Towards a Unifying Theory of Logical and Probabilistic Reasoning

Rolf Haenni
University of Berne
CH-3012 Berne, Switzerland
haenni@iam.unibe.ch

Abstract

Logic and probability theory have both a long history in science. They are mainly rooted in philosophy and mathematics, but are nowadays important tools in many other fields such as computer science and, in particular, artificial intelligence. Some philosophers studied the connection between logical and probabilistic reasoning, and some attempts to combine these disciplines have been made in computer science, but logic and probability theory are still widely considered to be separate theories that are only loosely connected. This paper introduces a new perspective which shows that logical and probabilistic reasoning are no more and no less than two opposite extreme cases of one and the same universal theory of reasoning called probabilistic argumentation.[1]

Keywords. Logical reasoning, probabilistic reasoning, uncertainty, ignorance, argumentation.

1 Introduction

Guessing the outcome of tossing a coin, given that the coin is known to be fair, is certainly different from guessing the outcome of tossing an unknown coin. Both are situations of uncertain reasoning, but with respect to the available information, the second one is less informative. We will refer to it as a situation of uncertain reasoning under (partial) *ignorance* or *ambiguity*.[2]

Initially, we may judge the two possible outcomes (*head* or *tail*) to be equally likely in both situations, but in the case of an unknown coin, one should be prepared to revise the judgment upon inspecting the coin or doing some experiments. As a consequence, if asked to bet money on tossing an unknown coin, people tend to either ask to see the coin or reject the bet, whereas betting on a fair coin is usually considered to be "safe". To know that the coin is fair provides thus a more solid ground for conclusions or decisions.

Classical probabilistic approaches to uncertain reasoning are based on the Bayesian probability interpretation, in which additive probabilities are used to represent degrees of belief of rational agents in the truth of statements. The problem is that probability theory, if applied in the classical way, cannot convey how much evidence one has. In the coin tossing example, this means that the same probability $\frac{1}{2}$ is assigned in all cases, although the two situations are not the same at all.

To remedy this problem, a number of techniques has been developed, e.g. *interval-valued probabilities* [15, 17, 26, 30], *second-order probabilities* [8, 16], *imprecise probabilities* [27, 28, 29], *Dempster-Shafer Theory* [3, 23], the *Transferable Belief Model* [25], the *Theory of Hints* [13], and many more. What most of these approaches have in common is that evaluating the truth of a statement yields two values, not just a single one like in classical probability theory. Depending on the context, the first value is either a *lower probability bound*, a *measure of belief*, or a *degree of support*, respectively, whereas the second one is either an *upper probability bound*, a *measure of plausibility*, or a *degree of possibility*. In all cases, it is possible to interpret the difference between the two values (i.e. the length of the corresponding interval) as a measure of ignorance relative to the hypothesis in question.

Alternatively, one can understand the two values to result from sub-additive *probabilities of provability* (or *epistemic probabilities*) [2, 7, 19, 20, 21, 24], for which the two respective values for a hypothesis and its complement do not necessarily sum up to one, i.e. where the ignorance is measured by the remaining gap.

[1] This paper is an extended version of a paper accepted at the ECSQARU'05 conference. It has been adapted for the particular interests of the SIPTA community.

[2] Some authors define *ambiguity* as the "uncertainty about probability, created by missing information that is relevant and could be known" [1, 4]. In this paper, we prefer to talk about *ignorance* as a general term for missing information.

1.1 The Basic Idea

This paper will further explore the "*probability of provability*" point of view. The idea is very simple. Consider two complementary hypothesis H and H^c, an let E denote the given evidence (what is known about the world). Instead of looking at posterior probabilities $P(H|E)$ and $P(H^c|E)$, respectively, we will look at the respective probabilities that H and H^c are *known* or *provable* in the light of E.[3] Note that under certain circumstances the evidence E may not be informative enough to prove either of them, i.e. we need to consider posterior probabilities of the events "H is provable", "H^c is provable", and "Neither of them is provable". Since these alternatives are exclusive and exhaustive, it is clear that

$$\begin{aligned}&P(H \text{ is provable}|E) +\\&P(H^c \text{ is provable}|E) +\\&P(\text{Neither of them is provable}|E) \;=\; 1.\end{aligned}$$

Furthermore, since $P(\text{Neither of them is provable}|E)$ may be positive, which implies $P(H \text{ is provable}|E) + P(H^c \text{ is provable}|E) \;\leq\; 1$, it becomes clear where the sub-additive nature of probabilities of provability comes from. It is a simple consequence of applying classical probability theory to a special class of so-called *epistemic events* [20, 21].

To illustrate this simple idea, suppose your friend Sheila flips a fair coin and promises to organize a party tomorrow night provided that the coin lands on *head*. Sheila does not say anything about what she is doing in case the coin lands on *tail*, i.e. she may or may not organize the party. What is the probability of provability that there will (or will not) be a party tomorrow night?

The given evidence E consists of two pieces. The first one is Sheila's promise, which may be appropriately encoded by a propositional sentence $head \rightarrow party$. The second one is the fact that the two possible outcomes of tossing a fair coin are known to be equally likely, which is encoded by a uniform distribution $P(head) = P(tail) = \frac{1}{2}$.

In this simple situation, we can look at *head* as a defeasible or hypothetical proof (or an argument) for *party*, and because it's the only such proof, we conclude that $P(party \text{ is provable}|E) \;=\; P(head) \;=\; \frac{1}{2}$. This means that the chance of being in a situation in which the party is certain to take place is 50%. On the other hand, because the given evidence does not allow any conclusions about $\neg party$, there is no hypothetical proof for $\neg party$, and we conclude that $P(\neg party \text{ is provable}|E) \;=\; 0$. It means that nothing speaks against the possibility that the party takes place, or in other words, the party is entirely plausible or possible.

[3] We will later give a precise definition of what is meant with "the probability that H (or H^c) is provable given E". An illustrative example is given below.

