it satisfies Equation 1:

$$\begin{aligned} Pr'(\alpha | \gamma_i) &= \frac{Pr'(\alpha, \gamma_i)}{Pr'(\gamma_i)} \\ &= \frac{\sum_{\omega \models \alpha, \gamma_i} Pr'(\omega)}{q_i} \end{aligned}$$

$$\begin{aligned} &= \frac{\sum_{\omega \models \alpha, \gamma_i} Pr(\omega) \frac{q_i}{Pr(\gamma_i)}}{q_i} \\ &= \frac{\sum_{\omega \models \alpha, \gamma_i} Pr(\omega)}{Pr(\gamma_i)} \\ &= \frac{Pr(\alpha, \gamma_i)}{Pr(\gamma_i)} \\ &= Pr(\alpha | \gamma_i). \end{aligned}$$

Therefore, for any event α , its probability under the new distribution Pr' is:

$$\begin{aligned} Pr'(\alpha) &= \sum_i Pr'(\alpha | \gamma_i) Pr'(\gamma_i) \\ &= \sum_i Pr(\alpha | \gamma_i) q_i \\ &= \sum_i q_i \frac{Pr(\alpha, \gamma_i)}{Pr(\gamma_i)}, \end{aligned}$$

which is the closed form for Jeffrey's rule. We now show an example of using Jeffrey's rule.

Example 1 (Due to Jeffrey) Assume that we are given a piece of cloth, where its color can be one of: green (c_g), blue (c_b), or violet (c_v). We want to know whether, in the next day, the cloth will be sold (s), or not sold ($\bar{s}$). Our original state of belief is given by the distribution Pr :

$$\begin{aligned} Pr(s, c_g) &= .12, & Pr(s, c_b) &= .12, & Pr(s, c_v) &= .32, \\ Pr(\bar{s}, c_g) &= .18, & Pr(\bar{s}, c_b) &= .18, & Pr(\bar{s}, c_v) &= .08. \end{aligned}$$

Therefore, our original state of belief on the color of the cloth (c_g, c_b, c_v) is given by the distribution $(.3, .3, .4)$. Assume that we now inspect the cloth by candlelight, and we want to revise our state of belief on the color of the cloth to the new distribution $(.7, .25, .05)$ using Jeffrey's rule. If we apply Jeffrey's rule (Equation 2), we get the new distribution Pr' :

$$\begin{aligned} Pr'(s, c_g) &= .28, & Pr'(s, c_b) &= .10, & Pr'(s, c_v) &= .04, \\ Pr'(\bar{s}, c_g) &= .42, & Pr'(\bar{s}, c_b) &= .15, & Pr'(\bar{s}, c_v) &= .01. \end{aligned}$$

3 Virtual Evidence and Pearl's Method

The problem of revising a probability distribution under uncertain evidence can be approached from a different perspective than that of probability kinematics. Specifically, when we have uncertain evidence about some mutually exclusive and exhaustive events $\gamma_1, \dots, \gamma_n$, we can interpret that evidence as *hard evidence* on some *virtual event* η , where the relevance of $\gamma_1, \dots, \gamma_n$ to η is uncertain. It is assumed that the virtual event η depends only on the events $\gamma_1, \dots, \gamma_n$ and, therefore, is independent of any other event α given γ_i :

$$Pr(\eta | \gamma_i, \alpha) = Pr(\eta | \gamma_i). \quad (3)$$

According to this approach, the uncertainty regarding evidence on $\gamma_1, \dots, \gamma_n$ is recast as uncertainty in the relevance of $\gamma_1, \dots, \gamma_n$ to the virtual event η . Specifically, the uncertainty is recast as the likelihood of γ_i given virtual evidence η : $Pr(\eta | \gamma_i)$, for $i = 1, \dots, n$.

