Proceedings of Machine Learning Research 58 (2016) 97-109

Rational Beliefs Real Agents Can Have – A Logical Point of View

Marcello D'Agostino MARCELLO.DAGOSTINO@UNIMI.IT
Department of Philosophy
University of Milan
20122 Milano, Italy

Tommaso Flaminio TOMMASO.FLAMINIO@UNINSUBRIA.IT
Department of Pure and Applied Sciences
University of Insubria
21100 Varese, Italy

Hykel Hosni HYKEL.HOSNI@UNIMI.IT
Department of Philosophy
University of Milan
20122 Milano, Italy.

Editor: Tatiana V. Guy, Miroslav Kárný, David Rios-Insua, David H. Wolpert

Abstract

The purpose of this note is to outline a framework for uncertain reasoning which drops unrealistic assumptions about the agents' inferential capabilities. To do so, we envisage a pivotal role for the recent research programme of *depth-bounded Boolean logics* (D'Agostino et al., 2013). We suggest that this can be fruitfully extended to the representation of rational belief under uncertainty. By doing this we lay the foundations for a prescriptive account of rational belief, namely one that *realistic agents*, as opposed to idealised ones, can feasibly act upon.

Keywords: Prescriptive rationality, tractability, logic-based probability, Bayesian norms

1. Introduction and motivation

Probability is traditionally the tool of choice for the quantification of uncertainty. Since Jacob Bernoulli's 1713 *Ars Conjectandi*, a number of arguments have been put forward to the effect that departing from a probabilistic assessment of uncertainty leads to *irrational* patterns of behaviour. This contributed to linking tightly the rules of probability with the defining norms of rationality, as fixed by the well known results of de Finetti (1974); Savage (1972). Lindley (2006) and Parmigiani and Inoue (2009) provide recent introductory reviews.

Over the past few decades, however, a number of concerns have been raised against the adequacy of probability as a norm of rational reasoning and decision making. Following the lead of Ellsberg (1961), whom in turn found himself on the footsteps of Knight (1921) and Keynes (1921), many decision theorists take issue with the idea that probability provides adequate norms for rationality. This is put emphatically in the title of Gilboa et al. (2012), a paper which has circulated for almost a decade before its publication. As a result, considerable formal and conceptual effort has gone into extending the scope of the probabilistic representation of uncertainty, as illustrated for instance by Gilboa and Marinacci (2013). Related to this, is the large family of *imprecise probability* models, and its decision-theoretic offsprings, which constitute the cutting edge of uncertain reasoning research, see e.g. Augustin et al. (2014).

One key commonality between "non Bayesian" decision theory and the imprecise probabilities approach is the fact they take issue with the identification of "rationality" and "probability" on representational grounds. For they insist on the counterintuitive consequences of assuming that the rational representation of uncertainty necessitates the Bayesian norms, and in particular that all uncertainty is to be represented probabilistically.

This note makes a case for adding a *logical* dimension to this ongoing debate. Key to this is a logical framing of probability. As recalled explicitly below, probability functions are normalised on classical tautologies. That is to say that a Bayesian agent is required to assign maximum degree of belief to *every* tautology of the propositional calculus. However classic results in computational complexity imply that the problem of deciding whether a given sentence is a tautology, exceeds, in general, what is considered to be *feasible*. Hence, probability imposes a norm of rationality which, under widely agreed hypotheses, realistic agents cannot be expected to meet. A related concern had already been put forward by Savage (1967), but this did't lead proponents of the Bayesian approach to take the issue seriously. This is precisely what the research outlined in this note aims to do.

By framing the question logically, we can offer a perspective on the problem which highlights the role of *classical* logic in determining the unwelcome features of canonical Bayesian rationality (Section 3). This suggests that a normatively reasonable account of rationality should to take a step back and rethink the logic in the first place.

The recently developed framework of Depth-Bounded Boolean logics (DBBLs) is particularly promising in this respect. By re-defining the meaning of logical connectives in terms of *information actually possessed by the agent* DBBLs give rise to a hierarchy of logics which (i) accounts for some key aspects of the asymmetry between knowledge and ignorance and (ii) provide computationally feasible approximations to classical logic. Section 4.2 reviews informally the core elements of this family of logics.

Finally, Section 5 outlines the applicability of this framework to probabilistic reasoning. In particular it points out how the hierarchy of DBBLs to can serve to define a hierarchy of prescriptively rational approximation of Bayesian rationality.