1.2 Goals and Overview

This paper introduces a new theory of uncertain reasoning under ignorance that deals with probabilities of provability. We will refer to it as the theory of *probabilistic argumentation*, in which probabilities of provability are alternatively called *degrees of support*. It turns out that probabilistic argumentation includes the two classical approaches to formal reasoning, namely logical and probabilistic reasoning, as special cases. To show how to link or unify logical and probabilistic reasoning is one of the primary goals of this paper. The key issue is to realize that the provability or the probability of a hypothesis are both degenerate cases of probabilities of provability. They result as opposite extreme cases, if one considers different sets of so-called *probabilistic variables* (see Sect. 4 for details).

The organization of this paper is bottom-up. We will start with discussing the basic concepts and properties of logical and probabilistic reasoning in Sect. 2 and 3, respectively. The goal is to emphasize the similarities and differences between these distinct types of formal reasoning. Probabilistic argumentation is introduced in Sect. 4. We will demonstrate that logical and probabilistic reasoning are contained as special cases. Finally, based on the notion of degrees of support, Sect. 5 describes uncertainty and ignorance as orthogonal aspects of our limited knowledge about the world.

2 Logical Reasoning

Logic has a long history in science dating back to the Greek philosopher Aristotle. For a long time, logic has primarily been a philosophical discipline, but nowadays it is also an important research topic in mathematics and computer science. The idea is to express knowledge by a set of sentences Σ which forms the *knowledge base*. The set $\mathcal{L}$ of all possible sentences is the corresponding *formal language*.

If another sentence $h \in \mathcal{L}$ represents a hypothesis of interest, we may want to know whether h is a *logical consequence* of Σ or not. Formally, we write $\Sigma \models h$ if h logically follows from Σ, and $\Sigma \not\models h$ otherwise. In other words, logical reasoning is concerned with the *provability* of h with respect to Σ.

If $\neg h$ denotes the complementary hypothesis of h (the sentence $\neg h$ is true whenever h is false and vice versa), then it is possible to evaluate the provability of $\neg h$ independently of the provability of h. As a consequence, we must distinguish the following four cases:[4]

- **LI**: $\Sigma \not\models h$, $\Sigma \not\models \neg h$,
  $\Rightarrow$ Σ does not allow any conclusions about h;
- **LT**: $\Sigma \models h$, $\Sigma \not\models \neg h$,
  $\Rightarrow$ h is true in the light of Σ, i.e. $\neg h$ is false;
- **LF**: $\Sigma \not\models h$, $\Sigma \models \neg h$,
  $\Rightarrow$ h is false in the light of Σ, i.e. $\neg h$ is true;
- **LC**: $\Sigma \models h$, $\Sigma \models \neg h$,
  $\Rightarrow$ the knowledge base is inconsistent or contradictory, i.e. $\Sigma \models \perp$.

Note that a definite answer is only given for LT or LF, whereas LI corresponds to a situation in which we are *ignorant* with respect to h. This means that the available information is insufficient to prove either h or its complement $\neg h$. In the particular case of $\Sigma = \emptyset$, this happens for all possible hypotheses $h \neq \top$. It reflects thus an important aspect of logical reasoning, namely that the absence of information does not allow meaningful conclusions. Finally, LC appears only if the knowledge base Σ is *inconsistent*. Inconsistent knowledge bases are usually excluded, and we will therefore consider them as *invalid*.

Suppose now that we are interested in a graded measure of provability and let's call it *degree of support* $dsp(h) \in [0,1] \cup \{undefined\}$. Intuitively, the idea of $dsp(h)$ is to express quantitatively the strength in which the available knowledge supports h. Provided that Σ is consistent, the support is maximal if h is a logical consequence of Σ, and it is minimal if Σ does not allow to infer h. In the light of this remark, we can define degree of support for logical reasoning in the following way:

$$dsp(h) \;:=\; \begin{cases} 0, \text{ if } \Sigma \not\models h \text{ and } \Sigma \not\models \perp, \\ 1, \text{ if } \Sigma \models h \text{ and } \Sigma \not\models \perp, \\ \mathit{undefined}, \text{ if } \Sigma \models \perp. \end{cases} \tag{1}$$

To further illustrate this, consider a set $V = \{X, Y\}$ of variables, corresponding sets Θ_X and Θ_Y of possible values, and an appropriate formal language $\mathcal{L}_V$ that includes statements about the variables X and Y. Vectors $\mathbf{x} = \langle x, y \rangle \in N_V$ with $N_V = \Theta_X \times \Theta_Y$ are then the possible interpretations of the language $\mathcal{L}_V$, and logical reasoning can be understood in terms of sets of interpretations instead of sets of sentences.

[4] L stands for "Logical Reasoning", T for "True", F for "False", I for "Ignorant", and C for "Contradictory".

With $N_V(\xi) \subseteq N_V$ we denote the set of all interpretations for which the sentence $\xi \in \mathcal{L}_V$ (or all sentences of a set $\xi \subseteq \mathcal{L}_V$) is true. The elements of $N_V(\xi)$ are also called *models* of ξ. Instead of Σ and h we will now work with the corresponding sets of models $E := N_V(\Sigma)$ and $H := N_V(h)$, respectively. By doing so, the consequence relation $\Sigma \models h$ translates into $E \subseteq H$. This allows us to define degree of support for logical reasoning in the following form:

$$dsp(H) \;:=\; \begin{cases} 0, \text{ if } \emptyset \neq E \not\subseteq H, \\ 1, \text{ if } \emptyset \neq E \subseteq H, \\ \mathit{undefined}, \text{ if } E = \emptyset. \end{cases} \tag{2}$$

With such a graded measure of provability in mind, we can think of the four different cases of logical reasoning according to Fig. 1. The inconsistent case LC is shown outside the unit square to indicate that degree of support is undefined.