We next show that the new distribution obtained after accepting the uncertain evidence on $\gamma_1, \dots, \gamma_n$, $Pr(\cdot | \eta)$, is:

$$Pr(\omega | \eta) = Pr(\omega) \frac{\lambda_i}{\sum_j Pr(\gamma_j) \lambda_j}, \text{ if } \omega \models \gamma_i, \quad (4)$$

where $\lambda_1, \dots, \lambda_n$ are ratios chosen such that:

$$\lambda_1 : \dots : \lambda_n = Pr(\eta | \gamma_1) : \dots : Pr(\eta | \gamma_n).$$

Hence, the specific likelihoods $Pr(\eta | \gamma_i)$ are not important here, but their ratios are. This is why this method usually specifies uncertain evidence using a set of likelihood ratios $\lambda_1, \dots, \lambda_n$ [Pearl, 1988]. The derivation of Equation 4 is based on the assumptions underlying the method of virtual evidence given by Equation 3:

$$\begin{aligned} Pr(\omega | \eta) &= \frac{Pr(\omega, \eta)}{Pr(\eta)} \\ &= \frac{Pr(\eta | \omega) Pr(\omega)}{\sum_j Pr(\eta | \gamma_j) Pr(\gamma_j)} \\ &= \frac{Pr(\eta | \gamma_i, \omega) Pr(\omega)}{\sum_j Pr(\eta | \gamma_j) Pr(\gamma_j)} \\ &= \frac{Pr(\eta | \gamma_i) Pr(\omega)}{\sum_j Pr(\eta | \gamma_j) Pr(\gamma_j)} \\ &= Pr(\omega) \frac{\lambda_i}{\sum_j Pr(\gamma_j) \lambda_j}, \text{ if } \omega \models \gamma_i. \end{aligned}$$

The last step is based on $Pr(\eta | \gamma_i) = k\lambda_i$, where k is a constant. For any event α , the new probability after accommodating the virtual evidence is:

$$\begin{aligned} Pr(\alpha | \eta) &= \sum_i Pr(\alpha, \gamma_i | \eta) \\ &= \sum_i Pr(\alpha, \gamma_i) \frac{\lambda_i}{\sum_j \lambda_j Pr(\gamma_j)} \\ &= \frac{\sum_i \lambda_i Pr(\alpha, \gamma_i)}{\sum_j \lambda_j Pr(\gamma_j)}, \end{aligned}$$

which is the closed form for Pearl's method.

The above revision method is a generalization of the method of virtual evidence proposed by Pearl [1988] in the context of belief networks. A belief network is a graphical probabilistic model, composed of two parts: a directed acyclic graph where nodes represent variables, and a set of conditional probability tables (CPTs), one for each variable [Pearl, 1988; Jensen, 2001]. The CPT for variable X with parents U defines a set of conditional probabilities of the form $Pr(x | u)$, where x is a value of variable X , and u is an instantiation of parents U . Suppose now that we have some virtual evidence bearing on variable Y , which has values $y_1, \dots, y_n$. This virtual evidence is represented in the belief network by adding a dummy node Z and a directed edge $Y \rightarrow Z$, where one value of Z , say z , corresponds to the virtual event η . This ensures the assumption of Equation 3, that that virtual event z is independent of every other event α given every y_i , i.e., $Pr(z | y_i, \alpha) = Pr(z | y_i)$, which follows from the independence semantics of belief networks [Pearl, 1988].

The uncertainty of evidence is quantified by the likelihood ratios: $Pr(z | y_1) : \dots : Pr(z | y_n) = \lambda_1 : \dots : \lambda_n$, which are specified in the CPT of variable Z . Finally, we incorporate the presence of the virtual event z by adding the observation $Z = z$ to the rest of evidence in the belief network. We now show a simple example.

Example 2 (Due to Pearl) We are given a belief network with two variables: A represents whether the alarm of Mr. Holmes' house goes off, and B represents whether there is a burglary. To represent the influence between the two variables, there is a directed edge $B \rightarrow A$. The CPTs of A and B are given by: $Pr(a | b) = .95$, meaning the alarm goes off if there is a burglary with probability .95; $Pr(a | \bar{b}) = .01$, meaning the alarm goes off if there is no burglary with probability .01; and $Pr(b) = 10^{-4}$, meaning on any given day, there is a burglary on any given house with probability 10^{-4} .