2. Bayesian rationality

In a number of areas, from Economics to the Psychology of reasoning and of course Statistics, probability has been defended as the norm of rational belief. Formally this can be seen to imply a normative role also for classical logic. So the Bayesian norms of rationality are best viewed as combination of probability and logic.

This allows us to distinguish two lines of criticisms against Bayesian rationality. First, it is often pointed out that probability washes out a natural asymmetry between knowledge and ignorance. Second, the intractability of classical logical reasoning is often suggested to deprive the normative theory of practical meaning. Both lines of criticisms can be naturally linked to the properties of *classical* logic.

2.1 Against the probability norm: the argument from information

Uncertainty has to do, of course, with not knowing, and in particular not knowing the outcome of an event of interest, or the value of a random variable. Ignorance has more subtle features, and is often thought of as our inability to quantify our own uncertainty. Knight (1921) gave this impalpable distinction an operational meaning in actuarial terms. He suggested the presence of ignorance is detected by the absence of a compete insurance market for the goods at hand. On the contrary, a complete insurance market provides an operational definition of *probabilistically quantifiable* uncertainty. Contemporary followers of Knight insist that not all uncertainty is probabilistically quantifiable and seek to introduce more general norms of rational belief and decision under "Knightian uncertainty" or "ambiguity".

A rather general form of the argument from information against Bayesian rationality is summarised by the following observation by Schmeidler (1989):

> The probability attached to an uncertain event does not reflect the heuristic amount of information that led to the assignment of that probability. For example, when the information on the occurrence of two events is symmetric they are assigned equal probabilities. If the events are complementary the probabilities will be 1/2 independent of whether the symmetric information is meager or abundant.

Gilboa (2009) interprets Schmeidler's observation as expressing a form of "cognitive unease", namely a feeling that the theory of subjective probability which springs naturally from Bayesian epistemology, is silent on one fundamental aspect of rationality (in its informal meaning). But why is it so? Suppose that some matter is to be decided by the toss of a coin. According to Schmeidler's line of argument, I should prefer tossing my own, rather than some one else's coin, on the basis, say of the fact that I have never observed signs of "unfairness" in my coin, whilst I just don't know anything about the stranger's coin. See also Gilboa et al. (2012); Gilboa (2009). This argument is of course reminiscent of the Ellsberg two-urns problem, which had been anticipated in Keynes (1921).

Similar considerations have been put forward in artificial intelligence and in the foundations of statistics. An early amendment of probability theory aimed at capturing the asymmetry between uncertainty and ignorance is known as the theory of Belief Functions (Shafer, 1976; Denoeux, 2016). Key to representing this asymmetry is the relaxation of the additivity axiom of probability. This in turn may lead to situations in which the *probabilistic excluded middle* does not hold. That is to say an agent could rationally assign belief less than 1 to the classical tautology $\theta \vee \neg\theta$. Indeed, as we now illustrate, the problem with normalising on tautologies is much more general.

2.2 Against the logic norm: the argument from tractability

Recall that classical propositional logic is *decidable* in the sense that for each sentence θ of the language there is an effective procedure to decide wether θ is a tautology or not. Such a procedure, however, is unlikely to be *feasible*, that is to say executable in practice. In terms of the theory of computational complexity this means that there is probably no algorithm running in polynomial time. So, a consequence of the seminal 1971 result by Stephen Cook, the tautology problem for classical logic is widely believed to be *intractable*.

If this conjecture is correct, we are faced with a serious foundational problem when imposing the normalisation of probability on tautology. For we are imposing agents constraints of rationality which they simply may never be able to satisfy.

It is remarkable that L.J. Savage had anticipated this problem with the Bayesian norms he centrally contributed to defining. To this effect he observed in Savage (1967) the following:

> A person required to risk money on a remote digit of π would have to compute that digit in order to comply fully with the theory though this would really be wasteful if the cost of computation were more than the prize involved. For the postulates of the theory imply that you should behave in accordance with the *logical implications of all that you know.* Is it possible to improve the theory in this respect, making allowance within it for the cost of thinking, or would that entail paradox [...] , as I am inclined to believe but unable to demonstrate ? If the remedy is not in changing the theory but rather in the way in which we attempt to use it, clarification is still to be desired. (Our emphasis)

Fifty years on, the difficulty pointed out by Savage failed to receive the attention it deserves. As the remainder of this note illustrates, however, framing the issue logically brings about significant improvements in our understanding of the key issues, paving the way for a tractable approximation of Bayesian rationality – or rational beliefs real agents can have.