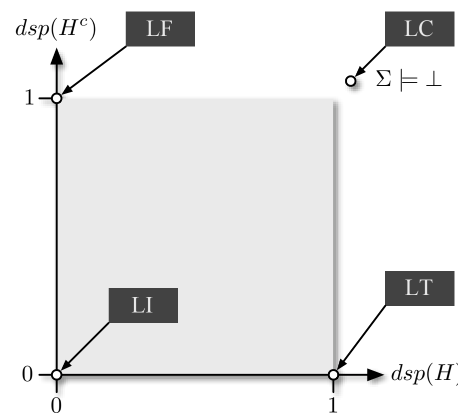

Figure 1: The four cases of logical reasoning.

Note that logical reasoning requires a *sub-additive* measure of provability, because LI is characterized by $dsp(H) = dsp(H^c) = 0$, which implies $dsp(H) + dsp(H^c) = 0$. Hence, sub-additivity turns out to be a natural property of logical reasoning, whereas probabilistic reasoning declines it (see Sect. 3).

Another important remark is to say that logical reasoning is *monotone* with respect to the available knowledge. This means that adding new sentences to Σ will never cause a transition from LT to LF or vice versa, i.e. if something is known the be true (false), it will never become false (true). The complete transition diagram for logical reasoning is shown in Fig. 2.

Note that human reasoning is not monotone at all. In fact, monotonicity is considered to be one of the major drawbacks of logical reasoning. This has led to the development of numerous non-monotone logics.

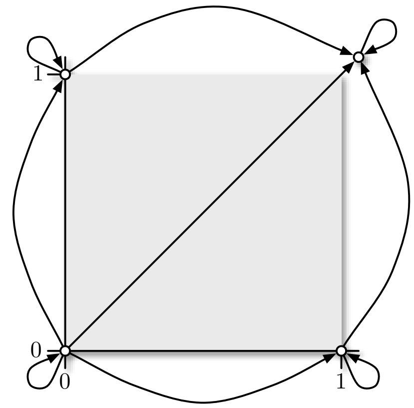

Figure 2: Logical reasoning is monotone.

3 Probabilistic Reasoning

The history of probability theory is not as old as the history of logic, but it also dates back to the 17th and 18th century. The goal is to measure the possible truth of a hypothesis H by a corresponding *posterior probability*

$$P'(H) \ = \ P(H|E) \ = \ \frac{P(H \cap E)}{P(E)} \tag{3}$$

for some observed *evidence* E. $P'(H) = 1$ means that H is true in the light of E, whereas $P'(H) = 0$ implies that H is false. Intermediate values between 0 and 1 represent all possible graded judgments between the two extremities. Probability theory is based on Kolomogorov's axioms [14], which, among other things, stipulates *additivity*. This means that $P'(H) + P'(H^c) = 1$ for all hypotheses H.

Usually, probabilistic reasoning starts with a multivariate model that consists of a *set of variables* V and a *prior probability distribution* $\mathbf{P}(V)$. A vector $\mathbf{x}$ that assigns a value to every variable of V is called *atomic event*. The hypothesis H and the evidence E are sets of such atomic events. With N_V we denote the set of all atomic events. $\mathbf{P}(V)$ can thus be regarded as an additive mapping that assigns values $P(\mathbf{x}) \in [0, 1]$ to all elements $\mathbf{x} \in N_V$.[5] With $\mathbf{P}'(V) = \mathbf{P}(V|E)$ we denote the corresponding posterior distribution for some given evidence $E \neq \emptyset$. Note that $P'(\mathbf{x}) = 0$ for all atomic events $\mathbf{x} \notin E$. The posterior distribution is undefined for $E = \emptyset$.

Prior distributions are often specified with the aid of *Bayesian networks* [19], in which the variables of the model correspond to the nodes of a directed acyclic graph. The benefits of Bayesian networks are at least twofold: first, they provide an economical way of specifying prior distributions for large sets V, and second, they play a central role for computing posterior probabilities efficiently.

For a given prior distribution $\mathbf{P}(V)$, it is possible to compute prior probabilities $P(H)$ by

$$P(H) \ = \ \sum_{\mathbf{x} \in H} P(\mathbf{x}). \tag{4}$$

Similarly, corresponding posterior probabilities are obtained for $E \neq \emptyset$ by

$$P'(H) \ = \ \sum_{\mathbf{x} \in H} P'(\mathbf{x}) \ = \ k \cdot \sum_{\mathbf{x} \in H \cap E} P(\mathbf{x}). \tag{5}$$

With $k = P(E)^{-1}$ we denote the *normalization* constant. For the particular case of $V = \{X, Y\}$, Fig. 3 illustrates the relationship between prior and posterior distributions.

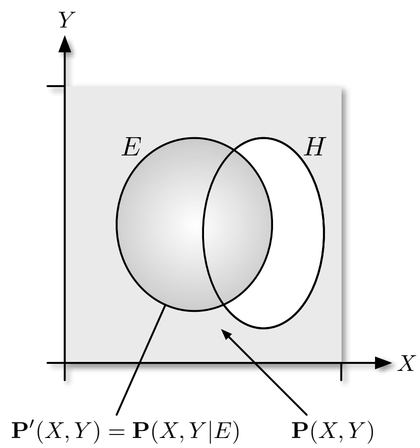

Figure 3: The relationship between prior and posterior distributions.

In order to make an analogy to the discussion of logical reasoning in Sect. 2, we can also distinguish four cases of probabilistic reasoning:[6]

- **PU**: $E \neq \emptyset$, $0 < P'(H) < 1$,
  $\Rightarrow$ $\mathbf{P}(V)$ and E do not allow a final judgment;
- **PT**: $E \neq \emptyset$, $P'(H) = 1$,
  $\Rightarrow$ H is true in the light of $\mathbf{P}(V)$ and E;
- **PF**: $E \neq \emptyset$, $P'(H) = 0$,
  $\Rightarrow$ H is false in the light of $\mathbf{P}(V)$ and E;
- **PC**: $E = \emptyset$, $P'(H) = \mathit{undefined}$,
  $\Rightarrow$ the evidence is invalid.