One day, Mr. Holmes' receives a call from his neighbor, Mrs. Gibbons, saying she may have heard the alarm of his house going off. Since Mrs. Gibbons suffers from a hearing problem, Mr. Holmes concludes that there is an 80% chance that Mrs. Gibbons did hear the alarm going off. According to the method of virtual evidence, this uncertain evidence can be specified by the virtual event η and the likelihood ratio: $Pr(\eta | a) : Pr(\eta | \bar{a}) = 4 : 1$. To incorporate the virtual evidence into the belief network, we add the variable Z , and the directed edge $A \rightarrow Z$, and specify the CPT of Z such that $Pr(z | a) : Pr(z | \bar{a}) = 4 : 1$. For example, we can assign $Pr(z | a) = .4$ and $Pr(z | \bar{a}) = .1$. After incorporating the virtual evidence by adding the observation $Z = z$ to the evidence, we can easily compute the answers to queries in the belief network. For example, the probability that there is a burglary at Mr. Holmes' house is now $Pr(b | z) \approx 3.85 \times 10^{-4}$.

4 Comparing the Revision Methods

From the illustrations of the two belief revision methods, Jeffrey's rule and Pearl's method of virtual evidence, we can see that a belief revision method can be broken into two parts: a formal method of specifying uncertain evidence, and a principle of belief revision that commits to a unique distribution among many which satisfy the uncertain evidence.

4.1 Pearl's method and Probability Kinematics

We now show that Pearl's method, like Jeffrey's rule, also obeys the principle of probability kinematics; what they differ in is how uncertain evidence is specified.

Suppose that a probability distribution Pr was revised using the method of virtual evidence, with likelihood ratios $\lambda_1, \dots, \lambda_n$ bearing on events $\gamma_1, \dots, \gamma_n$, obtaining the new distribution $Pr(\cdot | \eta)$. We can easily see that the revision satisfies the principle of probability kinematics, i.e. Equation 1:

$$\begin{aligned} Pr(\alpha | \gamma_i, \eta) &= \frac{Pr(\alpha, \gamma_i | \eta)}{Pr(\gamma_i | \eta)} \\ &= \frac{Pr(\alpha, \gamma_i) \frac{\lambda_i}{\sum_j Pr(\gamma_j) \lambda_j}}{Pr(\gamma_i) \frac{\lambda_i}{\sum_j Pr(\gamma_j) \lambda_j}} \end{aligned}$$

$$\begin{aligned}
&= \frac{Pr(\alpha, \gamma_i)}{Pr(\gamma_i)} \\
&= Pr(\alpha | \gamma_i).
\end{aligned}$$

Therefore, both Jeffrey's rule and Pearl's method obey the principle of probability kinematics for belief revision.

4.2 From Pearl's Method to Jeffrey's Rule

With the previous result, we now show how we can easily translate between the two methods of specifying uncertain evidence. For example, to translate from Pearl's method to Jeffrey's rule, we note that the new probabilities of events $\gamma_1, \dots, \gamma_n$ after applying virtual evidence, with likelihood ratios $\lambda_1, \dots, \lambda_n$, are given by:

$$\begin{aligned}
Pr(\gamma_i | \eta) &= \sum_{\omega \models \gamma_i} Pr(\omega | \eta) \\
&= \sum_{\omega \models \gamma_i} Pr(\omega) \frac{\lambda_i}{\sum_j Pr(\gamma_j) \lambda_j} \\
&= Pr(\gamma_i) \frac{\lambda_i}{\sum_j Pr(\gamma_j) \lambda_j}.
\end{aligned}$$

Suppose instead that we want to revise the distribution Pr using Jeffrey's rule, assuming that after accepting uncertain evidence, the probability of γ_i becomes:

$$q_i = Pr(\gamma_i) \frac{\lambda_i}{\sum_j Pr(\gamma_j) \lambda_j}.$$

Substituting the above probability q_i in Jeffrey's rule (Equation 2), we get:

$$\begin{aligned}
Pr'(\omega) &= Pr(\omega) \frac{Pr(\gamma_i) \sum_j \frac{\lambda_i}{Pr(\gamma_j) \lambda_j}}{Pr(\gamma_i)} \\
&= Pr(\omega) \frac{\lambda_i}{\sum_j Pr(\gamma_j) \lambda_j}, \text{ if } \omega \models \gamma_i,
\end{aligned}$$

which is exactly the distribution obtained by the method of virtual evidence (Equation 4). We now illustrate this translation by revisiting Example 2.