3. Logic, algebra and probability

A well-known representation result (see, e.g. Paris (1994)) shows that every probability function arises from distributing the unit mass across the 2^n atoms of the Boolean (Lindenbaum) Algebra generated by the propositional language $L = \{p_1, \dots p_n\}$, and conversely, that a probability function on L is completely determined by the values it takes on such atoms. Such a representation makes explicit the dependence of probability on classical logic. This has important and often underappreciated consequences. Indeed logic plays a twofold role in the theory of probability. First, logic provides the language in which events –the bearers of probability– are expresses, combined and evaluated. The precise details depend on the framework. See Flaminio et al. (2014) for a characterisation of probability on classical logic, and Flaminio et al. (2015) for the general case of Dempster-Shafer belief functions on many-valued events.

In measure-theoretic presentations of probability, events are identified with subsets of the field generated by a given sample space Ω. A popular interpretation for Ω is that of the elementary outcomes of some experiment, a view endorsed by A.N. Kolmogorov, who insisted on the generality of his axiomatisation. More precisely, let $\mathcal{M} = (\Omega, \mathcal{F}, \mu)$ a *measure space* where, $\Omega = \{\omega_1, \omega_2 \dots\}$ is the set of elementary outcomes, $\mathcal{F} = 2^\Omega$ is the field of sets ($\sigma-$algebra) over Ω. We call *events* the elements of $\mathcal{F}$, and $\mu : \mathcal{F} \to [0, 1]$ a *probability measure* if it is normalised, monotone and σ-additive, i.e.

(K1) $\mu(\Omega) = 1$

(K2) $A \subseteq B \Rightarrow \mu(A) \leq \mu(B)$

(K3) If $\{E\}_i$ is a countable family of pairwise disjoint events then $P(\bigcup_i E_i) = \sum_i P(E_i)$

The Stone representation theorem for Boolean algebras and the representation theorem for probability functions recalled above guarantee that the measure-theoretic axiomatisation of probability is equivalent to the logical one, which is obtained by letting a function from the language L to the real unit interval be a *probability function* if

(PL1) $\models \theta \Rightarrow P(\theta) = 1$

(PL2) $\models \neg(\theta \wedge \phi) \Rightarrow P(\theta \vee \phi) = P(\theta) + P(\phi)$.

Obvious as this logical "translation" may be, it highlights a further role for logic in the theory of probability, in addition that is to the linguistic one pointed out above. This role is best appreciated by focussing on the consequence relation $\models$ and can be naturally referred to as inferential.

In its measure-theoretic version, the normalisation axiom is quite uncontroversial. Less so, if framed in terms of classical tautologies, as in PL1. Indeed both arguments against Bayesian norms discussed informally above, emerge now formally. The first is to do with the fact that $\models$ interprets symmetrically "knowledge" and "ignorance" as captured by the fact that $\models \theta \vee \neg\theta$ is a tautology. Indeed similarly bothersome consequences follow directly from $PL1$ and $PL2$, namely

1. $P(\neg\theta) = 1 - P(\theta)$
2. $\theta \models \phi \Rightarrow P(\theta) \leq P(\phi)$

2 implies that if θ and ϕ are logically equivalent they get equal probability.

The argument from information recalled in Section 2.1 above clearly has its logical roots in the semantics of classical logic.

Similarly, the argument from tractability of Section 2.2 leads one into questioning the desirability of normalising probability on *any* classical tautology. Taken as a norm of rationality this requires agents to be capable of reasoning beyond what is widely accepted as feasible. Again, the unwelcome features of probability are rooted in classical logic.

A further, important, feature which emerges clearly in the logical presentation of probability is that uncertainty is resolved by appealing to the semantics of classical logic. This leads to the piecemeal identification of "events" with "sentences" of the logic. This identification, however, is not as natural as one may think.

On the one hand, an *event*, understood classically, either happens or not. A sentence expressing an event, on the other hand is evaluated in the binary set as follows

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

Hence, the probability of an event $P(\theta) \in [0, 1]$ measures the agent's degree of belief that the event did or will obtain. Finding this out is, in most applications, relatively obvious. However, as pointed out in Flaminio et al. (2014), a general theory of what it means for "states of the world" to "resolve uncertainty" is far from trivial.