[5] If some or all variables of V are continuous, then *density functions* are needed to specify the prior distribution. The discussion in this paper will be limited to discrete variables, but this is not a conceptual restriction.

[6] P stands for "Probabilistic Reasoning", T for "True", F for "False", U for "Uncertain", and C for "Contradictory".

Note that there is a one-to-one correspondence between the cases PT and LT, PF and LF, and PC and LC. In the sense that both PU and LI do now allow a final judgment with respect to H, they can also be considered as similar cases. But they differ from the fact that LI, in contrast to PU, represents a state of total ignorance.

To further illustrate the analogies and differences between logical and probabilistic reasoning, all possible pairs of values $P'(H)$ and $P'(H^c)$ are shown in the picture of Fig. 4. In accordance with the discussion about logical reasoning in Sect. 2, posterior probabilities $P'(H)$ and $P'(H^c)$ will now be considered as (additive) degrees of support $dsp(H)$ and $dsp(H^c)$, respectively.

As a consequence of the additivity requirement and for $E \neq \emptyset$, all possible pairs $dsp(H)$ and $dsp(H^c)$ lie on the diagonal between the upper left and the lower right corner of the unit square shown in Fig. 4. Case LI and logical reasoning in general are thus not covered by probabilistic reasoning. On the other hand, logical reasoning does not include PU. That's the fundamental difference between logical and probabilistic reasoning.

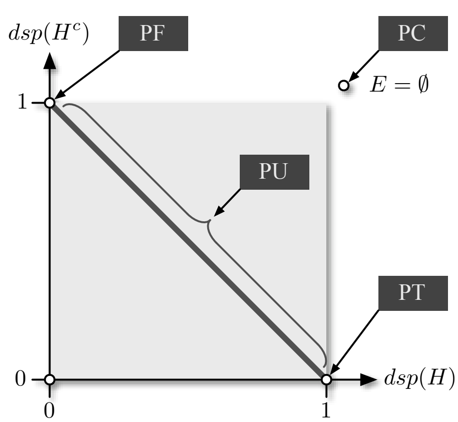

Figure 4: The four cases of probabilistic reasoning.

In comparison with logical reasoning, an important benefit of probabilistic reasoning is the fact that it is *non-monotone* with respect to the available knowledge. Except for $dsp(H) = 0$ and $dsp(H) = 1$, degree of support may thus increase or decrease when new evidence is added. As a consequence, even if a hypothesis is almost perfectly likely to be true (false), it may later turn out to be false (true). Figure 5 illustrates the possible transitions of degrees of support in the context of probabilistic reasoning.

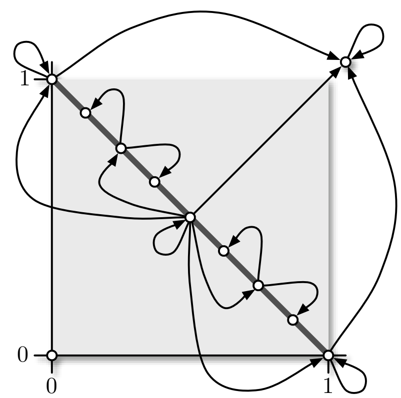

Figure 5: Probabilistic reasoning is non-monotone.

4 Probabilistic Argumentation

So far, we have tried to point out the similarities and differences between logical and probabilistic reasoning. In a nutshell, logical reasoning is sub-additive and monotone, whereas probabilistic reasoning is additive and non-monotone. Our goal now is to define a more general formal theory of reasoning that is sub-additive and non-monotone at the same time. In other words, the idea is to build a common roof for logical and probabilistic reasoning. We are thus interested in sub-additive and non-monotone degrees of support, which means that the judgment of a hypothesis may result in any point within the triangle shown in Fig. 6.

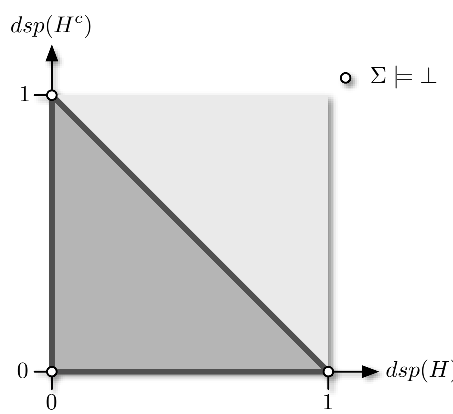

Figure 6: Sub-additive and non-monotone degrees of support.

In order to build such a unifying theory, we must first try to better understand the origin of the differences between logical and probabilistic reasoning. The key point to realize is the following: probabilistic reasoning presupposes the existence of a probability distribution over *all* variables, whereas logical reasoning does not deal with probability distributions at all, i.e. it presupposes a probability distribution over *none* of the variables involved. If we call the variables over which a probability distribution is known *probabilistic* (or *exogenous*), we can say that a probabilistic model consist of probabilistic variables only, whereas all variables of a logical model are *non-probabilistic* (or *endogenous*). From this point of view, the main difference between logical and probabilistic reasoning is the number of probabilistic variables. This simple observation turns out to be crucial for understanding the similarities and differences between logical and probabilistic reasoning.[7]

With this remark in mind, building a more general theory of reasoning is straightforward. The idea is to allow an arbitrary number of probabilistic variables. More formally, if $V = \{X_1, \ldots, X_n\}$ is the set of all variables involved, we will use $A \subseteq V$ to denote the subset of probabilistic variables. $\mathbf{P}(A)$ is the corresponding prior distribution over A. In the "Party Example" of Sect. 1, we can think of two Boolean variables *Coin* and *Party* with a given (uniform) prior distribution over *Coin*. This implies $V = \{Coin, Party\}$ and $A = \{Coin\}$.