Example 3 In Example 2, we applied the method of virtual evidence on Pr , by specifying the virtual evidence: $Pr(\eta | a) : Pr(\eta | \bar{a}) = \lambda_a : \lambda_{\bar{a}} = 4 : 1$. The original probabilities of a and $\bar{a}$ are given by: $Pr(a) = .010094$, and $Pr(\bar{a}) = .989906$. After applying the virtual evidence, the new probabilities of a and $\bar{a}$ are given by:

$$\begin{aligned}
Pr(a | \eta) &= Pr(a) \frac{\lambda_a}{Pr(a) \lambda_a + Pr(\bar{a}) \lambda_{\bar{a}}} \approx .039189; \\
Pr(\bar{a} | \eta) &= Pr(\bar{a}) \frac{\lambda_{\bar{a}}}{Pr(a) \lambda_a + Pr(\bar{a}) \lambda_{\bar{a}}} \approx .960811.
\end{aligned}$$

Alternatively, we can apply Jeffrey's rule to obtain the new distribution Pr' such that $Pr'(a) = Pr(a | \eta) \approx .039189$ and $Pr'(\bar{a}) = Pr(\bar{a} | \eta) \approx .960811$, and Pr' will be the same to the distribution $Pr(\cdot | \eta)$ obtained by virtual evidence.

4.3 From Jeffrey's Rule to Pearl's Method

We can also easily translate from Jeffrey's rule to Pearl's method. Specifically, the new probabilities of events $\gamma_1, \dots, \gamma_n$ after applying Jeffrey's rule are given by:

$$Pr'(\gamma_i) = q_i.$$

Suppose instead that we want to revise the distribution Pr using Pearl's method. We can do this by applying the method of virtual evidence, with likelihood ratios $\lambda_1, \dots, \lambda_n$, such that:

$$\lambda_1 : \dots : \lambda_n = \frac{q_1}{Pr(\gamma_1)} : \dots : \frac{q_n}{Pr(\gamma_n)}.$$

Therefore, $\lambda_i = kq_i/Pr(\gamma_i)$, where k is a constant. Substituting λ_i in Pearl's method (Equation 4), we get:

$$\begin{aligned}
Pr(\omega | \eta) &= Pr(\omega) \frac{\frac{kq_i}{Pr(\gamma_i)}}{\sum_j Pr(\gamma_j) \frac{kq_j}{Pr(\gamma_j)}} \\
&= Pr(\omega) \frac{q_i}{Pr(\gamma_i)}, \text{ if } \omega \models \gamma_i,
\end{aligned}$$

because $\sum_j q_j = 1$. This is exactly the distribution obtained by Jeffrey's rule (Equation 2). We now illustrate this translation by revisiting Example 1.

Example 4 In Example 1, we applied Jeffrey's rule on the original distribution Pr to obtain new distribution Pr' . We can also do this by applying the method of virtual evidence. We can interpret the inspection of the cloth by candlelight as virtual evidence η , where $Pr(\eta | c_g) : Pr(\eta | c_b) : Pr(\eta | c_v) = \lambda_g : \lambda_b : \lambda_v = 7 : 2.5 : .375$.¹ The probability distribution of the color of the cloth after applying virtual evidence is therefore given by:

$$\begin{aligned}
Pr(c_g | \eta) &= \frac{Pr(c_g) \lambda_g}{Pr(c_g) \lambda_g + Pr(c_b) \lambda_b + Pr(c_v) \lambda_v} = .7; \\
Pr(c_b | \eta) &= \frac{Pr(c_b) \lambda_b}{Pr(c_g) \lambda_g + Pr(c_b) \lambda_b + Pr(c_v) \lambda_v} = .25; \\
Pr(c_v | \eta) &= \frac{Pr(c_v) \lambda_v}{Pr(c_g) \lambda_g + Pr(c_b) \lambda_b + Pr(c_v) \lambda_v} = .05.
\end{aligned}$$

We can easily verify that this probability distribution $Pr(\cdot | \eta)$ obtained by virtual evidence is the same as the distribution Pr' obtained by Jeffrey's rule.

5 Interpreting Evidential Statements

The evidence specification protocols adopted by Jeffrey's rule and Pearl's method have been discussed by Pearl [2001], in relation to the problem of formally interpreting evidential statements. Consider the following statement as an example:

Looking at this evidence, I am willing to bet 2:1 that David is not the killer.