A more natural way of evaluating events arises taking an *information-based* interpretation of uncertainty resolution. The key difference with the previous, classical case, lies in the fact that this leads naturally to a *partial* evaluation of events, that is

$$v^i(\theta) = \begin{cases} 1 \text{ if I am informed that } \theta \\ 0 \text{ if I am informed that } \neg\theta \\ \bot \text{ if I am not informed about } \theta. \end{cases}$$

Quite obviously standard probability logic does *not* apply here, because the classical resolution of uncertainty has no way of expressing the $\bot$ condition.

As the next section shows, by looking for a logic which fixes this information asymmetry, we will also find a logic which deals successfully with the tractability problem.

4. An informational view of propositional logic

The main idea underlying the informational view of classical propositional logic is to replace the notions of "truth" and "falsity", by "informational truth" and "informational falsity", namely *holding the information* that a sentence φ is true, respectively false. Here, by saying that an agent a *holds* the information that φ is true or false we mean that this information (i) is *accepted* by a in the sense that a is ready to *act upon it*[1] (ii) it is *feasibly available* to a, in the sense that a has the means to obtain it in practice (and not only in principle); given the (probable) intractability of classical propositional logic this condition is not in general preserved by the corresponding consequence relation.

Clearly, these notions do not satisfy the informational version of the Principle of Bivalence: it may well be that for a given φ, we neither hold the information that φ is true, nor do we hold the information that φ is false. Knowledge and ignorance are not treated symmetrically under the informational semantics. However, in this paper we assume that they do satisfy the informational version of the Principle of Non-Contradiction: no agent can *actually* possess both the information that φ is true and the information that φ is false, as this could be deemed to be equivalent to possessing no definite information about φ.[2]

4.1 Informational semantics

We use the values 1 and 0 to represent, respectively, informational truth and falsity. When a sentence takes neither of these two defined values, we say that it is *informationally indeterminate*. It is technically convenient to treat informational indeterminacy as a third value that we denote by "$\bot$".[3] The three values are partially ordered by the relation $\preceq$

1. The kind of justification for this acceptance and whether or not the agent is human or artificial do not concern us here. Acceptance may include some (possibly non-conclusive) evidence that a deems sufficient for acceptance, or communication from some external source that a regards as reliable.
2. Notice that this assumption does not rule out the possibility of *hidden* inconsistencies in an agent's information state, but only of inconsistencies that can be feasibly detected by that agent. It is, however, possible to investigate paraconsistent variants of this semantics in which even this weak informational version of the Principle of Non-Contradiction is relaxed. This will be the subject of a subsequent paper.
3. This is the symbol for "undefined", the bottom element of the information ordering, not to be confused with the "falsum" logical constant.

$\wedge$	1	0	$\bot$
1	1	0	$\bot$
0	0	0	0
$\bot$	$\bot$	0	$\bot, 0$

$\vee$	1	0	$\bot$
1	1	1	1
0	1	0	$\bot$
$\bot$	1	$\bot$	$\bot, 1$

$\neg$	
1	0
0	1
$\bot$	$\bot$

Figure 1: Informational tables for the classical operators

such that $v \preceq w$ ("v is less defined than, or equal to, w") if, and only if, $v = \bot$ or $v = w$ for $v, w \in \{0, 1, \bot\}$.

Note that the old familiar truth tables for $\wedge, \vee$ and $\neg$ are still intuitively sound under this informational reinterpretation of 1 and 0. However, they are no longer exhaustive: they do not tell us what happens when one or all of the immediate constituents of a complex sentence take the value $\bot$. A remarkable consequence of this approach is that the semantics of $\vee$ and $\wedge$ becomes, as first noticed by Quine (1973, pp. 75–78), *non-deterministic*. In some cases an agent a may accept a disjunction $\varphi \vee \psi$ as true while abstaining on both components φ and ψ. To take Quine's own example, if I cannot distinguish between a mouse and chipmunk, I may still hold the information that "it is a mouse or a chipmunk" is true while holding no definite information about either of the sentences "it is a mouse" and "it is a chipmunk". In other cases, e.g. when the component sentences are "it is a mouse" and "it is in the kitchen" and I still hold no definite information about either, the most natural choice is to abstain on the disjunction. Similarly, a may reject a conjunction $\varphi \wedge \psi$ as false while abstaining on both components. To continue with Quine's example, I may hold the information that "it is a mouse and a chipmunk" is false, while holding no definite information about either of the two component sentences. But if the component sentences are "it is a mouse" and "it is in the kitchen" and I abstain on both, I will most probably abstain also on their conjunction. In fact, this phenomenon is quite common as far as the ordinary notion of information is concerned and the reader can figure out plenty of similar situations. Thus, depending on the "informational situation", when φ and ψ are both assigned the value $\bot$, the disjunction $\varphi \vee \psi$ may take the value 1 or $\bot$, and the conjunction $\varphi \wedge \psi$ may take the value 0 or $\bot$.

