

A LOGIC-BASED TRACTABLE APPROXIMATION OF PROBABILITY

PAOLO BALDI AND HYKEL HOSNI

ABSTRACT. We provide a logical framework in which a resource-bounded agent can be seen to perform approximations of probabilistic reasoning. Our main results read as follows. First we identify the conditions under which propositional probability functions can be approximated by a hierarchy of depth-bounded Belief functions. Second we show that under rather palatable restrictions, our approximations of probability lead to uncertain reasoning which, under the usual assumptions in the field, qualifies as tractable.

1. INTRODUCTION AND MOTIVATION

Paraphrasing Laplace’s famous dictum, probability depends both on the information we do possess and on the information we do not. Underlying this picture are two ideas: first, probability functions quantify uncertainty with the operational meaning of “rational” degrees of belief [9, 32, 28, 30] and second, information *resolves* fully uncertainty. In other words, probability provides a principled estimate of an agent’s uncertainty about events whose truth-conditions are fully decided by information as yet not possessed by the agent.

For a straightforward illustration which is also useful to fix some basic terminology, take an urn with n balls which only differ in colour, say yellow, red and green. Assume that an agent is further equipped with the information that $1/3$ of the balls are yellow, $1/3$ are red, $1/3$ green, and is asked to draw a ball from the urn. The agent is uncertain about the colour of the particular ball that will be drawn from the urn, but knows which conditions will need to obtain for, say the event “the ball is red” to have happened. Furthermore, she knows the distribution of the balls in the urn, and will thus readily assess the likelihood of each particular outcome of this kind of experiment, i.e. will rationally assess her degree of belief on the event e.g. “the ball is red”, $Bel(Red)$ for short, equal to the proportion of red balls on the urn, that is $Bel(Red) = 1/3$. The same will clearly apply for $Bel(Green)$ and $Bel(Yellow)$.

Now, let us imagine a second scenario, where an agent only knows that the balls in the urn are either yellow, red or green, without knowing their distribution. Here, the “Laplacian” way of looking at the problem will translate the lack of information about the composition of the urn into very precise probabilistic (as it turns out) information about the event of interest: since no further information on the composition is indeed available, then there is no reason to depart from the assumption that colours are uniformly distributed in the urn, which readily gives $Bel(Red) = 1/3$, and similarly for *Green* and *Yellow*.

Both scenarios sketched so far lead to the same assignment, but relying on different kinds of information. This is however problematic, as shown by the following example, where both types of situations coexist in the same context. Take another three-colours urn composed for $1/3$ of red balls and for $2/3$ of balls which are either yellow or green, in unknown proportion. Whilst the assignment of belief to *Red* will follow the pattern of the first scenario, and, using the provided information about the proportion of red balls will lead to $Bel(Red) = 1/3$, the assignment of belief to the other two colours would follow the pattern of the second scenario. A reasoning based on the lack of information will lead indeed to $Bel(Yellow) = Bel(Green) = 1/3$, and again, we will have $Bel(Yellow) = Bel(Green) = Bel(Red) = 1/3$, with an assignment motivated by two different kinds of information (and lack thereof).

Over the past century many have taken issue with the Laplacian line of argument above. Among them [24, 35, 39] had a profound and continuing impact on the foundations of uncertain reasoning, whereas [25, 13, 33, 23] brought related problems to the attention of decision theory. More recently, similar concerns have been raised in artificial intelligence from a variety of distinct angles, among which [11, 26] had significant impact. With regard to this growing literature, readers may find out more in the general exposition provided by [21] and an up-to-date overview [29].

Whilst the formal models which originated from the above criticisms may diverge significantly, a unifying feature which is easily found in all of them is the criticism of the additivity forced upon degrees of belief by the probabilistic representation of uncertainty. This, as will be developed in the remainder of this section, can be undesirable in those situations in which telling precisely the information an agent possesses from that which it doesn’t, is neither practically nor conceptually straightforward. The research reported in this paper aims at adding a logical dimension to this long-standing issue in the foundations of uncertain reasoning.

Unlike the above mentioned criticisms, we do not take issue directly with the additivity of the measure representing uncertainty. Rather, we focus on an important but very much neglected element in the analysis of the quantification/resolution view of uncertain reasoning, namely the fact that it requires certain *inferential abilities*. These, under suitable conditions, can give rise to *information*, which eventually leads the agent’s degree of belief converge to an additive measure – i.e. a probability. By tackling the problem at its logical roots we also hope to encourage logicians to join this very interesting multidisciplinary conversation of the very foundations of rational reasoning and decision-making.

To motivate the need for a logical approach to the problem, it is useful to point out explicitly the way *classical* logic constraints (finitely additive) probability functions. So take \vdash to be the classical consequence relation on a (finite) propositional language \mathcal{L} . A *probability function on \mathcal{L}* is a normalised and finitely additive function from the sentences of the language to the real unit interval, i.e. one satisfying the following conditions:

- (PL1) If $\vdash \theta$ then $P(\theta) = 1$,
- (PL2) If $\vdash \neg(\theta \wedge \varphi)$ then $P(\theta \vee \varphi) = P(\theta) + P(\varphi)$.

The Stone representation theorem for Boolean algebras and the representation theorem for probability functions (see e.g. [30]) guarantee that the Kolmogorov measure-theoretic axiomatisation of probability is equivalent to the logical one above. Framing the axioms logically however puts forward the role of \vdash in giving probability functions an unwelcome inability to distinguish between “probabilistic knowledge” and “probabilistic ignorance”. To see this, recall that $\vdash \theta \vee \neg\theta$, and that (PL1) and (PL2) immediately give us

$$P(\neg\theta) = 1 - P(\theta). \quad (1)$$

This amounts to saying that in the framework we have no way of representing a very common situation: the agent doesn’t know anything about θ . Instead of allowing the agent to represent its lack of opinions about the event θ , (1) effectively translates absence of information about θ into information about its negation. Whilst this is easily seen to be inappropriate in many realistic scenarios, to the best of our knowledge [8] was the first paper to point out that the root of the problem, and hence a robust way to fix it, was not to be found (exclusively) in the axioms of probability functions, but at the more fundamental level of the underlying consequence relation.

This is apparent by focusing on another immediate consequence of (PL1), (PL2):

$$\theta \vdash \varphi \text{ implies } P(\theta) \leq P(\varphi), \quad (2)$$

which expresses the monotonicity of probability functions with respect to \vdash . This property is mathematically convenient, but puts on the agent a very heavy burden which owes to its intractability. Indeed, a logical deduction from θ to φ might be highly nontrivial and hard to find, making the application of (2) a constraint of rationality that realistic agents may not be in a position to comply with. Many authors, including [31, 19] have pointed out that imposing that “rational degrees of belief” should be deductively closed is too tall an order, but they did not tackle the problem from the logical point of view. This is what this paper attempts to do, by drawing on the theory of Depth-bounded Boolean logic (DBBL) investigated in [7, 6].

A characteristic feature of DBBL is its informational nature. Whereas connectives in classical logic are defined in terms of truth-values, DBBL provide an informational view of logical consequence which, as a by-product, also provides a tractable approximation of it. Thus DBBL provide the ideal setting to tackle the shortcomings of the probabilistic quantification of uncertainty highlighted by (1) and (2) above.

The central idea behind the (semantic) approach to DBBL is to distinguish two kinds of information. The first is information which the agent possesses explicitly or can trivially infer from it. For definiteness, it is the kind of information that an agent holds when she holds the information that a conjunction is true, i.e. that both conjuncts are true. Put otherwise, the information that both conjuncts are true is immediately available to the agent possessing the information that the conjunction is true. The second kind of information is more subtle and is *obtained* by making a specific kind of hypothetical reasoning. This is closely related to classical reasoning by cases, with the additional requirement that its disjunctive premises always partition the set of (classical) valuations.

To illustrate this kind of reasoning, let us consider a puzzle, originally due to Hector Levesque, and which has been largely discussed in the psychology of reasoning literature, see e.g. Chapter 6 of [38].

Example 1. *Experimental subjects are faced with the following test:*

Jack is looking at Anne, and Anne is looking at George. Jack is married, but George is not. Is a married person looking at an unmarried person?

- A) *Yes*
- B) *No*
- C) *Cannot be determined.*

It is reported by [38] that over 80% of subjects give a wrong answer, the vast majority of which being C).

Here is a plausible explanation for this systematic mistake, which was certainly engineered in the problem. The subject is provided with some explicit information, which in first-order logic notation can be represented as follows:

$$\text{Looks}(\text{Jack}, \text{Anne}) \quad (3)$$

$$\text{Looks}(\text{Anne}, \text{George}) \quad (4)$$

$$\text{Married}(\text{Jack}) \quad (5)$$

$$\neg \text{Married}(\text{George}). \quad (6)$$

Hence there is no information which is provided about the marital status of Ann, which probably leads most respondents to this frequent mistake.

The correct answer is instead A), which can be arrived at by performing a particular kind of hypothetical reasoning based on the fact that

$$\text{either } \text{Married}(\text{Anne}) \text{ or } \neg \text{Married}(\text{Anne}). \quad (\text{Biv})$$

Those mutually exclusive disjuncts clearly exhaust the set of classical models. So we get two cases, both of which decide positively the question of interest, namely

$$\exists x \exists y \varphi(x, y) \quad \text{where } \varphi(x, y) := \text{Looks}(x, y) \wedge \text{Married}(x) \wedge \neg \text{Married}(y).$$

For we have, from (Biv) and the information in (3)-(6) that either

$$\gamma := \text{Married}(\text{Anne}) \wedge \neg \text{Married}(\text{George}) \wedge \text{Looks}(\text{Anne}, \text{George})$$

or

$$\gamma' := \neg \text{Married}(\text{Anne}) \wedge \text{Married}(\text{Jack}) \wedge \text{Looks}(\text{Jack}, \text{Anne}),$$

holds, with $\varphi(\text{Anne}, \text{George})$ following immediately from γ , and $\varphi(\text{Jack}, \text{Anne})$ following immediately from γ' .

This example singles out two senses of “information”. The first is conveyed by (3)-(6) which is readily seen as information actually held by the experimental subjects. In addition, subjects can be seen to attain logical conclusions which exceed the information actually held, by means of hypothetical reasoning. This essentially amounts to applying a single instance of the reasoning by cases on a pair of jointly exhaustive alternatives - as pointed out in (Biv), either Anne is married or she is not.