This statement can be formally interpreted using either protocol. For example, if α denotes the event "David is not the killer," this statement can be interpreted in two ways:

¹The λ 's are chosen such that $\lambda_i = kPr'(c_i)/Pr(c_i)$, where k is a constant.

1. After accepting the evidence, the probability that David is not the killer becomes twice the probability that David is the killer: $Pr'(\alpha) = 2/3$ and $Pr'(\bar{\alpha}) = 1/3$.
2. The probability that I will see this evidence η given that David is not the killer is twice the probability that I will see it given that David is the killer: $Pr(\eta | \alpha) : Pr(\eta | \bar{\alpha}) = 2 : 1$.

The first interpretation translates directly into a formal piece of evidence, Jeffrey's style, and can be characterized as an "All things considered" interpretation since it is a statement about the agent's final beliefs, which are a function of both his prior beliefs and the evidence [Pearl, 2001]. On the other hand, the second interpretation translates directly into a formal piece of evidence, Pearl's style, and can be characterized as a "Nothing else considered" interpretation since it is a statement about the evidence only [Pearl, 2001].

The two interpretations can lead to contradictory conclusions about the evidence. For example, if we use the "Nothing else considered" approach to interpret our statement, we will conclude that the evidence is against David being the killer. However, if we use the "All things considered" interpretation, it is not clear whether the evidence is for or against, unless we know the original probability that David is the killer. If, for example, David is one of four suspects who are equally likely to be the killer, we will originally have: $Pr(\alpha) = 3/4$. Therefore, this evidence has actually increased the probability that David is the killer! Because of this, Pearl argues for the "Nothing else considered" interpretation, as it provides a summary of the evidence and the evidence alone, and discusses how people tend to use betting odds to quantify their beliefs even when they are based on the evidence only [Pearl, 2001].

Example 2 provides another opportunity to illustrate the subtlety involved in interpreting evidential statements. The evidential statement in this case is "Mr. Holmes concludes that there is an 80% chance that Mrs. Gibbons did hear the alarm going off." Interpreting this statement using the "All things considered" approach gives us the conclusion that $Pr'(a) : Pr'(\bar{a}) = 4 : 1$, where a denotes the event that the alarm has gone off. This interpretation assumes that the 4 : 1 ratio applies to the posterior beliefs in a and $\bar{a}$, after Mr. Holmes has accommodated the evidence provided by Mrs. Gibson. However, in Example 2, this statement was given a "Nothing else considered" interpretation, as by Pearl [1988, Page 44-47], where the 4 : 1 ratio is taken as a quantification of the evidence strength. That is, the statement is interpreted as $Pr(\eta | a) : Pr(\eta | \bar{a}) = 4 : 1$, where η stands for the evidence. In fact, the two interpretations will lead to two different probability distributions and, hence, give us different answers to further probabilistic queries. For example, if we use the "All things considered" approach in interpreting this evidential statement, the probability of having a burglary will be $Pr'(b) = 7.53 \times 10^{-3}$, which is much larger than the probability we get using the "Nothing else considered" approach in Example 2, which is 3.85×10^{-4} .

From the discussions above, the formal interpretation of evidential statements appears to be a non-trivial task, which can be sensitive to context and communication protocols. Re-

gardless of how this is accomplished though, we need to stress that the process of mapping an informal evidential statement into a revised probability distribution involves three distinct elements:

1. One must adopt a formal method for specifying evidence.
2. One must interpret the informal evidential statement, by translating it into a formal piece of evidence.
3. One must apply a revision, by mapping the original probability distribution and formal piece of evidence into a new distribution.

Our main point here is that Jeffrey's rule and Pearl's method employs the same belief revision principle, i.e. probability kinematics. Moreover, although they adopt different, formal methods for specifying evidence, one can translate between the two methods of specification. Finally, one may take the view that Jeffrey's rule and Pearl's method constitute commitments to how informal evidential statements should be interpreted, making the first and second elements the same, but we do not take this view here. Instead, we regard the issue of interpretation as a third, orthogonal dimension which is best addressed independently.