As a consequence of this informational interpretation, the traditional truth-tables for the $\vee$, $\wedge$ and $\neg$ should be replaced by the "informational tables" in Figure 1, where the value of a complex sentence, in some cases, is not uniquely determined by the value of its immediate components.[4] A non-deterministic table for the informational meaning of the Boolean conditional can be obtained in the obvious way, by considering $\varphi \to \psi$ as having the same meaning as $\neg\varphi \vee \psi$ (see D'Agostino, 2015, p. 82).

4. In his (Quine, 1973) Quine calls them "verdict tables" and the values are "assent", "dissent" and "abstain". This non-deterministic semantics was subsequently and independently re-proposed (with no apparent connection with the intuitive interpretation given by Quine) by Crawford and Etherington (1998) who claimed without proof that it provides a characterization of unit resolution (a tractable fragment of resolution that requires formulae to be translated into clausal form). The general theory of non-deterministic semantics for logical systems has been brought to the attention of the logical community and extensively investigated (with no special connection with tractablity) by Arnon Avron and co-authors (see Avron and Zamansky (2011) for an overview).

4.2 Depth-bounded Boolean logics

In (D'Agostino et al., 2013) and (D'Agostino, 2015) it is shown that the informational semantics outlined in the previous section provides the basis to define an infinite hierarchy of *tractable* deductive systems (with no syntactic restriction on the language adopted) whose upper limit coincides with classical propositional logic. As will be clarified in the sequel the tractability of each layer is a consequence of the shift from the classical to the informational interpretation of the logical operators (that is the same throughout the hierarchy) and on an upper bound on the nested use of "virtual information", i.e. information that the agent *does not actually hold*, in the sense specified in the previous section.

Definition 1 *A* 0*-depth information state is a valuation V of the formulae in L that agrees with the informational tables.*

Note that, given the non-determinism of the informational tables, the valuation V is not uniquely determined by an assigment of values to the atomic sentences. For example the valuation V_1 that assigns $\bot$ to both p and q and $\bot$ to $p \vee q$ is as admissible as the valuation V_2 that still assigns $\bot$ to both p and q, but 1 to $p \vee q$. Let S_0 be the set of all 0-depth information states.

Definition 2 *We say that φ is a 0-depth consequence of a finite set Γ of sentences, and write $\Gamma \vDash_0 \varphi$, when*

$$(\forall V \in S_0)\, V(\Gamma) = 1 \Longrightarrow V(A) = 1.$$

We also say that Γ is 0-depth inconsistent*, and write $\Gamma \vDash_0$ if there is no $V \in S_0$ such that $V(\Gamma) = 1$.*

It is not difficult to verify that $\vDash_0$ is a Tarskian consequence relation, i.e., it satisfies reflexivity, monotonicity, transitivity and substitution invariance.

In fact, it can be shown that we do not need to consider valuations of the whole language L but can restrict our attention to the subformulae of the formulae that occur as premises and conclusion of the inference under consideration. Let us call *search space* any finite set Λ of formulae that is closed under subformulae, i.e., if φ is a subformula of a formula in Λ, $\varphi \in \Lambda$.