The hierarchy of Depth-bounded Boolean logic arises by bounding the number k of allowed nested iterations of hypothetical inferential steps of the kind just illustrated. If $k = 0$, then only information actually held by the agent

can be used as the premises of a logical deduction, yielding the 0-depth consequence relation denoted by \vdash_0 . So, an agent confined to a 0-depth consequence relation, would be unable to solve the problem of Example 1, which is readily solved above at depth 1.

For $0 \leq k < m$ it can be shown that $\vdash_k \subset \vdash_m$, whereas $\lim_{k \rightarrow \infty} \vdash_k = \vdash$. This justifies interpreting the hierarchy of DBBLs as an approximation to classical logic, which is indeed attained when the agent is allowed an unbounded use of hypothetical information along the lines illustrated. As we shall point out in due course, keeping this bound in place also results in each element \vdash_k of the sequence being tractable, in the sense of computational complexity.

While various approaches in the literature [15, 27] have similarly focused on feasible approximation of logical systems, what sets apart DBBLs from those is its clear information-based semantics: this provides indeed an uncertainty resolution mechanism which is naturally more adequate than a truth-based semantics, for tackling our motivating questions. Recall that an *event*, understood classically, is a binary random variable. This is captured by the (classical) valuation

$$v(\theta) = \begin{cases} 1 & \text{if the event obtained} \\ 0 & \text{otherwise,} \end{cases}$$

which represents the event's indicator function. In this classical picture, uncertainty is resolved by a logical valuation, which is a total function from the language to the binary set $\{0, 1\}$. In probabilistic modelling, this amounts to the usual assumption concerning the “deterministic nature” of the phenomena whose uncertainty is being quantified. See e.g. [10] for an authoritative point of view along these lines.

By contrast, an information-based uncertainty resolution accounts explicitly for the possibility that the agent is not informed about a particular event. This can be illustrated by decorating the valuation symbol used above with a superscript i , meant to suggest “information”:

$$v^i(\theta) = \begin{cases} 1 & \text{if the agent is informed that } \theta \text{ obtained} \\ 0 & \text{if the agent is informed that } \neg\theta \text{ obtained} \\ * & \text{if the agent is not informed about } \theta . \end{cases}$$

So in Example 1 above, it is natural to represent the fact that the marital status of Ann is unspecified by $v^i(\text{Married}(\text{Anne})) = *$. However, a single binary branching leads to the hypothetical resolution of this uncertainty, with one branch deciding $v^i(\text{Married}(\text{Anne}))$ positively, and the other negatively.

The key difference between v and an evaluation such as v^i lies in the fact that the latter allows for a bounded number k of iterations of hypothetical uncertainty resolution, and, in particular when $k = 0$ may provide only a *partial* evaluation of the sentences/events. This in turn leads to a “non-deterministic” semantics [7] which is instrumental to avoiding the automatic translation of the agent’s ignorance about θ to information about $\neg\theta$.

In analogy to DBBLs, this paper investigates a hierarchy of *Depth-bounded Belief functions* which approximate probability functions and asymptotically coincide with them. Our construction is inspired by the theory of Dempster-Shafer Belief functions [35], as suggested by our choice of terminology. As in DS-theory, none of the belief functions in our hierarchy is constrained by additivity (PL2 above), except for the one attained in the limit. More precisely, given two incompatible events φ and ψ , a depth-bounded belief function might give $B_k(\varphi \vee \psi)$ a value in $[0, 1]$ which is distinct from $B_k(\varphi) + B_k(\psi)$. Going back to our comments on (1) above, the non-additive nature of our belief functions is necessary to ensure a distinct representation of ‘knowledge’ and ‘ignorance’. However, unlike DS-theory (and related proposals) it is achieved primarily by replacing classical logic as the uncertainty resolution mechanism – i.e. by changing ‘the logic of events’ – rather than by just relaxing the axiomatic properties of the uncertainty measure leaving the underlying logic classical, as put forward in the position paper [8].

Finally, a very interesting feature of the framework to be introduced below is in our view this: endowing the agent with higher logical abilities (as captured by the index of the relation \vdash_k) allows her to obtain increasingly tighter approximations of $B_k(\varphi) + B_k(\psi)$ by $B_k(\varphi \vee \psi)$. This puts forward a seemingly novel approach, a logic-based one, to Ellsberg-like problems, as illustrated by Example 4 below.

As a welcome byproduct, major reasoning tasks which are based on any member B_k of our hierarchy will turn out to be computationally tractable, in a sense to be clarified in Section 4.

1.1. Plan of the paper and summary of contributions. Section 2 presents the key features of Depth-bounded boolean logics in a form which suits our present needs. Since this section is instrumental to making this paper self-contained, we refer the interested reader to [7, 6] for the full logical details and [5] for the philosophical underpinnings of the theory of DBBL. Section 3 introduces the belief functions which provide, through our main results, an approximation to probability functions. The resulting approximation will be immune, by logical design, to the problem of turning the absence of information about an event into information about its negation.

Section 4 investigates the complexity of our approximation and identifies the reasoning tasks with depth-bounded belief functions which are feasible (under the usual assumptions which will be briefly recalled). Section 5 collects some final remarks and points to the directions for future work on this topic.

Our main results read as follows. Theorem 1 shows that each probability function can be approximated by a hierarchy of Depth-bounded Belief functions, while, conversely, Theorem 2 singles out the conditions under which our Depth-bounded Belief functions actually determine a probability in the limit. Theorems 3, 4, and 5 show that under rather palatable restrictions, the depth-bounded functions introduced here are an adequate tool to tackle the well-known unfeasibility of logic-based uncertain reasoning [30, 37].

The framework and results presented below extend significantly those of [1], which is therefore superseded by the present paper. In the related [2] we put forward a logic-based framework which served as a direct inspiration for the one introduced below, and whose aim was the approximation of DS Belief functions. So the main result of [2] is that the hierarchy of approximating measures asymptotically coincides with DS Belief functions. Note that elements in the hierarchy of depth-bounded belief functions in [2] only provide a logical model of agents with increasingly higher inferential ability, which nonetheless has no impact on the resolution of uncertainty. This is reflected formally by the fact that, even in the limit, the Depth-bounded belief functions in [2] need not be additive. This work, in contrast, puts the hierarchy of depth-bounded belief functions in correspondence with an increasingly higher ability to resolve uncertainty beyond the information actually held by the agent. This leads in the limit to *additive* measures. Furthermore, [2] provides no analysis of the computational complexity of Depth-bounded belief functions (though it is conjectured in there that they are actually intractable, in the most general case).

2. BOUNDED SEQUENCES AND BOUNDED LOGICS

We begin with some conventions and terminology. We will denote by \mathcal{L} a finite propositional language, with variables p, q etc., connectives \neg, \wedge, \vee , and the constant \perp , standing for any contradiction. The set of sentences generated recursively from \mathcal{L} will be denoted by \mathcal{SL} . Lowercase Greek letters denote sentences, and uppercase Greek letters sets of sentences. Given any sentence $\varphi \wedge \psi, \varphi \vee \psi$ in \mathcal{SL} , the sentences φ, ψ are said to be *the immediate subsentences* of $\varphi \wedge \psi, \varphi \vee \psi$, respectively. Similarly, for any sentence $\neg\varphi$,

its immediate subsentence is φ . A propositional variable has no immediate subsentences. The set of *subsentences* of φ is denoted by $S(\varphi)$ and is the smallest set containing φ and closed under immediate subsentences – ditto for $S(\Gamma)$, where Γ is a set of sentences. Finally $At_{\mathcal{L}}$ denotes the set of *atoms*, i.e. the maximal (classically) consistent conjunctions of literals from \mathcal{L} .

Trees, and sequences thereof, will represent the information intervening in depth-bounded reasoning. In our framework the root of a tree is labelled with the information initially held by the agent. If no such information is available to the agent, the root will be labelled by $*$. Otherwise it will be labelled by a set of sentences from \mathcal{SL} . In what follows, the leaves of a tree T are denoted by $Le(T)$, and the *depth of a node* α in a tree is the length of the path from the root of the tree to α .

Branching is where the core construction of the other kind of information, the hypothetical one, leading to uncertainty resolution, takes place, along the line of our first definition:


Definition 1. *Let $\gamma \in \mathcal{SL} \cup \{*\}$. A depth-bounded tree sequence $(T_k)_{k \in \mathbb{N}}$ (DBT-sequence for short) rooted in γ is defined recursively as follows:*

- T_0 is a tree having as only node its root γ .
- for $k > 0$, T_k is obtained by branching at least one node α of depth $k - 1$ in T_{k-1} with two nodes $\alpha \wedge \beta$ and $\alpha \wedge \neg\beta$, for some $\beta \in \mathcal{SL}$.

A finite depth-bounded tree sequence is any initial segment $(T_k)_{k < n}$ of a depth-bounded sequence $(T_k)_{k \in \mathbb{N}}$, for some $n \in \mathbb{N}$.

A few consideration on the definition above are in place. First, the formula $\beta \in \mathcal{SL}$, used for the branching in the Definition above, is not fixed, but different formulas can be used for each different parent node and each iteration of the construction of the T_k s. Second, in case γ is the empty information $*$, we will take T_1 to be a tree rooted in $*$, with two branchings β and $\neg\beta$, for some sentence $\beta \in \mathcal{SL}$. Finally, note that at least one formula $\alpha \in Le(T_k)$ has to be of depth k , but this need not be the case for all the formulas in $Le(T_k)$. Consider indeed the following.

Example 2.



Both the tree above on the left, that we will denote by T_2 , and that on the right, that we will denote by S_2 , are trees of depth 2, obtained by branching

a tree (say $T_1 = S_1$) composed of $*$ and children nodes $p, \neg p$. However, T_2 is obtained by branching each formula in $Le(T_1)$, so that each node in $Le(T_2)$ has depth 2. On the other hand, S_2 is obtained by expanding only the node p , of depth 1, in $Le(S_1)$. The node $\neg p \in Le(S_1)$ is not expanded, so that we will also have $\neg p \in Le(S_2)$, while $\neg p$ is still of depth 1. Note that it will not suffice to expand the latter node to obtain a tree of depth 3, but rather at least one of the nodes of depth 2 in $Le(S_2)$ has to be expanded.