Clearly, logical reasoning is characterized by $A = \emptyset$ and probabilistic reasoning by $A = V$, but we are now interested in the general case of arbitrary sets of probabilistic variables. We will refer to it as the theory of *probabilistic argumentation*.[8] The connection between probabilistic argumentation and the classical fields of logical and probabilistic reasoning is depicted in Fig. 7.

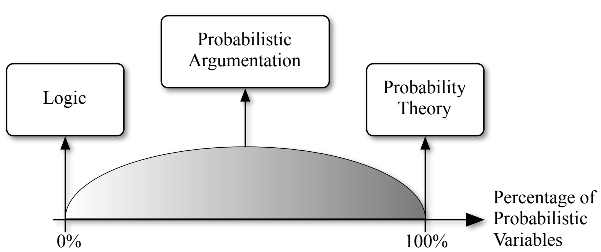

Figure 7: Different sets of probabilistic variables.

In the general context of arbitrary sets of probabilistic variables, the goal is to define degree of support as a sub-additive and non-monotone measure of uncertainty and ignorance. We suppose that our knowledge or evidence is encoded by a set of sentences $\Sigma \subseteq \mathcal{L}_V$, which determines then a set of possible atomic events $E := N_V(\Sigma) \subseteq N_V$. It is assumed, that the true state of the world is exactly one element of E. A quadruple $\mathcal{A} = (V, A, \mathbf{P}(A), \Sigma)$ is called *probabilistic argumentation system*. In the following, we will focus on sets E of atomic events rather than sets Σ of sentences. Accordingly, we will consider hypotheses $H \subseteq N_V$ rather than corresponding sentences $h \in \mathcal{L}_V$.

The definition of degree of support will be based on two observations. The first one is the fact that the set $E \subseteq N_V$, which restricts the set of possible atomic events relative to V, also restricts the atomic events relative to A. The set of all atomic events with respect to A is denoted by N_A. Its elements $\mathbf{s} \in N_A$ are called *scenarios*. $P(\mathbf{s})$ denotes the corresponding prior probability of a scenario $\mathbf{s}$. By projecting E from N_V to N_A, we get the set $E^{\downarrow A} \subseteq N_A$ of possible scenarios that are consistent with E. This means that exactly one element of $E^{\downarrow A}$ corresponds to the true state of the world. All other scenarios $\mathbf{s} \notin E^{\downarrow A}$ are impossible. This allows us to condition $\mathbf{P}(A)$ on $E^{\downarrow A}$, i.e. we need to replace the prior distribution $\mathbf{P}(A)$ by a posterior distribution $\mathbf{P}'(A) = \mathbf{P}(A|E^{\downarrow A})$. This is illustrated in Fig. 8 for the particular case of $V = \{X, Y\}$ and $A = \{X\}$.

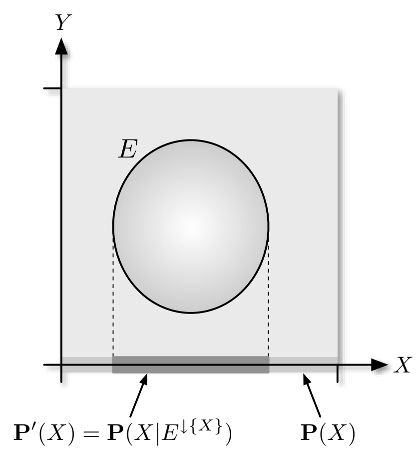

Figure 8: Prior and posterior distributions over probabilistic variables.

The second observation goes in the other direction, that is from N_A to N_V. Let's assume that a certain scenario $\mathbf{s} \in N_A$ is the true scenario. This reduces the set of possible atomic events with respect to V from E to

$$E|\mathbf{s} \ := \ \{\mathbf{x} \in E : \mathbf{x}^{\downarrow A} = \mathbf{s}\}. \qquad (6)$$

Such a set is called *conditional knowledge base* or *conditional evidence* given the scenario $\mathbf{s}$. It contains all atomic events of E that are compatible with $\mathbf{s}$. This idea is illustrated in Fig. 9 for $V = \{X, Y\}$, $A = \{X\}$, and three scenarios $\mathbf{s}_0$, $\mathbf{s}_1$, and $\mathbf{s}_2$. Note that we have $E|\mathbf{s} \neq \emptyset$ for every consistent scenario $\mathbf{s} \in E^{\downarrow A}$, e.g. for $\mathbf{s}_1$ and $\mathbf{s}_2$. Otherwise, as in the case of $\mathbf{s}_0$, we have $E|\mathbf{s} = \emptyset$

[7] The literature on how to combine logic and probability is huge, but the idea of distinguishing probabilistic and non-probabilistic variables seems to be a new one.

[8] Previous work on probabilistic argumentation focuses on propositional languages $\mathcal{L}_P$, see [12, 9, 10].

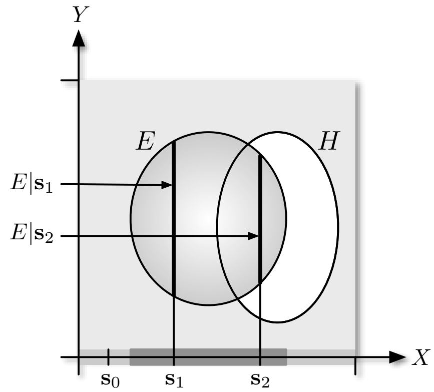

Figure 9: Evidence conditioned on various scenarios.