Definition 3 *A 0-depth information state over a search space Λ is a valuation V of Λ that agrees with the informational tables.*

Let S_0^{Λ} be the set of all 0-depth information states over a search space Λ. Given a finite set Δ of formulae, let us write Δ^* to denote the search space consisting of all the subformulae of the formulae in Δ. Then, it can be shown that:

Theorem 4 $\Gamma \vDash_0 \varphi$ *if and only if* $(\forall V \in S_0^{(\Gamma\cup\{\varphi\})^*})\, V(\Gamma) = 1 \Longrightarrow V(A) = 1$*. Moroever, $\Gamma \vDash_0$ if and only if there is no $V \in S_0^{(\Gamma\cup\{\varphi\})^*}$ such that $V(\Gamma) = 1$.*

On the basis of the above result, in (D'Agostino et al., 2013) it is shown that $\vDash_0$ is *tractable*:

Theorem 5 *Whether or not* $\Gamma \vDash_0 \varphi$ *(*Γ *is* 0*-depth inconsistent) can be decided in time* $O(n^2)$ *where* n *is the total number of occurrences of symbols in* $\Gamma \cup \{\varphi\}$ *(in* Γ*).*

A simple proof system that is sound and complete with respect to $\vDash$ is shown in (D'Agostino et al., 2013; D'Agostino, 2015) in the form of a set of introduction and elimination rules (in the fashion of Natural Deduction) that are only based on *actual information*, i.e., information that is held by an agent, with no need for *virtual information*, i.e., simulating information that does not belong to the current information state, as happens in case-reasoning or in some ways of establishing a conditional (as in the introduction rule for the conditional in Gentzen-style natural deduction).

The subsequent layers of the hierarchy depend on fixing an upper bound on the depth at which the nested use of virtual information is allowed.

Let $\sqsubseteq$ be the partial ordering of 0-depth information states (over a given search space) defined as follows: $V \sqsubseteq V'$ if and only if V' is a refinement of V or is equal to V, that is, for every formula φ in the domain of V and V', $V(\varphi) \neq \bot$ implies that $V'(\varphi) = V(\varphi)$.

Definition 6 *Let* V *be a 0-depth information state over a search space* Λ*.*

- $V \Vdash_0 \varphi$ *if and only if* $V(\varphi) = 1$
- $V \Vdash_{k+1} \varphi$ *if and only if*

$$(\exists \psi \in \Lambda)(\forall V' \in S_0^{\Lambda})\, V \sqsubseteq V' \text{ and } V'(\psi) \neq \bot \Longrightarrow V' \Vdash_k \varphi.$$

Here $\Vdash_j$, with $j \in \mathbb{N}$, is a kind of "forcing" relation and the shift from one level of depth to the next is determined by simulating refinements of the current information state in which the value of some $\psi \in \Lambda$ is defined (either 1 or 0) and checking that in either case the value of φ is forced to be 1 at the immediately lower depth. Such use of a definite value for ψ, that is not even implicitly contained in the current information state V of the agent, is what we call *virtual information.*

Definition 7 *A* k*-depth information state over a search space* Λ *is a valuation* V *of* Λ *that agrees with the informational tables and is closed under the forcing relation* $\Vdash_k$*.*

Let S_k^{Λ} be the set of all k-depth information states over Λ.

Definition 8 *We say that* φ *is a* k*-depth consequence of* Γ*, and write* $\Gamma \vDash_k \varphi$ *if*

$$(\forall V \in S_k^{(\Gamma \cup \{\varphi\})^*})\, V(\Gamma) = 1 \Longrightarrow V(\varphi) = 1.$$

We also say that Γ *is* k*-depth inconsistent, and write* $\Gamma \vDash_k$ *, if there no* $V \in S_k^{(\Gamma \cup \{\varphi\})^*}$ *such that* $V(\Gamma) = 1$*.*

It can also be shown that $\Gamma \vDash_k \varphi$ if and only if there is a finite sequence $\psi_1, \ldots, \psi_n$ such that $\psi_n = \varphi$ and for every element ψ_i of the sequence, either (i) ψ_i is a formula in Γ or (ii) $V \Vdash_k \psi_i$ for all $V \in S_0^{(\Gamma \cup \{\varphi\})^*}$ such that $V\{\psi_1, \ldots, \psi_{i-1}\} = 1$.

Unlike $\vDash_0$, $\vDash_k$ is not a Tarskian consequence relation, but gets very close to being such, for $\vDash_k$ satisfies reflexivity, monotonicity, substitution invariance and the following restricted version of transitivity in which the "cut formula" is required to belong to the search space defined by the deduction problem under consideration.

$$(\forall \psi \in (\Gamma \cup \{\varphi\})^*)\ \Gamma \vDash_k \psi \text{ and } \Delta, \psi \vDash_k \varphi \Longrightarrow \Gamma, \Delta \vDash_k \varphi. \qquad \text{(Bounded Transitivity)}$$

In (D'Agostino et al., 2013) it is shown that $\vDash_k$ is tractable for every fixed k.