The characteristic feature of Definition 1 is that trees form an ordered *sequence*, which can be interpreted as a hierarchy of levels of hypothetical uncertainty resolution branchings, of the kind illustrated by Example 1. Each successive level is obtained through a new binary branching aimed at deciding a yet undecided sentence and its negation. In this sense each new level of hypothetical uncertainty resolution is achieved by means of a constrained version of reasoning by cases. Put in yet another way, step $k + 1$ in the recursive construction of $(T_k)_{k \in \mathbb{N}}$ amounts to the binary expansion of some k -depth leaves by new children nodes, each being labeled by a pair of mutually inconsistent pieces of information. Hence depth-bounded sequences can be seen as possible explorations of how uncertainty may be resolved by information which is *not* held by the agent, starting from the information which the agent does hold. Clearly, agents whose reasoning is constrained by distinct bounds may not resolve the relevant uncertainty in the same way. This provides the logical basis to add in Section 3 an important and as yet under appreciated dimension to the the Laplacian dictum to the effect that probability depends in part on “knowledge” and in part on “ignorance”.

Going back to the construction in Definition 1, note that if a leaf contains an atom as a subsentence and is not contradictory, then *all* uncertainty in the sentences of interest has been resolved, i.e. the information reached coincides with a classical truth-value. However, and this is our main interest here, the agent will not normally possess sufficient resources to get there, and hence will be forced to content herself with an approximation of such truth-value. As we prove in the next section, the greater the depth of the tree-sequence, the better the agent will approximate a probabilistic estimate of such an uncertainty.

We must now integrate our framework with a sensible way of applying the recursive step in Definition 1, i.e. for selecting only limited, *useful* amount of branchings. A criterion of “usefulness” emerges rather naturally: given a sentence φ in classical logic that an agent might want to decide, the agent checks which leaves α classically derive φ , i.e. $\alpha \vdash \varphi$. However, this is clearly not viable for resource-bounded agents. First, owing to the intractability of

classical logic. Second, as pointed out in [6] and informally in the introduction to this paper, a classical derivation $\alpha \vdash \varphi$ might itself make “hidden” use of hypothetical information, beyond that encoded in α . This would be undesirable given our goal of accounting for the distinct roles played by actual and hypothetical information in the construction of the agent’s uncertainty quantification. Hence we resort to 0-depth inference, introduced in [7]. This logic, as mentioned in the introduction, does allow any use of hypothetical information, and, as a welcome byproduct, it is computationally tractable.

So, in preparation for Proposition 1, which collects the key properties of depth-bounded sequences, we briefly recall now the basic features of 0-depth inference from a proof-theoretic point of view. The presentation is entirely based on the IntElim (Introduction and Elimination) rules in Table 1. We note in passing that this logic, and indeed the whole hierarchy of depth-bounded boolean logics which stems from it, is sound and complete with respect to a non-deterministic, information-based semantics. See [6, 7] for more details.

$\frac{\varphi \quad \psi}{\varphi \wedge \psi} (\wedge \mathcal{I})$	$\frac{\neg \varphi}{\neg(\varphi \wedge \psi)} (\neg \wedge \mathcal{I}1)$	$\frac{\neg \psi}{\neg(\varphi \wedge \psi)} (\neg \wedge \mathcal{I}2)$
$\frac{\neg \varphi \quad \neg \psi}{\neg(\varphi \vee \psi)} (\neg \vee \mathcal{I})$	$\frac{\varphi}{\varphi \vee \psi} (\vee \mathcal{I}1)$	$\frac{\psi}{\varphi \vee \psi} (\vee \mathcal{I}2)$
$\frac{\varphi \quad \neg \varphi}{\perp} (\perp \mathcal{I})$	$\frac{\varphi}{\neg \neg \varphi} (\neg \neg \mathcal{I})$	
$\frac{\varphi \vee \psi \quad \neg \varphi}{\psi} (\vee \mathcal{E}1)$	$\frac{\varphi \vee \psi \quad \neg \psi}{\varphi} (\vee \mathcal{E}2)$	$\frac{\neg(\varphi \vee \psi)}{\neg \varphi} (\neg \vee \mathcal{E}1)$
$\frac{\neg(\varphi \vee \psi)}{\neg \psi} (\neg \vee \mathcal{E}2)$	$\frac{\varphi \wedge \psi}{\varphi} (\wedge \mathcal{E}1)$	$\frac{\varphi \wedge \psi}{\psi} (\wedge \mathcal{E}2)$
$\frac{\neg(\varphi \wedge \psi) \quad \varphi}{\neg \psi} (\neg \wedge \mathcal{E}1)$	$\frac{\neg(\varphi \wedge \psi) \quad \psi}{\neg \varphi} (\neg \wedge \mathcal{E}2)$	
$\frac{\neg \neg \varphi}{\varphi} (\neg \neg \mathcal{E})$	$\frac{\perp}{\varphi} (\perp \mathcal{E})$	

TABLE 1. The standard propositional Introduction and Elimination rules for \neg, \wedge, \vee .

Definition 2. For $\Gamma \cup \{\varphi\} \subseteq \mathcal{SL}$, let $\Gamma \vdash_0 \varphi$ if and only if there is a sequence of sentences $\varphi_1, \dots, \varphi_m$ such that $\varphi_m = \varphi$ and each sentence φ_i is either in Γ or obtained by an application of the rules in Table 1 on the sentences φ_j with $j < i$.

So whilst $p \wedge q \vdash_0 p$, it is not the case that $* \vdash_0 p \vee \neg p$. More generally, it can be shown that, contrary to classical logic, \vdash_0 has no tautologies at all, i.e. for no φ , do we have $* \vdash_0 \varphi$.

The remainder of this section illustrates the pivotal role of 0-depth derivation in the construction of bounded sequences, starting with the definition of *maximal sequences* which pins down the most useful form of uncertainty resolution.

As usual, we say that γ decides φ if either $\gamma \vdash_0 \varphi$ or $\gamma \vdash_0 \neg\varphi$. It is convenient to pin down the elements of a set Γ which decide the sentence φ , i.e. we write

$$\text{dec}_\Gamma(\varphi) = \{\alpha \in \Gamma \mid \alpha \text{ decides } \varphi\}.$$

This allows us to define maximality and closure in trees and sequences, as follows.

Definition 3. Let $\{T_k\}_{k \in \mathbb{N}}$ be a DBT-sequence and $\varphi \in \mathcal{SL}, \Pi \subseteq \mathcal{SL}$.

- A tree T_k is said to be $\{\varphi\}$ -maximal if there exists no tree T'_k of depth k with the same root, such that $|\text{dec}_{Le(T'_k)}(\varphi)| > |\text{dec}_{Le(T_k)}(\varphi)|$.
- A tree T_k is said to be $\{\varphi\}$ -closed if $\text{dec}_{Le(T_k)}(\varphi) = Le(T_k)$, i.e. all the formulas $\alpha \in Le(T_k)$ decide φ .
- A sequence $\{T_k\}_{k \in \mathbb{N}}$ is Π -closed if, for each $\varphi \in \Pi$ there is a $k \in \mathbb{N}$ such that T_k is $\{\varphi\}$ -closed.
- A finite sequence $\{T_k\}_{k \leq n}$ is Π -maximal(closed) if for each $\varphi \in \Pi$, there is a $j \leq n$ such that T_j is $\{\varphi\}$ -maximal(closed).

Clearly, closed trees are in particular maximal. While a closed tree represents the ideal case in which an agent has resolved all the uncertainty relatively to a formula of interest, we can think of maximality as a “second-best”. For DBT sequences, it can be shown that the two concepts actually coincide in the infinite case (this is the reason why we do not include a notion of maximality for them), while they differ for the finite case. Correspondingly, our approximation results, i.e. Theorem 2 and Theorem 1 make use of assumptions on closed sequences, while those in Section 4, where finite sequences are concerned, are limited to maximal ones.

Example 3. Consider the formula $(p \vee \neg p) \vee q$, and the following depth 1 trees:



The tree on the left-hand side is $\{(p \vee \neg p) \vee q\}$ -closed, for we have both $p \vdash_0 (p \vee \neg p) \vee q$ and $\neg p \vdash_0 (p \vee \neg p) \vee q$. However, it is not the case that $\neg q$ on the tree on the right-hand-side decides $(p \vee \neg p) \vee q$. Hence the tree on the right-hand-side is not $\{(p \vee \neg p) \vee q\}$ -closed, and in particular not maximal.

Note that closeness (let alone maximality) does not imply uniqueness. This is easily shown by taking again the 1 depth trees above, which are both $\{(p \vee \neg p) \vee (q \vee \neg q)\}$ -closed.

So far, we have only used 0-depth consequence, which is only the starting point of the hierarchy of Depth-Bounded Boolean Logics. The remaining logics in the hierarchy were defined in [7] as follows.

Definition 4. For each $k > 0$ and set of sentences $\Gamma \cup \{\varphi\} \subseteq \mathcal{SL}$, we let $\Gamma \vdash_k \varphi$ if and only if there is a $\beta \in S(\Gamma \cup \{\varphi\})$ such that $\Gamma, \beta \vdash_{k-1} \varphi$ and $\Gamma, \neg\beta \vdash_{k-1} \varphi$.

We refer to the original works [7, 6] for further details on the key properties of Depth-Bounded Boolean Logics which include the following:

- $\vdash_0 \subset \vdash_1 \subset \dots \subset \vdash_k \subset \dots$, so the depth-bounded consequence relations form a hierarchy;
- the limit for $k \rightarrow \infty$ of \vdash_k is \vdash , i.e. in the limit, the hierarchy of DBBLs coincides with classical logic;
- for each k , \vdash_k has a polynomial-time decision procedure.

The next Proposition shows that k -depth logics and maximal trees are closely connected. Before that, let us recall that, as some classical inferences are available only at higher level of the hierarchy of DBBLs, the same holds for recognizing contradictions. In principle, it might be the case that a piece of information α is contradictory ($\alpha \vdash \perp$), while this might not be recognizable at 0-depth ($\alpha \not\vdash_0 \perp$). Ruling out this case leads us to a even tighter connection between k -depth trees and the corresponding k -depth logic. In the following we will say that a tree T_k is *free of deep contradictions* iff for each $\alpha \in Le(T_k)$, having $\alpha \vdash \perp$ implies that $\alpha \vdash_0 \perp$.