Consider now a consistent scenario $\mathbf{s} \in E^{\downarrow A}$ for which $E|\mathbf{s} \subseteq H$. This means that H is a logical consequence of $\mathbf{s}$ and E, and we can thus see $\mathbf{s}$ as a defeasible or hypothetical proof for H in the light of E. We must say *defeasible*, because it is unknown whether $\mathbf{s}$ is the true scenario or not. In other words, H is only *supported* by $\mathbf{s}$, but not entirely proved. The set of all *supporting scenarios* is denoted by

$$\begin{aligned} SP(H) &:= \{\mathbf{s} \in E^{\downarrow A} : E|\mathbf{s} \subseteq H\} \\ &= \{\mathbf{s} \in N_A : \emptyset \neq E|\mathbf{s} \subseteq H\}. \end{aligned} \tag{7}$$

In the example of Fig. 9, the hypothesis H is supported by $\mathbf{s}_2$, but not by $\mathbf{s}_0$ or $\mathbf{s}_1$ ($\mathbf{s}_0$ is inconsistent). Note that $SP(\emptyset) = \emptyset$ and $SP(N_V) = E^{\downarrow A}$. Sometimes, the elements of $SP(H)$ and $SP(H^c)$ are also called *arguments* and *counter-arguments* of H, respectively. This is where the name of this theory originally comes from.

The set of supporting scenarios is the key notion for the definition of degree of support. In fact, because every supporting scenario $\mathbf{s} \in SP(H)$ contributes to the possible truth of H, we can measure the strength of such a contribution by the posterior probability $P'(\mathbf{s})$, and the total support for H corresponds to the sum

$$dsp(H) := P'(SP(H)) = \sum_{\mathbf{s} \in SP(H)} P'(\mathbf{s}) \tag{8}$$

over all elements of $SP(H)$. Note that $dsp(H)$ defines an ordinary (additive) probability measure in the classical sense of Kolmogorov, that is $P'(SP(H)) + P'(SP(H)^c) = 1$ holds for all $H \subseteq N_V$. However, because the sets $SP(H)$ and $SP(H^c)$ are not necessarily complementary, we have $dsp(H) + dsp(H^c) \leq 1$ as required. Degrees of support should therefore be understood as sub-additive *posterior probabilities of provability*. Provided that $E \neq \emptyset$, they are well-defined for all possible hypotheses $H \subseteq N_V$, that is even in cases in which the prior distribution $\mathbf{P}(A)$ does not cover all variables. This is a tremendous advantage over classical probabilistic reasoning, which presupposes the existence of a prior distribution over all variables.

An alternative to considering degrees of support of complementary hypotheses is to define so-called *degrees of possibility* by $dps(H) := 1 - dsp(H^c)$. Hypotheses are then judged by a pairs of values $dsp(H) \leq dps(H)$. Note that there is a strong connection to $Bel(H)$ and $Pl(H)$ in the context of Dempster-Shafer Theory [10].

To complete this section, we will briefly investigate how the classical fields of logical and probabilistic reasoning fit into this general theory of probabilistic argumentation.

Logical Reasoning is characterized by $A = \emptyset$. This has a number of consequences. First, it implies that the set of possible scenarios $N_A = \{\langle\rangle\}$ consists of a single element $\langle\rangle$, which represents the empty vector of values. This means that $P(\langle\rangle) = 1$ is the only possible prior distribution. Second, if we assume $E \neq \emptyset$, we get $E^{\downarrow A} = \{\langle\rangle\} = N_A$ and thus $P'(\langle\rangle) = 1$. Furthermore, we have $E|\langle\rangle = E$, which allows us to rewrite (7) as

$$SP(H) = \begin{cases} \{\langle\rangle\}, & \text{for } \emptyset \neq E \subseteq H, \\ \emptyset, & \text{otherwise.} \end{cases} \tag{9}$$

Finally, $P(\langle\rangle) = 1$ implies $dsp(H) = 1$ for $\emptyset \neq E \subseteq H$, $dsp(H) = 0$ for $\emptyset \neq E \not\subseteq H$, and $dsp(H) = \textit{undefined}$ for $E = \emptyset$. This corresponds with the definition of degree of support for logical reasoning in (2).

Probabilistic Reasoning is characterized by $A = V$. This means that the sets N_A and N_V are identical, and it implies

$$E|\mathbf{s} = \begin{cases} \{\mathbf{s}\}, & \text{for } \mathbf{s} \in E, \\ \emptyset, & \text{otherwise.} \end{cases} \tag{10}$$

From this it follows that $SP(H) = H \cap E$, and it allows us to rewrite (8) as

$$dsp(H) = \sum_{\mathbf{s} \in H \cap E} P'(\mathbf{s}) = \sum_{\mathbf{s} \in H} P'(\mathbf{s}), \tag{11}$$

which corresponds to the definition of for degrees of support (posterior probabilities) in the case of probabilistic reasoning, see (5).

5 Uncertainty vs. Ignorance

The previous section has pointed out the crucial role of the probabilistic variables in the context of probabilistic argumentation. We can consider them as representatives for *atomic sources of uncertainty* for which further details are supposed to be inaccessible. The corresponding prior distribution $\mathcal{P}(A)$ summarizes or quantifies the uncertainty stemming from them. In other words, the best we can expect to know about an atomic source of uncertainty is a prior distribution. Tossing a fair coin, for example, is such an atomic source of uncertainty. It is a complex physical process whose exact details are inaccessible. But with $P(head) = P(tail) = \frac{1}{2}$ it is possible to summarize the uncertainty involved.

As a consequence of using prior probabilities as a basic ingredient of probabilistic argumentation, the best we can expect from evaluating a hypothesis H are additive degrees of support for H and H^c, respectively. Such a situation corresponds to the case of ordinary posterior probabilities. Because we cannot expect more than this, we can say that our ignorance with respect to H is *minimal* for $dsp(H)+dsp(H^c) = 1$ (true for PU, PT, PF, LT, and LF). On the other hand, $dsp(H) = dsp(H^c) = 0$ reflects a situation in which our ignorance with respect to H is *maximal* (true for LI), that is H as well as H^c are totally unsupported by the given knowledge. Note that ignorance is always relative to a hypothesis: it may well be that the available knowledge is very informative with respect to a hypothesis H, whereas another hypothesis H' is not affected at all. In the light of these remarks, we can formally define *degree of ignorance* by

$$dig(H) \ := \ 1 - dsp(H) - dsp(H^c), \tag{12}$$

which allows to characterize PU, PT, PF, LT, and LF by $dig(H) = 0$ and LI by $dig(H) = 1$. Note that the the evaluation of a hypothesis H may result in any intermediate degree of ignorance. We can thus use $dig(H)$ to gradually measure the amount of the available knowledge that is relevant for the evaluation of H. In this sense, $dig(H)$ also quantifies various levels of more or less complete knowledge bases. Figure 10 illustrates this idea.