6 Commutativity of Iterated Revisions

We now discuss the problem of commutativity of iterated revisions, that is, whether the order in which we incorporate uncertain evidence matters.²

It is well known that iterated revisions by Jeffrey's rule are not commutative [Diaconis & Zabell, 1982]. As a simple example, assume that we are given a piece of uncertain evidence which suggests that the probability of event α is .7, followed by another piece of uncertain evidence which suggests that the probability of α is .8. After incorporating both pieces in that order, we will believe that the probability of α is .8. However, if the reversed order of revision is employed, we will believe that the probability of α is .7. In general, even if we are given pieces of uncertain evidence on different events, iterated revisions by Jeffrey's rule are not commutative.

This was viewed as a problematic aspect of Jeffrey's rule for a long time, until clarified recently by Wagner [2002]. First, Wagner observed and stressed that the evidence specification method adopted by Jeffrey is suitable for the "All things considered" interpretation of evidential statements. Moreover, he argued convincingly that when evidential statements carry this interpretation, they must not be commutative to start with. So the lack of commutativity is not a problem of the revision method, but a property of the method used to specify evidence. In fact, Wagner suggested a third method for specifying evidence based on Bayes factors [Good, 1950; 1983; Jeffrey, 1992], which leads to commutativity. Specifically, if Pr and Pr' are two probability distributions, and

²There is a key distinction between iterated revisions using certain evidence versus uncertain evidence. In the former case, pieces of evidence may be logically inconsistent, which adds another dimension of complexity to the problem [Darwiche & Pearl, 1997], leading to different properties and treatments.

γ_i and γ_j are two events, the *Bayes factor* (or odds factor), $\mathcal{F}_{Pr', Pr}(\gamma_i : \gamma_j)$, is defined as the ratio of new-to-old odds:

$$\mathcal{F}_{Pr', Pr}(\gamma_i : \gamma_j) \stackrel{\text{def}}{=} \frac{Pr'(\gamma_i)/Pr'(\gamma_j)}{Pr(\gamma_i)/Pr(\gamma_j)}.$$

Given this notion, one can specify uncertain evidence on a set of mutually exclusive and exhaustive events $\gamma_1, \dots, \gamma_n$ by specifying the Bayes factor for every pair of events γ_i and γ_j . One interesting property of this method of specification is that Bayes factors do not constrain the probability distribution Pr , i.e., any evidence specified in this way is compatible with every distribution Pr .³ Hence, they are suitable for a “Nothing else considered” interpretation of evidential statements.

Interestingly enough, Wagner [2002] showed that when evidence is specified using Bayes factors and revisions are accomplished by probability kinematics, belief revision becomes commutative.⁴ This was an illustration that Jeffrey’s rule is indeed composed of two independent elements: an evidence specification method, and a revision method.

In fact, we now show that the evidence specification method adopted by Pearl’s method corresponds to the method of specifying evidence by Bayes factors. This has a number of implications. First, it shows that revisions by the virtual evidence method are commutative. Second, it provides an alternate, more classical, semantics for the virtual evidence method. Finally, it shows again that Jeffrey’s rule and Pearl’s method both revise distributions using probability kinematics.

Specifically, suppose that we revise a probability distribution Pr using the method of virtual evidence, with likelihood ratios $\lambda_1, \dots, \lambda_n$ bearing on events $\gamma_1, \dots, \gamma_n$, and suppose that we get distribution Pr' as a result. We then have:

$$\begin{aligned} \mathcal{F}_{Pr', Pr}(\gamma_i : \gamma_j) &= \frac{Pr'(\gamma_i)/Pr'(\gamma_j)}{Pr(\gamma_i)/Pr(\gamma_j)} \\ &= \frac{Pr(\gamma_i)\lambda_i / (\sum_k Pr(\gamma_k)\lambda_k)}{Pr(\gamma_j)\lambda_j / (\sum_k Pr(\gamma_k)\lambda_k)} \\ &= \frac{Pr(\gamma_i)}{Pr(\gamma_j)} \cdot \frac{\lambda_i}{\lambda_j}. \end{aligned}$$

That is, if we specify evidence by Bayes factor λ_i/λ_j for every pair of events γ_i and γ_j , and then revise distribution Pr using probability kinematics, we obtain the same distribution that arises from applying the virtual evidence method.