Proposition 1. Let $\varphi \in \mathcal{SL} \cup \{*\}$, $\psi \in \mathcal{SL}$. If $\varphi \vdash_k \psi$, then the following hold:

- (a) There exists a tree T_k of depth k rooted in φ , such that $\alpha \vdash_0 \psi$, for each $\alpha \in Le(T_k)$.

- (b) *There exists a tree T_k of depth k rooted in any sentence $\gamma \in \mathcal{SL} \cup \{*\}$, such that, for each $\alpha \in Le(T_k)$, if $\alpha \vdash_0 \varphi$, then $\alpha \vdash_0 \psi$.*
- (c) *For any $\{\psi\}$ -maximal tree T_k of depth k rooted in φ and free of deep contradictions, we have that $\alpha \vdash_0 \psi$, for each $\alpha \in Le(T_k)$.*
- (d) *For any $\{\neg\varphi \vee \psi\}$ -maximal tree T_{k+1} of depth $k+1$ rooted in a sentence γ and free of deep contradictions, we have that $\alpha \vdash_0 \psi$, for each $\alpha \in Le(T_{k+1})$ such that $\alpha \vdash_0 \varphi$.*

Proof. (a). Let us construct a tree T_k of depth k as follows. The root is φ . Now, since $\varphi \vdash_k \psi$, by Definition 4, there is a β such that $\varphi, \beta \vdash_{k-1} \psi$ and $\varphi, \neg\beta \vdash_{k-1} \psi$. We thus branch the root with two nodes $\varphi \wedge \beta$, $\varphi \wedge \neg\beta$. We repeat the procedure as follows: for any $1 \leq j < k$, assuming that, for each node α at depth j , we have $\alpha \vdash_{k-j} \psi$, we branch the node α with the two formulas β and $\neg\beta$ such that $\alpha, \beta \vdash_{k-j-1}$ and $\alpha, \neg\beta \vdash_{k-j-1} \psi$ (such a β exists, by Definition 4).

Proceeding by induction, we then obtain that any node α at depth $0 \leq j \leq k$ in the tree thus constructed is such that $\alpha \vdash_{k-j} \psi$; hence, in particular, for any leaf $\alpha \in Le(T_k)$ (i.e. formula at depth k), $\alpha \vdash_0 \psi$.

(b) We consider the same tree as in (a), but with root γ , instead of φ . Let us write any leaf in this tree as $\gamma \wedge \beta$ and any leaf in the tree in (a) as $\varphi \wedge \beta$. By (a), we have that $\varphi \wedge \beta \vdash_0 \psi$ for any β . Whenever $\gamma \wedge \beta \vdash_0 \varphi$, since by $(\wedge\mathcal{E})$, $\gamma \wedge \beta \vdash_0 \beta$, we also have $\gamma \wedge \beta \vdash_0 \varphi \wedge \beta$ (by $\wedge\mathcal{I}$). Hence, by the transitivity of \vdash_0 , we obtain our claim, namely that $\gamma \wedge \beta \vdash_0 \psi$, whenever $\gamma \wedge \beta \vdash_0 \varphi$.

(c) The tree built in (a) is such that $\alpha \vdash_0 \psi$ for each $\alpha \in Le(T_k)$, hence it is a $\{\psi\}$ -closed tree. Any other $\{\psi\}$ -maximal tree T'_k will thus also be in particular $\{\psi\}$ -closed. Consider now a $\{\psi\}$ -closed tree T'_k which is free of deep-contradictions: each leaf will decide ψ . Assume that, for some leaf, say $\varphi \wedge \beta \in Le(T'_k)$, it holds that $\varphi \wedge \beta \vdash_0 \neg\psi$. This means in particular that $\varphi \wedge \beta \vdash_k \neg\psi$. But then, since $\varphi \vdash_k \psi$, we also have $\varphi \wedge \beta \vdash_k \psi$. Hence, $\varphi \wedge \beta \vdash_k \perp$, and by the assumption that T'_k is free of deep contradictions (see Definition 3), we will have $\varphi \wedge \beta \vdash_0 \perp$. Hence, by rule $(\perp\mathcal{E})$, we have again $\varphi \wedge \beta \vdash_0 \psi$. This finally shows that each leaf in a $\{\psi\}$ -maximal tree derives at 0-depth ψ .

(d) Assume $\varphi \vdash_k \psi$. Note that $\neg\varphi \vdash_0 \neg\varphi \vee \psi$ by the rule $(\vee\mathcal{I}1)$, hence in particular $\neg\varphi \vdash_k \neg\varphi \vee \psi$. On the other hand, since $\varphi \vdash_k \psi$, again by the rule $(\vee\mathcal{I}1)$, we get $\varphi \vdash_k \neg\varphi \vee \psi$. This shows, by Definition 4 that $\vdash_{k+1} \neg\varphi \vee \psi$. Applying (c), we obtain that for any $\{\neg\varphi \vee \psi\}$ -maximal tree T_{k+1} rooted in γ , we have $\alpha \vdash_0 \neg\varphi \vee \psi$ for each $\alpha \in Le(T_{k+1})$. Now, let $\alpha \in Le(T_{k+1})$

be such that $\alpha \vdash_0 \varphi$. By the rule $\neg\neg\mathcal{I}$ applied to φ , we obtain $\alpha \vdash_0 \neg\neg\varphi$. Recalling that $\alpha \vdash_0 \neg\varphi \vee \psi$, applying the rule $\vee\mathcal{E}1$ to $\neg\varphi \vee \psi$ and $\neg\neg\varphi$, we finally obtain $\alpha \vdash_0 \psi$.

□

3. APPROXIMATING PROBABILITY FUNCTIONS

We now turn to our main concern, namely showing that finitely additive probability functions $P: \mathcal{SL} \rightarrow [0, 1]$ can be approximated by sequences of uncertainty measures which we will call *k-depth belief functions*.

Depth-bounded trees capture the hypothetical uncertainty resolution of *one* particular piece of information. However the probabilistic quantification of uncertainty typically involve the simultaneous hypothetical exploration of more pieces of information.

Suppose, for instance, that an agent is informed that half of the balls in a urn are either yellow or green ($Y \vee G$), and the remaining half are either red or white ($R \vee W$). Here, if an agent is willing to assign probabilities to the propositional variables Y, G, R, W , she needs to simulate information that she does not possess, concerning the piece of information $Y \vee G$ on the one hand, and $R \vee W$, on the other. For instance, on the basis of the symmetry of the problem, she might assign equal weights to R and W , and to Y and G . This would result in a probabilistic assignment P such that $P(Y) = P(G) = P(R) = P(W) = 1/4$. Once more, let us stress that the assignment $P(Y) = P(G) = P(R) = P(W) = 1/4$ is different in nature from the information $P(Y \vee G) = P(R \vee W) = 1/2$: the former was arrived at by the agent, by taking into account different hypothetical scenarios and weighting them, while the latter was information actually provided to her from the start. This setup will be formally addressed, using our framework, in Example 5.

The information actually held by an agent, as that on $Y \vee G$ and $R \vee W$, in the above example, will be collected in the *support set*, which we denote by *Supp*. By this we do not mean that an agent possesses information on whether the sentences in *Supp* are the case, but rather that only on these sentences the agent can initially quantify her uncertainty probabilistically. Each formula of *Supp* will act then as the root of a DBT-sequence, aimed at resolving the uncertainty therein contained.

With this motivation we now generalise the framework of Section 2 from sequences of depth-bounded trees to *depth-bounded forests*. To do this, it will be convenient to assume that *Supp* is nonempty and contains at least

one consistent sentence. As before, the symbol $*$ will be used to denote the special case in which the agent holds no initial information at all, i.e. we let in that case $Supp = \{*\}$. In the following $inc(\Gamma)$ will denote the set $\{\alpha \in S \mid \alpha \vdash_0 \perp\}$ for any $\Gamma \subseteq \mathcal{SL}$. Unless otherwise stated, we shall thus have $inc(Supp) \subset Supp \subseteq \mathcal{SL} \cup \{*\}$.

Definition 5. For each $\gamma \in Supp$, let $\{T_k(\gamma)\}_{k \in \mathbb{N}}$ be a DBT-sequence rooted in γ . Then:

- A depth-bounded forest sequence (DBF-sequence, for short) $(F_k)_{k \in \mathbb{N}}$ based on $Supp$ is obtained by letting $F_k = \cup_{\gamma \in Supp} T_k(\gamma)$, for each $k \in \mathbb{N}$;
- A finite depth-bounded forest sequence is any initial segment $(F_k)_{k \leq n}$ of a depth-bounded sequence $(F_k)_{k \in \mathbb{N}}$, for some $n \in \mathbb{N}$.

The notion of maximality and closure (see Definition 3 above) adapts immediately to forests as follows.

Definition 6. For $\varphi \in \mathcal{SL}, \Pi \subseteq \mathcal{SL}$ we say that

- a forest F_k based on $Supp$ is $\{\varphi\}$ -closed (maximal) if each tree $T_k(\gamma)$ in F_k is $\{\varphi\}$ -closed(maximal) for each $\gamma \in Supp$.
- a DBF-sequence $\{F_k\}_{k \in \mathbb{N}}$ based on $Supp$ is Π -closed if each DBT-sequence $\{T_k(\gamma)\}_{k \in \mathbb{N}}$ in $\{F_k\}_{k \in \mathbb{N}}$ is Π -closed, for each $\gamma \in Supp$.
- a finite DBF-sequence $\{F_k\}_{k \leq n}$ based on $Supp$ is Π -closed (maximal) if each finite DBT-sequence $\{T_k(\gamma)\}_{k \leq n}$ is Π -closed (maximal), for each $\gamma \in Supp$.

We can now start working our way up to the approximation of probability functions.

Definition 7 (Probability mass function). For $\Gamma \subseteq \mathcal{SL} \cup \{*\}$, say that $m: \mathcal{P}(\Gamma) \rightarrow [0, 1]$ is a probability mass function over Γ if m is a finitely additive measure, such that $m(\Gamma) = 1$, $m(\emptyset) = 0$, and $m(inc(\Gamma)) = 0$.