As one would expect, $dig(H)$ typically decreases when new evidence arrives, but in general degree of ignorance is *non-monotone*. This is due to the fact that new information may independently affect the two sets $SP(H)$ and $SP(H^c)$, as well as the normalization constant and therewith the posterior distribution.

Independently of the actual degree of ignorance, we can argue that the uncertainty with respect to a hypothesis H is *maximal*, whenever $dsp(H) = dsp(H^c)$. This reflects a case in which H and H^c are equally supported by the given knowledge and the prior distribution. The two extreme cases of maximal uncertainty are $dsp(H) = dsp(H^c) = 0$ (maximal ignorance) and $dsp(H) = dsp(H^c) = 0.5$ (minimal ignorance). On the other hand, the uncertainty is *minimal*

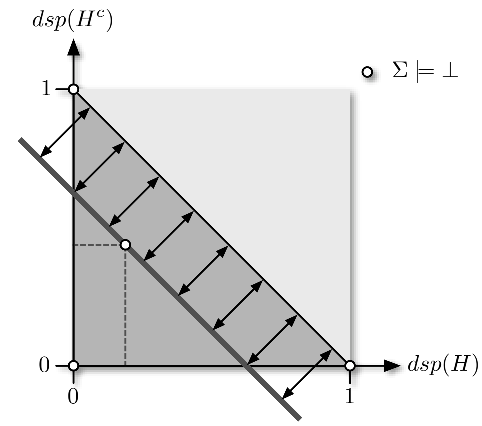

Figure 10: Various levels of ignorance.

if H is known to be either true or false. Various levels of (un-) certainty are depicted in Fig. 11. The case of maximal uncertainty corresponds to the diagonal from the lower left corner to the center of the unit square.

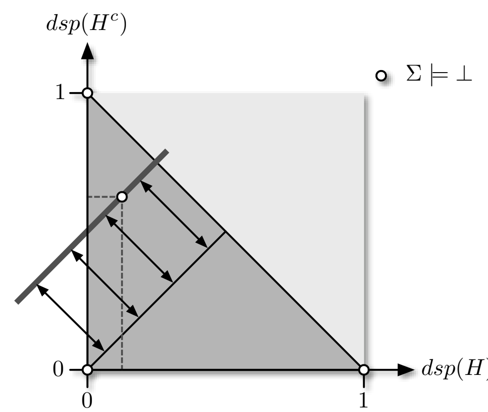

Figure 11: Various levels of uncertainty

To make this point more clear, consider four different situations of tossing a coin: a) the coin is known to be fair; b) the coin is known to be faked with heads on both sides; c) the coin is known to be faked with tails on both sides; d) the coin is unknown. In each case, we consider two possible hypotheses $H = \{head\}$ and $H^c = \{tail\}$. Corresponding probabilistic argumentation systems will then produce the following results:

a) $dsp(\{head\}) = dsp(\{tail\}) = 0.5$;
b) $dsp(\{head\}) = 1$, $dsp(\{tail\}) = 0$;
c) $dsp(\{head\}) = 0$, $dsp(\{tail\}) = 1$;
d) $dsp(\{head\}) = 0$, $dsp(\{tail\}) = 0$.

Note that these are the four extreme cases of probabilistic argumentation: the ignorance is minimal for a), b), and c), and maximal for d), whereas the uncertainty is minimal for b) and c) and maximal for a) and d).

This discussion demonstrates that uncertainty and ignorance should be considered as distinct phenomenons. In fact, as illustrated by Fig. 10 and Fig. 11, uncertainty and ignorance are *orthogonal* measures and reflect different aspects of our limited knowledge about the world. This important observation has been widely disregarded by probabilistic reasoning or Bayesianism in general.

What are the benefits of properly distinguishing between uncertainty and ignorance using a sub-additive measure of support? First, by taking into account the possibility of lacking knowledge or missing data, it provides a more realistic and more complete picture with regard to the hypothesis to be evaluated in the light of the given knowledge. Second, if reasoning deals as a preliminary step for decision making, a proper measure of ignorance is useful to decide whether the available knowledge justifies an immediate decision. The idea is that high degrees of ignorance imply low confidence in the results due to lack of information. On the other hand, low degrees of ignorance result from situations where the available knowledge forms a solid basis for a decision. Therefore, decision making should always consider the additional option of postponing the decision until enough information is gathered.[9] The study of decision theories with the option of further deliberation is a current research topic in philosophy of economics [5, 11, 22]. Apart from that, *decision making under ignorance* is a relatively unexplored discipline. For an overview of attempts in the context of Dempster-Shafer theory we refer to [18]. A more detailed discussion of such a general decision theory is beyond the scope of this paper.

6 Conclusion

This paper introduces the theory of probabilistic argumentation as a general theory of reasoning under ignorance. The key concept of the theory is the notion of degree of support, a sub-additive and non-monotone measure of uncertainty and ignorance. Degrees of support are (posterior) probabilities of provability. Obviously, this includes the notion of provability from logical reasoning, as well as the notion of probability from probabilistic reasoning. The two classical approaches to automated reasoning – logical and probabilistic reasoning – are thus special cases of probabilistic argumentation. The parameter that makes them distinct is the number of probabilistic variables. Probabilistic argumentation is more general in the sense that it allows any number of probabilistic variables.