7 Bounding Belief Change Induced by Probability Kinematics

One important question relating to belief revision is that of measuring the extent to which a revision disturbs existing

³This is not true if we use ratios of probabilities instead of ratios of odds. For example, if $Pr'(\alpha) = 2Pr(\alpha)$, we must have $Pr(\alpha) \leq .5$ because $Pr'(\alpha) \leq 1$ [Wagner, 2002].

⁴Wagner shows not only that the representation of uncertain evidence using Bayes factors is sufficient for commutativity, but in a large number of cases, necessary.

beliefs. We have recently proposed a distance measure defined between two probability distributions which can be used to bound the amount of belief change induced by a revision [Chan & Darwiche, 2002]. We review this measure next and then use it to provide some guarantees on any revision which is based on probability kinematics.⁵

Definition 1 [Chan & Darwiche, 2002] Let Pr and Pr' be two probability distributions over the same set of worlds ω . We define a measure $D(Pr, Pr')$ as follows:

$$D(Pr, Pr') \stackrel{\text{def}}{=} \ln \max_{\omega} \frac{Pr'(\omega)}{Pr(\omega)} - \ln \min_{\omega} \frac{Pr'(\omega)}{Pr(\omega)},$$

where $0/0$ is defined as 1.

This distance measure can also be expressed using the Bayes factor:

$$D(Pr, Pr') = \ln \max_{\omega_i, \omega_j} \mathcal{F}_{Pr', Pr}(\omega_i : \omega_j).$$

This measure satisfies the three properties of distance: positiveness, symmetry, and the triangle inequality. It is useful to compute this distance measure between two probability distributions as it allows us to bound the difference in beliefs captured by them.

Theorem 1 [Chan & Darwiche, 2002] Let Pr and Pr' be two probability distributions over the same set of worlds. Let α and β be two events. We then have:

$$e^{-D(Pr, Pr')} \leq \frac{O'(\alpha | \beta)}{O(\alpha | \beta)} \leq e^{D(Pr, Pr')},$$

where $O(\alpha | \beta) = Pr(\alpha | \beta)/Pr(\bar{\alpha} | \beta)$ is the odds of event α given β with respect to Pr , and $O'(\alpha | \beta) = Pr'(\alpha | \beta)/Pr'(\bar{\alpha} | \beta)$ is the odds of event α given β with respect to Pr' .⁶ The bound is tight in the sense that for every pair of distributions Pr and Pr' , there are events α and β such that:

$$\frac{O'(\alpha | \beta)}{O(\alpha | \beta)} = e^{D(Pr, Pr')}; \quad \frac{O'(\bar{\alpha} | \beta)}{O(\bar{\alpha} | \beta)} = e^{-D(Pr, Pr')}.$$

According to Theorem 1, if we are able to compute the distance measure between the original and revised distributions, we can get a tight bound on the new belief in any conditional event given our original belief in that event. The following theorem computes this distance measure for belief revision methods based on probability kinematics.

Theorem 2 If Pr' comes from Pr by probability kinematics on $\gamma_1, \dots, \gamma_n$, the distance measure between Pr and Pr' is given by:

$$D(Pr, Pr') = \ln \max_i \frac{Pr'(\gamma_i)}{Pr(\gamma_i)} - \ln \min_i \frac{Pr'(\gamma_i)}{Pr(\gamma_i)}.$$

⁵The results in this section are reformulations of previous results [Chan & Darwiche, 2002], and are inspired by a new understanding of Jeffrey’s rule and Pearl’s method as two specific instances of revision based on probability kinematics, and the understanding of Pearl’s method in terms of Bayes factors.

⁶Of course, we must have $Pr(\beta) \neq 0$ and $Pr'(\beta) \neq 0$ for the odds to be defined.

Theorem 2 allows us to compute the distance measure for revisions based on Jeffrey's rule and virtual evidence.