Note that in the interest of readability we write $m(\alpha)$ in place of $m(\{\alpha\})$, unless confusion may arise. We shall also refer to probability mass functions simply as “mass functions”. When applied to the set $Supp$, the mass function defined above captures, in analogy with Dempster-Shafer Belief functions, the uncontroversial probabilistic quantification of uncertainty which is licensed by the actual information possessed by the agent. While in Shafer’s theory [36] agents can fill in potential gaps in their “evidence” by using judgment, the key idea of our proposal is to provide a logical framework to constrain this judgment formation. In particular, depth-bounded forests

provide a principled way to allow agents to go beyond the information they actually possess, by exploring hypotheses for the resolution of uncertainty. The characteristic feature of our construction is that this simulation will be bounded by a *fixed* sequence.

Definition 8 (DBM-sequence). *Let $\{F_k\}_{k \in \mathbb{N}}$ be a DBF-sequence based on $Supp$, which is free of deep contradictions. We say that $\mathcal{F} = (F_k, m_k)_{k \in \mathbb{N}}$ is a depth-bounded mass sequence (DBM-sequence for short) if each m_k is a mass function over $Le(F_k)$, and for each $k > 0$, and leaf $\alpha \in Le(F_{k-1})$, the following holds:*

- $m_k(\alpha) = m_{k-1}(\alpha)$, if $\alpha \in Le(F_k)$.
- $m_k(\{\alpha \wedge \beta, \alpha \wedge \neg \beta\}) = m_{k-1}(\alpha)$, if α has children nodes $\alpha \wedge \beta, \alpha \wedge \neg \beta$.

The condition on deep contradictions excludes the possibility that sentences might result consistent at depth $k-1$, while turning out inconsistent at depth k (see [2, 7]). If this were the case, the inconsistent sentences would have to be assigned mass 0, and their original mass should be redistributed among the consistent formulas. We do not consider that more general case, since it would complicate our definitions, without any major conceptual gains for our main results.

Note that the initial mass function m_0 is defined over $Le(F_0)$, which is equal to the initial support set $Supp$. The agent further refines the quantification by distributing the mass assigned to each leaf at depth k , among its children nodes, which are the leaves at depth m_{k+1} . The leaves of depth-bounded forests for $k > 0$ keep track of the information that has been probabilistically quantified by the agent.

To extend this to an arbitrary sentence φ , the agent looks at the set of sentences whose uncertainty has been assessed (the leaves) and that allows her to 0-derive φ . Dually, the agent can also consider the set of sentences whose uncertainty has been assessed (the leaves) and that do not allow a 0-derivation of $\neg \varphi$. This is again reminiscent of the construction of Dempster-Shafer Belief functions.

Indeed, for any sentence $\varphi \in \mathcal{SL}$, given a DBF-sequence $\{F_k\}_{k \in \mathbb{N}}$ of depth-bounded forests, the set

$$b_k(\varphi) = \{\alpha \in Le(F_k) \mid \alpha \vdash_0 \varphi, \alpha \not\vdash_0 \perp\}$$

captures the belief mass that an agent commits to a sentence φ *after* simulating uncertainty resolution of depth k whereas

$$pl_k(\varphi) = \{\alpha \in Le(F_k) \mid \alpha \not\vdash_0 \neg \varphi\}$$

similarly captures k -depth plausibility. This is the sense in which our construction puts the judgment component of Belief functions under logical constraints.

Definition 9. Let $(F_k, m_k)_{k \in \mathbb{N}}$ be a DBM-sequence. The k -depth Belief function B_k and the k -depth Plausibility function Pl_k are defined by letting:

$$B_k(\varphi) = m_k(b_k(\varphi)) \text{ and } Pl_k(\varphi) = m_k(pl_k(\varphi)),$$

respectively.

Let us now put our framework to work through a classic example. Though anticipated by [24, 25] the example has gained prominence in the analysis of [13], who used it to argue normatively against the adequacy of the probabilistic representation of uncertainty in the face of partly resolving uncertainty, or ambiguity. In addition to illustrating the basic mechanism of our framework, the example also sheds new light on how the logical framework may contribute to the analysis of foundational issues in decision theory.

Example 4 (Ellsberg’s “one urn” problem). Let \mathcal{L} be a finite language with the propositional variables $\{Y, G, R\}$, interpreted as the ball is Yellow, Green, Red, respectively. The problem is that of a single draw from an urn and the “Laplacian information” is to the effect that $2/3$ of the balls are either yellow or green with the remaining $1/3$ being red. This information is logically represented by the conjunction γ of the sentences in the set

$$\{Y \rightarrow (\neg G \wedge \neg R), R \rightarrow (\neg G \wedge \neg Y), G \rightarrow (\neg R \wedge \neg Y), Y \vee G \vee R\}.$$

The problem clearly suggests letting $Supp = \{(Y \vee G) \wedge \gamma, R \wedge \gamma\}$.

The first step in our model now requires us to quantify our uncertainty based only on actual information, which is readily done by letting $m_0((Y \vee G) \wedge \gamma) = 2/3$ and $m_0(R \wedge \gamma) = 1/3$. This implies

$$\begin{aligned} B_0(Y \vee G) &= 2/3 \\ B_0(R) &= 1/3 \\ B_0(\gamma) &= 1 \\ B_0(Y) = B_0(G) &= 0. \end{aligned}$$

Since the interesting part of this problem is the process of analyzing the initial information and assigning probabilities to its sentences, it is natural to restrict the attention to sequences which are $S(Supp)$ -maximal. One can easily check that, for the node $R \wedge \gamma$ in $Supp$, we already have $R \wedge \gamma \vdash_0 \neg Y$, $R \wedge \gamma \vdash_0 \neg G$ and $R \wedge \gamma \vdash_0 R$. The node thus decides $S(Supp)$, and needs not be expanded at depth 1. However, at depth 1 the node $(Y \vee R) \wedge \gamma$ will be

suitably expanded. Note that there are only two possible $S(\text{Supp})$ -maximal forests at this depth:

$$\begin{array}{ccc} (Y \vee G) \wedge \gamma & & R \wedge \gamma \\ \wedge & & \\ Y & & \neg Y \end{array}$$

and

$$\begin{array}{ccc} (Y \vee G) \wedge \gamma & & R \wedge \gamma \\ \wedge & & \\ G & & \neg G \end{array}$$

In both cases all the leaves of the forest at depth 1 decide $S(\text{Supp})$, hence the forest is $S(\text{Supp})$ -closed, and in particular maximal. At this depth, the symmetry of the available, mutually exclusive, options may lead the agent to assign $m_1((Y \vee G) \wedge \gamma \wedge Y) = m_1((Y \vee G) \wedge \gamma \wedge \neg Y) = 1/3$, though this is not the unique available option). This in turn yields

$$\begin{aligned} B_1(Y) &= 1/3 \\ B_1(G) &= 1/3 \\ B_1(R) &= 1/3. \end{aligned}$$

Note that this is now compatible with the probabilistic quantification of uncertainty, which can be seen to have been obtained from information actually held by the agent, via a single iteration of hypothetical uncertainty resolution.

Example 5. The variant of Ellsberg's problem sketched at the beginning of this Section arises by considering the propositional variables $\{Y, G, R, W\}$, a sentence γ representing background information, and taking $\text{Supp} = \{(Y \vee G) \wedge \gamma, (W \vee R) \wedge \gamma\}$.

The quantification of uncertainty based on the information actually possessed by the agent is rendered by letting

$$m_0((Y \vee G) \wedge \gamma) = 1/2$$

and

$$m_0((W \vee R) \wedge \gamma) = 1/2.$$

This implies

$$\begin{aligned} B_0(Y \vee G) &= 1/2 \\ B_0(W \vee R) &= 1/2 \\ B_0(\gamma) &= 1 \\ B_0(Y) = B_0(G) = B_0(W) = B_0(R) &= 0. \end{aligned}$$

At depth 1 the nodes $(Y \vee G) \wedge \gamma$ and $(W \vee R) \wedge \gamma$ will be suitably expanded, for instance as follows:

$$\begin{array}{ccc} (Y \vee G) \wedge \gamma & & (W \vee R) \wedge \gamma \\ \wedge & & \wedge \\ Y \quad \neg Y & & W \quad \neg W \end{array}$$

At this point, at depth 1 all the leaves of the forest decide $S(\text{Supp})$. Again, by the symmetry of the problem, we might consider $m_1((Y \vee G) \wedge \gamma \wedge Y) = m_1((Y \vee G) \wedge \gamma \wedge \neg Y) = 1/4$ and $m_1((W \vee R) \wedge \gamma \wedge W) = m_1((W \vee R) \wedge \gamma \wedge \neg W) = 1/4$ which in turn yields:

$$\begin{aligned} B_1(Y) &= 1/4 \\ B_1(G) &= 1/4 \\ B_1(R) &= 1/4 \\ B_1(W) &= 1/4. \end{aligned}$$

According to the main result of this Section, and indeed of the paper, k -depth Belief functions are shown to approximate their additive counterparts. In preparation to it we will illustrate the close relation between k -depth Belief functions and k -depth deductions. We begin by adapting to the present framework Proposition 1 of [2]. The main difference between the two formalisms is that the k -depth belief functions here are based on a *fixed* choice of hypothetical information, encoded in the underlying DBF-sequence, while those in [2] may take into account any such choices at once. This owes to the rather different research questions which motivate the two papers. However, the two formalisms are reconciled by adding suitable maximality assumptions to the DBF-sequence defining our functions.

Proposition 2. *Let $\mathcal{F} = (F_k, m_k)_{k \in \mathbb{N}}$ be a DBM-sequence based on $\text{Supp} \subseteq \mathcal{SL} \cup \{*\}$. The following hold:*

- (1) $\vdash_k \varphi$ implies $B_k(\varphi) = 1$, for any F_k which is $\{\varphi\}$ -maximal and free of deep contradictions.
- (2) $\vdash_k \neg\varphi$ implies $B_k(\varphi) = 0$, for any F_k that is $\{\varphi\}$ -maximal and free of deep contradictions.
- (3) $\varphi \vdash_k \psi$ implies
 - $B_k(\varphi) \leq B_k(\psi)$, for some F_k .
 - $B_{k+1}(\varphi) \leq B_{k+1}(\psi)$, for any F_k which is $\{\neg\varphi \vee \psi\}$ -maximal and free of deep contradictions.
- (4) $B_k(\bigvee_{i=1}^n \varphi_i) \geq \sum_{\emptyset \neq S \subseteq \{1, \dots, n\}} (-1)^{|S|-1} B_k(\bigwedge_{i \in S} \varphi_i)$ for each $\varphi_1, \dots, \varphi_n \in \mathcal{SL}$.