The significance and the consequences of this paper are manyfold. In the field of automated reasoning, we consider probabilistic argumentation as a new foundation that unifies the existing approaches of logical and probabilistic reasoning. This will have a great impact on the understanding and the numerous applications of automated reasoning within and beyond AI. An important basic requirement for this is a generalized decision theory that takes into account the possibility of lacking information or missing data. Because probabilistic argumentation generally clarifies the relationship between logic and probability theory, it will also put some new light on topics from other areas such as philosophy, mathematics, or statistics. In fact, promising preliminary work shows that looking at statistical inference from the perspective of probabilistic argumentation helps to eliminate the discrepancies between the classical and the Bayesian approach to statistics.

Acknowledgements

Thanks to Michael Wachter for helpful comments. This research is supported by the Swiss National Science Foundation, Project No. PP002–102652.

References

[1] C. Camerer and M. Weber. Recent developments in modelling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty*, 50:325–370, 1992.

[2] L. J. Cohen. *The Probable and the Provable.* Clarendon Press, Oxford, 1977.

[3] A. P. Dempster. A generalization of Bayesian inference. *Journal of the Royal Statistical Society*, 30:205–247, 1968.

[4] D. Dequech. Fundamental uncertainty and ambiguity. *Eastern Economic Journal*, 26(1):41–60, 2000.

[5] I. Douven. Decision theory and the rationality of further deliberation. *Economics and Philosophy*, 18(2):303–328, 2002.

[9] It's like in real life: people do not like decisions under ignorance. In other words, people prefer betting on events they know about. This psychological phenomenon is called *ambiguity aversion* and has been experimentally demonstrated by Ellsberg [6]. His observations are rephrased in *Ellsberg's paradox*, which is often used as an argument against decision-making on the basis of subjective probabilities.

[6] D. Ellsberg. Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics*, 75:643–669, 1961.

[7] R. Fagin and J. Y. Halpern. Uncertainty, belief, and probability. *Computational Intelligence*, 7(3):160–173, 1991.

[8] H. Gaifman. A theory of higher order probabilities. In J. Y. Halpern, editor, *TARK'86, 1st Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 275–292, Monterey, CA, 1986. Morgan Kaufmann.

[9] R. Haenni, J. Kohlas, and N. Lehmann. Probabilistic argumentation systems. In J. Kohlas and S. Moral, editors, *Handbook of Defeasible Reasoning and Uncertainty Management Systems, Volume 5: Algorithms for Uncertainty and Defeasible Reasoning*, pages 221–288. Kluwer, Dordrecht, 2000.

[10] R. Haenni and N. Lehmann. Probabilistic argumentation systems: a new perspective on Dempster-Shafer theory. *International Journal of Intelligent Systems (Special Issue: the Dempster-Shafer Theory of Evidence)*, 18(1):93–106, 2003.

[11] D. Kelsey and J. Quiggin. Theories of choice under ignorance and uncertainty. *Journal of Economic Surveys*, 6(2):133–53, 1992.

[12] J. Kohlas, D. Berzati, and R. Haenni. Probabilistic argumentation systems and abduction. *Annals of Mathematics and Artificial Intelligence*, 34(1–3):177–195, 2002.

[13] J. Kohlas and P. A. Monney. *A Mathematical Theory of Hints. An Approach to the Dempster-Shafer Theory of Evidence*, volume 425 of *Lecture Notes in Economics and Mathematical Systems*. Springer-Verlag, 1995.

[14] A. N. Kolmogorov. *Foundations of the Theory of Probability*. Chelsea Publishing Company, New York, 1950.

[15] H. E. Kyburg. Higher order probabilities and intervals. *International Journal of Approximate Reasoning*, 2:195–209, 1988.

[16] H. E. Kyburg. Higher order probabilities. In L. N. Kanal, T. S. Levitt, and J. F. Lemmer, editors, *UAI '87: 3rd Conference on Uncertainty in Artificial Intelligence*, pages 15–22. Elsevier, 1989.

[17] H. E. Kyburg. Interval-valued probabilities. In G. de Cooman and P. Walley, editors, *The Imprecise Probabilities Project*. IPP Home Page, available at http://ippserv.rug.ac.be, 1998.

[18] H. T. Nguyen and E. A. Walker. On decision making using belief functions. In R. R. Yager, J. Kacprzyk, and M. Pedrizzi, editors, *Advances in the Dempster-Shafer Theory of Evidence*, pages 311–330. John Wiley and Sons, New York, 1994.

[19] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, CA, 1988.

[20] E. Ruspini, J. Lowrance, and Th. Strat. Understanding evidential reasoning. *International Journal of Approximate Reasoning*, 6(3):401–424, 1992.

[21] E. H. Ruspini. The logical foundations of evidential reasoning. Technical Report 408, AI Center, SRI International, Menlo Park, CA, 1986.

[22] D. Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57(3):571–587, 1989.

[23] G. Shafer. *The Mathematical Theory of Evidence*. Princeton University Press, 1976.

[24] Ph. Smets. Probability of provability and belief functions. *Journal de la Logique et Analyse*, 133:177–195, 1991.

[25] Ph. Smets and R. Kennes. The transferable belief model. *Artificial Intelligence*, 66:191–234, 1994.

[26] B. Tessem. Interval probability propagation. *International Journal of Approximate Reasoning*, 7:95–120, 1992.

[27] P. Walley. *Statistical Reasoning with Imprecise Probabilities*. Monographs on Statistics and Applied Probability 42. Chapman and Hall, London, UK, 1991.

[28] P. Walley. Towards a unified theory of imprecise probability. *International Journal of Approximate Reasoning*, 24(2–3):125–148, 2000.

[29] K. Weichselberger. The theory of interval-probability as a unifying concept for uncertainty. *International Journal of Approximate Reasoning*, 24(2–3):149–170, 2000.

[30] R. R. Yager and V. Kreinovich. Decision making under interval probabilities. *International Journal of Approximate Reasoning*, 22:195–215, 1999.