Corollary 1 *If Pr' comes from Pr by applying Jeffrey's rule to events $\gamma_1, \dots, \gamma_n$, with posterior probabilities $q_1, \dots, q_n$, the distance measure between Pr and Pr' is given by:*

$$D(Pr, Pr') = \ln \max_i \frac{q_i}{Pr(\gamma_i)} - \ln \min_i \frac{q_i}{Pr(\gamma_i)}.$$

Corollary 2 *If Pr' comes from Pr by applying the method of virtual evidence, with likelihood ratios $\lambda_1, \dots, \lambda_n$ bearing on $\gamma_1, \dots, \gamma_n$, the distance measure between Pr and Pr' is given by:*

$$D(Pr, Pr') = \ln \max_i \lambda_i - \ln \min_i \lambda_i.$$

The significance of Corollaries 1 and 2 is that we can compute the distance measure easily in both cases. For Jeffrey's rule, we can compute the distance measure by knowing only the prior and posterior probabilities of events $\gamma_1, \dots, \gamma_n$. For Pearl's method, we can compute the distance measure by knowing only the likelihood ratios $\lambda_1, \dots, \lambda_n$. For both revision methods, the distance measure can be computed in constant time from the uncertain evidence, and we can guarantee a bound on the belief change due to the revision principle of probability kinematics, without explicitly knowing the revised probability distribution.

We close this section by showing that the principle of probability kinematics is optimal in a very precise sense: it commits to a probability distribution that minimizes our distance measure.

Theorem 3 *If Pr' comes from Pr by probability kinematics on $\gamma_1, \dots, \gamma_n$, it is optimal in the following sense. Among all possible distributions that agree with Pr' on the probabilities of events $\gamma_1, \dots, \gamma_n$, Pr' is the closest to Pr according to the distance measure given by Definition 1.*

8 Conclusion

In this paper, we analyzed two main methods for revising probability distributions given uncertain evidence: Jeffrey's rule and Pearl's method of virtual evidence. We showed that the two methods use the same belief revision principle, i.e. probability kinematics, with their difference being only in the manner in which they specify uncertain evidence, and showed how to translate between the two methods for specifying evidence. We also discussed the much debated problem of interpreting evidential statements. Moreover, we showed that the method of virtual evidence can be reformulated in terms of Bayes factors, which implies a number of results, including the commutativity of revisions based on this method. Finally, we showed that revisions based on probability kinematics are optimal in a very specific way, and pointed to a distance measure for bounding belief change triggered by any revision based on probability kinematics. Our bounds included Jeffrey's rule and Pearl's method as special cases.

Acknowledgments

This work has been partially supported by NSF grant IIS-9988543, MURI grant N00014-00-1-0617, and DiMI grant

00-10065. We also wish to thank Carl Wagner, Judea Pearl, and the late Richard Jeffrey, for their valuable comments and discussions.

References

- [Chan & Darwiche, 2002] Hei Chan and Adnan Darwiche. A distance measure for bounding probabilistic belief change. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence*, pages 539–545, Edmonton, Canada, July 2002.
- [Darwiche & Pearl, 1997] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artificial Intelligence* 87(1–2):1–29, January 1997.
- [Diaconis & Zabell, 1982] Persi Diaconis and Sandy L. Zabell. Updating subjective probability. *Journal of the American Statistical Association* 77:822–830, 1982.
- [Good, 1950] Irving J. Good. *Probability and the Weighing of Evidence*. Charles Griffin, London, 1950.
- [Good, 1983] Irving J. Good. *Good Thinking: the Foundations of Probability and Its Applications*. University of Minnesota Press, Minneapolis, 1983.
- [Jeffrey, 1965] Richard C. Jeffrey. *The Logic of Decision*. McGraw-Hill, New York, 1965. (2nd edition) University of Chicago Press, Chicago, 1983. (Paperback correction) 1990.
- [Jeffrey, 1992] Richard C. Jeffrey. *Probability and the Art of Judgement*. Cambridge University Press, Cambridge, 1992.
- [Jensen, 2001] Finn V. Jensen. *Bayesian Networks and Decision Graphs*. Springer-Verlag, New York, 2001.
- [Pearl, 1988] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Francisco, 1988.
- [Pearl, 2001] Judea Pearl. On two pseudo-paradoxes in Bayesian analysis. *Annals of Mathematics and Artificial Intelligence* 32:171–177, 2001.
- [Wagner, 2002] Carl Wagner. Probability kinematics and commutativity. *Philosophy of Science* 69:266–278, 2002.