Proof. (1) and (2) are immediate consequences of Proposition 1 (b). The first item of (3) is an immediate consequence of Proposition 1(c), and the second of Proposition 1(d)

(4). By a straightforward adaptation of Proposition 1 in [2]. \square

Remark 1. *Proposition 2 makes explicit that suitable maximality assumptions play a fundamental role in establishing the close connection between k -depth Belief functions of Definition 9 and depth-bounded logics. The assumption requires that the agent chooses the “right” pieces of hypothetical information: i.e. those which decide the sentence of interest in as many cases (i.e. leaves of the forest) as possible. This condition makes our results normatively interpretable: once the limitation of the inferential ability of the agent are taken into account (the depth), we are still requiring that her belief follows constraints dictated by the corresponding (k -depth) logic.*

We will now show that, under suitable assumptions, DBM-sequences determine a probability, in the limit. We begin with a lemma adapted from [3], where it was originally presented in a qualitative setting.

Lemma 1. *Let $\mathcal{F} = (F_k, m_k)_{k \in \mathbb{N}}$ be a DBM-sequence which is free of deep contradictions. For any formula φ , if there is a $j \in \mathbb{N}$ such that F_j is $\{\varphi\}$ -closed, then $B_k(\varphi) = B_j(\varphi)$ for any $k \geq j$.*

Proof. Let F_j be $\{\varphi\}$ -closed, i.e. assume that each $\alpha \in Le(F_j)$ decides φ . Let us focus the set of leaves α such that $\alpha \vdash_0 \varphi$, i.e. the set $b_j(\alpha)$. Now, for any $k \geq j$, the leaves which are descendants of α will be of the form $\alpha \wedge \beta$. Hence, by applying rule $\wedge\mathcal{E}$ in Table 1, and using the transitivity of \vdash_0 we will have $\alpha \wedge \beta \vdash_0 \varphi$, for each such $\alpha \wedge \beta$. A similar argument goes through for those remaining $\alpha \in Le(F_j)$ such that $\alpha \vdash \neg\varphi$. We get thus that $b_k(\varphi) = \{\alpha \wedge \beta \in Le(F_k) \mid \alpha \vdash_0 \varphi\}$. Moreover, since \mathcal{F} is free of deep contradictions, we have that $m_j(\alpha) = m_k\{\alpha \wedge \beta \mid \alpha \wedge \beta \text{ descendant of } \alpha \in Le(F_j)\}$ (see Definition 8) for each $\alpha \in b_j(\varphi)$. Hence, we obtain $m_k(b_k(\varphi)) = m_j(b_j(\varphi))$, and thus $B_k(\varphi) = B_j(\varphi)$ for each $k \geq j$. \square

We can now state and prove our main results. First, we show that (additive) probability functions can be approximated by k -depth Belief functions. Conversely we identify the conditions under which k -depth Belief functions give rise to probability functions in the limit.

Theorem 1. *Let $P: \mathcal{SL} \rightarrow [0, 1]$ be a probability function. Then there is a DBM-sequence \mathcal{F} based on $\text{Supp} = \{*\}$ such that*

$$P(\varphi) = \lim_{k \rightarrow \infty} B_k(\varphi),$$

for any $\varphi \in \mathcal{SL}$.

Something indeed quite stronger will be proved, namely that there is an $n \in \mathbb{N}$ such that $P(\varphi) = B_n(\varphi)$. Recall that we denote by $At_{\mathcal{L}}$ the set of maximally (classical) consistent conjunctions of literals from a language \mathcal{L} with n propositional variables.

Proof. Define a DBM-sequence $\mathcal{F} = (F_k, m_k)_{k \in \mathbb{N}}$ based on $*$ such that $Le(F_n)$ coincides with $At_{\mathcal{L}}$ and $m_n(\alpha) = P(\alpha)$ for each $\alpha \in At_{\mathcal{L}} = Le(F_n)$. Note that once we fix m_n , Definition 8 forces us to uniquely determine all the m_k for $k < n$. Moreover, it can be easily shown that, if $\alpha \in At_{\mathcal{L}}$, then for any $\varphi \in \mathcal{SL}$, we have that $\alpha \vdash \varphi$ iff $\alpha \vdash_0 \varphi$. Hence for any $\varphi \in \mathcal{SL}$

$$P(\varphi) = \sum_{\substack{\alpha \in At_{\mathcal{L}} \\ \alpha \vdash \varphi}} P(\alpha) = \sum_{\substack{\alpha \in At_{\mathcal{L}} \\ \alpha \vdash_0 \varphi}} P(\alpha) = m_n(b_n(\varphi)) = B_n(\varphi).$$

Moreover, at depth n , the uncertainty about all the propositional variables in \mathcal{L} will have been resolved and F_n will be $\{\varphi\}$ -closed. Hence, by Lemma 1 we will have $B_k(\varphi) = B_n(\varphi) = P(\varphi)$ for any $k \geq n$, as required. \square

Remark 2. *Note that the sequence constructed in the proof above is in particular an \mathcal{SL} -closed sequence, which at a certain depth decides all the propositional variables, hence all the sentences, in our finite language. The argument extends also to a countable language. To see this, note that by Lemma 1, for each sentence $\varphi \in \mathcal{SL}$ there is a $j_\varphi \in \mathbb{N}$ such that $P(\varphi) = B_{j_\varphi}(\varphi)$, and $B_k(\varphi) = B_{j_\varphi}(\varphi)$ for each $k \geq j_\varphi$. Hence, what is peculiar to the countable case is the fact that the index n at which $B_n(\varphi)$ equals $P(\varphi)$ may be distinct for distinct elements of \mathcal{SL} . Whether \mathcal{L} is finite or countable, $P(\varphi) = \lim_{k \rightarrow \infty} B_k(\varphi)$.*

Theorem 2. *Let $\mathcal{F} = (F_k, m_k)_{k \in \mathbb{N}}$ be a DBM-sequence with $\{F_k\}_{k \in \mathbb{N}}$ a \mathcal{SL} -closed DBF-sequence free of deep contradictions. Then, the function $P: \mathcal{SL} \rightarrow [0, 1]$, obtained by letting*

$$P(\varphi) = \lim_{k \rightarrow \infty} B_k(\varphi)$$

for any $\varphi \in \mathcal{SL}$, is a probability.

Proof. To prove that P is a probability function, we need to show that it satisfies the two properties (PL1) and (PL2) in the introduction.

For (PL1), assume $\vdash \varphi$ for some $\varphi \in \mathcal{SL}$. Then, since the DBLLs approximate \vdash [7] there is a i such that $\vdash_i \varphi$. On the other hand, since $\{F_k\}_{k \in \mathbb{N}}$ is \mathcal{SL} -closed, there is a $\{\varphi\}$ -closed j -depth forest F_j free of deep contradictions. Hence, by Proposition 2 we will then have $B_j(\varphi) = 1$, and by Lemma 1, $B_n(\varphi) = 1$ for each $n \geq \max\{j, i\}$.

Let us now assume that $\vdash \neg(\varphi \wedge \psi)$. We want to show that $P(\varphi \vee \psi) = P(\varphi) + P(\psi)$. As before, for some i , we will have that $B_i(\varphi \wedge \psi) = 0$, and the same will hold for any $n \geq i$.

Now, since \mathcal{F} is \mathcal{SL} -closed, there will be a j such F_j is $\{\varphi\}$ -closed. and a k such that F_k is $\{\psi\}$ -closed. Now, let us consider any $n \geq \max\{i, j, k\}$. First, since $n \geq i$, we have $B_n(\varphi \wedge \psi) = m_n(b_n(\varphi \wedge \psi)) = 0$. Let us now consider the set $b_n(\varphi \vee \psi)$. Since $n \geq \max\{j, k\}$, F_n will be both $\{\varphi\}$ -closed and $\{\psi\}$ -closed, hence we will have

$$\begin{aligned} b_n(\varphi \vee \psi) &= \{\alpha \in Le(F_n) \mid \alpha \vdash_0 \varphi \vee \psi\} \\ &= \{\alpha \in Le(F_n) \mid \alpha \vdash_0 \varphi, \alpha \vdash_0 \psi\} \\ &\quad \cup \{\alpha \in Le(F_n) \mid \alpha \vdash_0 \varphi, \alpha \vdash_0 \neg\psi\} \\ &\quad \cup \{\alpha \in Le(F_n) \mid \alpha \vdash_0 \neg\varphi, \alpha \vdash_0 \psi\} \end{aligned} \tag{7}$$

where all three sets are clearly disjoint. On the other hand, we have

$$\{\alpha \in Le(F_n) \mid \alpha \vdash_0 \varphi, \alpha \vdash_0 \psi\} = \{\alpha \in Le(F_j) \mid \alpha \vdash_0 \varphi \wedge \psi\}$$

and $m_n(\{\alpha \in Le(F_n) \mid \alpha \vdash_0 \varphi \wedge \psi\}) = m_n(b_n(\varphi \wedge \psi)) = 0$. From this and Equation 7, it follows that $m_n(b_n(\varphi \vee \psi))$ equals:

$$\begin{aligned} m_n\{\alpha \in Le(F_n) \mid \alpha \vdash_0 \varphi, \alpha \vdash_0 \neg\psi\} + m_n\{\alpha \in Le(F_n) \mid \alpha \vdash_0 \neg\varphi, \alpha \vdash_0 \psi\} \\ = m_n(b_n(\varphi)) + m_n(b_n(\psi)) \end{aligned}$$

From the latter, we finally get $B_n(\varphi \vee \psi) = B_n(\varphi) + B_n(\psi)$. This will hold for any $n \geq \max\{i, j, k\}$, hence we finally obtain our desired result, i.e. $P(\varphi \vee \psi) = P(\varphi) + P(\psi)$. \square

4. WHEN k -DEPTH BELIEF FUNCTIONS ARE TRACTABLE

We now turn to identifying the conditions under which the approximations granted by Theorem 1 are computationally feasible.

Definition 10. A sequence $\mathcal{F} = (F_k)_{k \in \mathbb{N}}$ of depth-bounded forests with initial support $Supp$ is said to be:

- uniform iff, for each k , all the nodes in the trees in F_k contain the same hypothetical information, with the exception of the root.

- analytic iff it is $S(\text{Supp})$ -maximal, and all the nodes in the trees in F_k contain as hypothetical information only formulas in $S(\text{Supp})$.
- A Π -maximal uniform(analytic) sequence is a sequence which is Π -maximal among the uniform(analytic) sequences.

Note that the restriction to uniform sequences is immaterial for the result in Theorem 1, since the approximating sequence built in the proof is based on $\{*\}$, hence it is trivially uniform.

Example 6. Let $\text{Supp} = \{p \vee q, s \vee r\}$. The 1-depth forest

$$\begin{array}{cc} (p \vee q) & (s \vee r) \\ \wedge & \wedge \\ p & \neg p \end{array}$$

is uniform, while the 1-depth forest

$$\begin{array}{cc} (p \vee q) & (s \vee r) \\ \wedge & \wedge \\ p & \neg s \end{array}$$

is not. Moreover, the forests in Example 3 are all trivially uniform, since they reduce to just single trees, while the forests at Depth 1 in Example 4 are not uniform.

Let us now consider two core reasoning tasks with our depth-bounded belief functions. Following [30, 16, 22, 14], which investigate the complexity of reasoning with probability functions, we assume that an agent is provided with n linear constraints on the sentences $\varphi_1, \dots, \varphi_p$, of the form:

$$\sum_{j=1}^p d_{ij} B(\varphi_j) \leq z_i \quad i = 1, \dots, n \quad d_{ij}, z_i \in \mathbb{Q}. \quad (8)$$

Let us equivalently rewrite (8) as

$$\sum_{j=1}^{p_1} a_{ij} B(\varphi_j) + w_i \leq \sum_{j=1}^{p_2} b_{ij} B(\varphi_j) + v_i \quad (9)$$

where $i = 1, \dots, n$ $a_{ij}, b_{ij}, w_i, v_i \in \mathbb{N}, p_1 + p_2 = p$. The GENPSAT problem (see e.g. [4]) consists in finding out whether there is a probability function over \mathcal{SL} satisfying such a system. It is well-known that GENPSAT, as well as other minor variants involving linear constraints, is NP-complete [30, 16]. For our present purposes, a variation on the problem will be more appropriate: we will rather aim to find out whether there are depth-bounded Belief and Plausibility functions, which can be seen as approximations of a

probability function satisfying the given linear constraints. We will thus consider decision problems which stand to our k -depth logics and k -depth belief functions as the GENPSAT problem stands to classical logic and classical probability functions.

Let us start from the 0-depth case:

GENSAT₀ Problem

INPUT: The set of p sentences and n linear constraints in (9).

PROBLEM: Find a 0-depth belief function over $Supp = \{\varphi_1, \dots, \varphi_p\}$ such that for each $i \in \{1, \dots, n\}$:

$$\sum_{j=1}^{p_1} a_{ij} B_0(\varphi_j) + w_i \leq \sum_{j=1}^{p_2} b_{ij} Pl_0(\varphi_j) + v_i$$

where $a_{ij}, b_{ij}, w_i, v_i \in \mathbb{N}$.

By Definition 9 the problem boils down to finding a solution for the following n system of linear inequalities in unknowns $m_0(\varphi_1), \dots, m_0(\varphi_p)$:

$$\begin{aligned} \sum_{j=1}^{p_1} a_{ij} \sum_{\substack{k=1, \dots, p_1 \\ \varphi_k \vdash_0 \varphi_j}} m_0(\varphi_k) + w_i &\leq \sum_{j=1}^{p_2} b_{ij} \sum_{\substack{k=1, \dots, p_2 \\ \varphi_k \not\vdash_0 \neg \varphi_j}} m_0(\varphi_k) + v_i & i = 1, \dots, n \\ m_0(\varphi_j) &= 0 & \text{if } \varphi_j \vdash_0 \perp. \end{aligned} \quad (10)$$

Let us recall that $|inc(Supp)| = |\{\alpha \in Supp \mid \alpha \vdash_0 \perp\}|$ stands for the count of inconsistent sentences in the support.

Theorem 3. *The GENSAT₀ problem is solvable in $\mathbf{PTIME}(|S(Supp)| + n)$.*

Proof. By results in [7], finding out whether $\varphi_k \vdash_0 \varphi_j$, $\varphi_k \not\vdash_0 \neg \varphi_j$ and whether $\varphi_j \vdash_0 \perp$ requires time polynomial in $|S(Supp)|$. Note that the system (10) has size $p \times (n + p + 1 + |inc(Supp)|)$, so finding a solution requires polynomial time as well. Hence our problem is in $\mathbf{PTIME}(|S(Supp)| + n)$. \square

Let us now consider the problem of finding out whether there is a k -depth belief function, for a given $k > 0$, which approximates a probability function satisfying the constraints in (9). The idea is to fix k and an initial $Supp$, and then consider all the $S(Supp)$ -maximal depth-bounded forests F_k , in order to check whether there is a depth-bounded Belief function B_k defined on such forests, which satisfies the constraints. We will take as initial support the set $Supp = \{\varphi_1, \dots, \varphi_p\}$ and assume that the forests are uniform. Note that this assumption prevents the number of forests to grow exponentially

with $|Supp|$. Let us now fix a $k > 0$ and consider the following problem:

GENSAT_k Problem.

INPUT: The n constraints in (9).

PROBLEM: Find out whether there is a finite DBM-sequence $\mathcal{F} = (F_j, m_j)_{j \leq k}$, where $\{F_j\}_{j \leq k}$ is a uniform, analytic, finite DBF-sequence based on $Supp = \{\varphi_1, \dots, \varphi_p\}$, such that

$$\sum_{j=1}^{p_1} a_{ij} B_k(\varphi_j) + w_i \leq \sum_{j=1}^{p_2} b_{ij} Pl_k(\varphi_j) + v_i$$

where $i = 1, \dots, n$ $a_{ij}, b_{ij}, w_i, v_i \in \mathbb{N}, p_1 + p_2 = p$.

Note that, also in this case with $k > 0$, we have still taken $Supp = \{\varphi_1, \dots, \varphi_p\}$. Since these are the sentences involved in the formulation of the constraints, we can take them to be the initial information provided to an agent, i.e. the information that can be probabilistically quantified by the agent without the use of further hypothetical information.

Recalling Definition 9, solving the problem $GENSAT_k$ amounts to finding an F_k based on $Supp = \{\varphi_1, \dots, \varphi_p\}$, such that the following system, in the unknowns $m_k(\alpha)$, for $\alpha \in Le(F_k)$, has a solution:

$$\begin{aligned} \sum_{j=1}^{p_1} a_{ij} \sum_{\alpha \in b_k(\varphi_j)} m_k(\alpha) + w_i &\leq \sum_{j=1}^{p_2} b_{ij} \sum_{\alpha \in pl_k(\varphi_j)} m_k(\alpha) + v_i && \text{for each } i = 1, \dots, n \\ m_k(\alpha) &\geq 0 && \text{for each } \alpha \in Le(F_k) \\ \sum_{\alpha \in Le(F_k)} m_k(\alpha) &= 1 \\ m_k(\alpha) &= 0 && \text{if } \alpha \vdash_0 \perp \end{aligned} \tag{11}$$

Theorem 4. $GENSAT_k$ is solvable in $\mathbf{PTIME}(|S(Supp)| + n)$.

Proof. Consider the number of possible forests of depth k , based on $Supp$. Since we assume that the forests are uniform, this reduces to the number of possible trees, which in turn reduces to the possible choices of formulas for expanding each node. The number of such choices at depth k will be $2^k - 1$ and, since we restrict to analytic sequences, we will only consider sentences in $S(Supp)$ as candidates. We thus obtain that the number of possible forests at depth k is bounded above by $|S(Supp)|^{2^k-1}$.

We then select among those $|S(Supp)|^{2^k-1}$ different forests, those which are maximal with respect to $S(Supp)$, see Definition 3. This is verified by

running, for each forest, and for each formula φ in $S(Supp)$, the quadratic algorithm e.g. in [7], to verify whether the leaves decide φ .

For each remaining selected forest F_k of depth k , the number of leaves in $Le(F_k)$ will be bounded above by $|Supp| \cdot 2^k$, which is linear in $|Supp|$ (once k is fixed, 2^k is constant).

We then need to check whether the set of sentences in $Le(F_k)$ may satisfy system (11). Each such system has size bounded by $2^k \times (n + 2^k + 1 + |inc(Le(F_k))|)$. Hence we have polynomially many systems of polynomial size. Finally, since solving each such system still requires polynomial time (by standard linear algebraic methods), we obtain our claim. \square

The problems considered so far are decision problems, and in particular variants of the well-known SAT problem and GENPSAT problem for propositional satisfiability, and (generalized) probabilistic satisfiability, respectively. Another class of problems is also quite relevant in the area of probabilistic logic, as presented in [20] and in subsequent developments [18]. The idea is to provide a probabilistic generalization of inference, by addressing the following problem:

Let $\varphi_1, \dots, \varphi_p, \varphi$ be sentences. Take any probability function P such that $P(\varphi_1), \dots, P(\varphi_p)$ obey a given set of constraints. What kind of constraints follow on $P(\varphi)$?

In [20, 18] the constraints imposed on the probabilities of $\varphi_1, \dots, \varphi_p$ are typically intervals, and the goal is to identify the “smallest” interval $[a, b] \subseteq [0, 1]$ such that $P(\varphi) \in [a, b]$, whenever P satisfies the constraints on the premises. This reduces in the classical setting to a linear programming problem. We will consider here a more general problem, allowing constraints as in (9), and searching for solutions based on depth-bounded belief functions. We have then the following:

B_k INF Problem

INPUT: The n constraints in (9), and an arbitrary sentence φ .

PROBLEM: Compute the smallest interval $[a, b]$ such that, for any B_k, Pl_k which satisfy the constraints in (9) as in the GENSAT $_k$ Problem, B_k, Pl_k satisfy $a \leq B_k(\varphi), Pl_k(\varphi) \leq b$.

As we now show, we can just adapt the classical techniques based on linear programming, and Theorem 4 in this section, to obtain that also the

functional problem $B_k\text{INF}$ is solvable in polynomial time. Again, we will let $\text{Supp} = \{\varphi_1, \dots, \varphi_m\}$,

Theorem 5. *The problem $B_k\text{INF}$ is solvable in $\mathbf{FPTIME}(|S(\text{Supp})| + S(\varphi) + n)$.*

Proof. By the argument in the proof of Theorem 4, we identify polynomially many analytic uniform DBF-sequences based on $\text{Supp} = \{\varphi_1, \dots, \varphi_m\}$. For each such system our problem is solved by considering the following linear programming problem.

$$\begin{aligned}
& \textbf{minimize} && \sum_{\alpha \in b_k(\varphi)} m_k(\alpha) \\
& \text{subject to} && \\
& \sum_{j=1}^{p_1} a_{ij} \sum_{\alpha \in b_k(\varphi_j)} m_k(\alpha) + w_i \leq \sum_{j=1}^{p_2} b_{ij} \sum_{\alpha \in pl_k(\varphi_j)} m_k(\alpha) + v_i && \text{for each } i = 1, \dots, n \\
& m_k(\alpha) \geq 0 && \text{for each } \alpha \in Le(F_k) \\
& \sum_{\alpha \in Le(F_k)} m_k(\alpha) = 1 \\
& m_k(\alpha) = 0 && \text{if } \alpha \vdash_0 \perp
\end{aligned} \tag{12}$$

and similarly

$$\begin{aligned}
& \textbf{maximize} && \sum_{\alpha \in pl_k(\varphi)} m_k(\alpha) \\
& \text{subject to} && \\
& \sum_{j=1}^{p_1} a_{ij} \sum_{\alpha \in b_k(\varphi_j)} m_k(\alpha) + w_i \leq \sum_{j=1}^{p_2} b_{ij} \sum_{\alpha \in pl_k(\varphi_j)} m_k(\alpha) + v_i && \text{for each } i = 1, \dots, n \\
& m_k(\alpha) \geq 0 && \text{for each } \alpha \in Le(F_k) \\
& \sum_{\alpha \in Le(F_k)} m_k(\alpha) = 1 \\
& m_k(\alpha) = 0 && \text{if } \alpha \vdash_0 \perp
\end{aligned} \tag{13}$$

Solving the two problems will provide our desired bounds for φ . This can be achieved in polynomial time invoking algorithms for linear programming [34], since the size of the systems to be solved is polynomial as well in the size of $\varphi_1, \dots, \varphi_p, \varphi$. \square

5. CONCLUSIONS

The results of this paper suggest that there should be no methodological tension between requiring that

- i) models of uncertain reasoning should be aimed at realistic, rather than idealised agents;
- ii) the quantification of uncertainty should use all the information actually possessed by an agent and
- iii) the quantification of uncertainty should keep track of the reasoning which leads to the hypothetical resolution of uncertainty – i.e. using information which the agent does not actually possess but which can be obtained by suitably exploiting reasoning by cases.

In consonance with the general view outlined in [8] the results illustrated in this paper suggest that a key to achieving such a realistic model of uncertain reasoning lies in grounding it in logics which are sensitive to information, rather than classical truth-conditions.

Not only our main results may lead to a reconciliation between additive and non-additive approaches to uncertain reasoning, but they also suggest that logic may play a much more prominent role in uncertain reasoning compared to what is seen in current practice. Whilst AI has long benefitted methodologically from logical approaches, Economic theory has done this to a much lesser extent. Building on Examples 4 and 5 future work will focus on this promising application of logical methods in uncertain reasoning.

As to AI applications, future work will be devoted first to the problem of determining, for a specific application-driven problem, the bounds k for which the uncertainty quantification provided by B_k turns out to be adequate. In addition, better algorithms for the problems in Section 4 are desirable and will be investigated. Let us note that the proofs in Section 4 contain essentially brute-force search algorithms, where the search space is bounded by the fixed required depth, and only this prevents an exponential explosion. This leaves open the investigation of algorithms which are more tightly calibrated to exploring only the necessary maximal forests. In future work, we will develop such algorithms and provide suitable time estimations

Finally, as we pointed out in Remark 1 the notion of maximal hypothetical uncertainty resolution makes our framework interpretable normatively. Moreover the tractability results of Section 4 suggest that reasoning with depth-bounded belief functions may be within the means of realistic agents. Hence, the framework introduced in this paper may be taken as a benchmark

for the empirical analysis of what experimental subjects should be expected to do when performing uncertain reasoning tasks.

ACKNOWLEDGMENTS

This research was funded by the Department of Philosophy “Piero Martinetti” of the University of Milan under the Project “Departments of Excellence 2018-2022” awarded by the Ministry of Education, University and Research (MIUR). We are grateful to Marcello D’Agostino for very useful conversations on the theory of Depth-Bounded Boolean Logics which motivated a substantial part of this research. We are also very grateful to two anonymous reviewers for their unrestrained criticisms, which certainly lead this to becoming a much better paper.

REFERENCES

- [1] P. Baldi, M. D’Agostino, and H. Hosni. Depth-Bounded Approximations of Probability. *Communications in Computer and Information Science*, 1239 CCIS:607–621, 2020.
- [2] P. Baldi and H. Hosni. Depth-bounded Belief functions. *International Journal of Approximate Reasoning*, 123:26–40, 2020.
- [3] P. Baldi and H. Hosni. Logical approximations of qualitative probability. In *Proceedings of the 12th International Symposium on Imprecise Probability: Theories and Applications (ISIPTA)*, 2021.
- [4] C. Caleiro, F. Casal, and A. Mordido. Generalized probabilistic satisfiability. *Information and Computation*, 332:39–56, 2017.
- [5] M. D’Agostino and L. Floridi. The enduring scandal of deduction. *Synthese*, 167(2):271–315, 2009.
- [6] M. D’Agostino. An informational view of classical logic. *Theoretical Computer Science*, 606:79–97, 2015.
- [7] M. D’Agostino, M. Finger, and D.M. Gabbay. Semantics and proof-theory of depth bounded Boolean logics. *Theoretical Computer Science*, 480:43–68, 2013.
- [8] M. D’Agostino, T. Flaminio, and H. Hosni. Rational beliefs real agents can have – A logical point of view. *Journal of Machine Learning Research: Workshops and Conference and Proceedings*, 58:1–13, 2016.
- [9] B. de Finetti. Sul significato soggettivo della probabilità. *Fundamenta Mathematicae*, 17:289–329, 1931.
- [10] B. de Finetti. *Theory of Probability, Vol 1*. John Wiley and Sons, 1974.
- [11] D. Dubois and H. Prade. Possibility theory , probability theory and multiple-valued: A clarification. *Annals of Mathematics and Artificial Intelligence*, 32:35–66, 2001.
- [12] D. Dubois and H. Prade. *Possibilistic Logic - An Overview*, in D. M. Gabbay and J. H. Siekmann and J. Woods (eds.), *Computational Logic*, Volume 9 of the Handbook of The History of Logic. Elsevier, 2014.
- [13] D. Ellsberg. Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics*, 75(4):643–669, 1961.

- [14] R. Fagin, J. Y. Halpern, and N. Megiddo. A logic for reasoning about probabilities. *Information and Computation*, 1-2(87):78–128, 1990.
- [15] M. Finger and R. Wasserman. The universe of propositional approximations. *Theoretical Computer Science*, 355(2):153–166, 2006.
- [16] M. Finger and G. De Bona. Probabilistic satisfiability: algorithms with the presence and absence of a phase transition. *Annals of Mathematics and Artificial Intelligence*, 75(3-4):351–389, 2015.
- [17] I. Gilboa, A. Postlewaite, and D. Schmeidler. Is It Always Rational To Satisfy Savage’s Axioms? *Economics and Philosophy*, 25(03):285, 2009.
- [18] R. Haenni, J.W. Romeijn, J. Williamson, and G. Wheeler. *Probabilistic Logics and Probabilistic Networks*. Springer, 2011.
- [19] Ian Hacking. Slightly more realistic personal probability. *Philosophy of Science*, 34(4):311–325, 1967.
- [20] T. Hailperin. *Sentential Probability Logic. Origins, Development, Current Status, and Technical Applications*. Academic Press, 1996.
- [21] J.Y. Halpern. *Reasoning About Uncertainty*. MIT Press, 2003.
- [22] P. Hansen and B. Jaumard. Probabilistic satisfiability. In *Handbook of Defeasible Reasoning and Uncertainty Management Systems.*, volume 5, page 321–367. Dordrecht: Kluwer, 2000.
- [23] J-Y. Jaffray. Coherent bets under partially resolving uncertainty and belief functions. *Theory and Decision*, 26(2):99–105, 1989.
- [24] J.M. Keynes. *Treatise on Probability*. Harper & Row, 1921.
- [25] F.H. Knight. *Risk, uncertainty and profit*. Beard Books Inc, 1921.
- [26] H.E. Kyburg and C.M. Teng. *Uncertain Inference*. Cambridge University Press, 2001.
- [27] G. Lakemeyer and H.J. Levesque. A First-Order Logic of Limited Belief Based on Possible Worlds *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR 2020)*, 624–635, 2020.
- [28] D. V. Lindley. Scoring Rules and the Inevitability of Probability. *International Statistical Review*, 50(1):1–11, 1982.
- [29] P. Marquis, O. Papini, and H. Prade. *A Guided Tour of Artificial Intelligence Research 1: Knowledge Representation, Reasoning and Learning*. Springer, 2020.
- [30] J.B. Paris. *The uncertain reasoner’s companion: A mathematical perspective*. Cambridge University Press, 1994.
- [31] L. J. Savage. Difficulties in the theory of personal probability. *Philosophy of Science*, 34(4):305–310, 1967.
- [32] L.J. Savage. *The Foundations of Statistics*. Dover, 2nd edition, 1972.
- [33] D. Schmeidler. Subjective Probability and Expected Utility without Additivity. *Econometrica*, 57(3):571–587, 1989.
- [34] A. Schrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, 1987.
- [35] G. Shafer. *A mathematical theory of evidence*. Princeton University Press, 1976.
- [36] G. Shafer. Perspectives on the theory and practice of belief functions. *Int. J. Approx. Reasoning*, 4(5-6):323–362, 1990.
- [37] M. Sipser. *Introduction to the theory of computation*. PWS Publishing Company, 1997.
- [38] K. Stanovich *What Intelligence Tests Miss: The psychology of Rational Thought* Yale University Press, 2009
- [39] P. Walley. *Statistical Reasoning with Imprecise Probabilities*. Wiley, 1991.
- [40] J. Williamson. *In defence of objective Bayesianism*. Oxford University Press, 2010.

DEPARTMENT OF PHILOSOPHY, UNIVERSITY OF MILAN

Email address: `paolo.baldi,hykel.hosni@unimi.it`