Theorem 9 *Whether or not $\Gamma \vDash_k \varphi$ (Γ is k-depth inconsistent) can be decided in time $O(n^{2k+2})$, where n is the total number of occurrences of symbols in $\Gamma \cup \{\varphi\}(\Gamma)$.*

Observe that, by definition, if $\Gamma \vDash_j \varphi$ (Γ is j-depth inconsistent), then $\Gamma \vDash_k \varphi$ (Γ is k-depth inconsistent) for every $k > j$. Classical proposition logic is the limit of the sequence of the depth-bounded consequence relations $\vDash_k$ as $k \to \infty$.

A proof system for each of the k-depth approximations is obtained by adding to the introducton and elimination rules for $\vDash_0$ a single structural rule that reflects the use of virtual information in Definition 6, and bounding the depth at which nested applications of this rule are allowed (see (D'Agostino et al., 2013; D'Agostino, 2015) for the details and a discussion of related work).

5. Towards a prescriptive theory of Bayesian rationality

Let us briefly recap. By framing probability logically we are able to locate the source of a number of important criticisms which are commonly held up against Bayesian rationality in classical logic. The theory of Depth-Bounded Boolean logics meets some of those objections, and gives us an informational semantics leading to a hierarchy of tractable approximations of classical logic. The logical axiomatisation of probability recalled above naturally suggests to investigate which notion of rational belief is yielded once $\models$ is replaced with $\models_k$ in PL1-PL2 above.

This gives us a natural desideratum, namely to construct a family of rational belief measures B_i from L to $[0, 1]$, $i \in \mathbb{N}$ acting as the analogues of probability functions on Depth-bounded logics. Since DBLs coincide, in the limit, with classical propositional logic, our desideratum is then the construction of a hierarchy of belief measures $B_0, \ldots, B_k \ldots$ which asymptotically coincides with probability, i.e. such that for all sentences θ, $B_\infty(\theta) = P(\theta)$.

Each element in the resulting hierarchy would then be a natural candidate to providing a logically rigorous account of a *prescriptive* model of rational belief, in the sense of Bell et al. (1988): every agent whose deductive capabilities are bounded by $\models_k$ must quantify, on pain of irrationality, uncertainty according to B_k.

There is an obvious link between the interpretation of disjunction given by the non-deterministic informational semantics discussed in Section 4.1 and the behaviour of this logical connective in quantum logic. As is well-known, in quantum logic a proposition θ can be represented as a closed subspace M_θ of the Hilbert space $\mathcal{H}$ under consideration. The disjunction $\varphi \vee \psi$ is not represented by the union of M_φ and M_ψ, for in general the union of two closed subspaces is not a closed subspace, but by $M_\varphi \sqcup M_\psi$, i.e. as the smallest closed subspace including both M_φ and M_ψ. So, as is the case for the informational interpretation of disjunction given by the non-deterministic semantics discussed above, a disjunction $\varphi \vee \psi$ in quantum logic may be true even if neither of the disjuncts are true, since $M_\varphi \sqcup M_\psi$ may contain vectors that are not contained in $M_\varphi \cup M_\psi$. On this point see

(Aerts, 2000) and (Dalla Chiara et al, 2004). The negative part of the analogy concerns the behaviour of conjunction which in quantum logic is interpreted as $M_\varphi \cap M_\psi$, so that if a conjunction is false, at least one of the two conjuncts must be false, which departs from the informational intepretation of this operator given by our non-deterministic table. We also point out that this connection between the non-deterministic semantics of Depth-bounded Boolean Logics and the semantics of Quantum Logic opens to a natural parallel between our desideratum and quantum probabilities. This is reinforced by recent experimental findings in the cognitive sciences (Pothos and Busemeyer, 2013; Oaksford, 2014) suggesting that some features of Bayesian quantum probability (Pitowsky, 2003) provide accurate descriptions of experimental subjects.

The key step towards achieving our goal will be of course to define the sense in which we take any B_k to be a *rational* belief measure. The task, as it can be easily figured out, is far from trivial. Though encouraging, our preliminary results suggest that much work is still to be done in this direction. At the same time they suggest that the consequences of such a fully-fledged framework will be far reaching, as it will provide significant steps towards identifying norms of rationality *realistic agents can abide to*.

Aknowledgements. The authors would like to thank the two referees for the careful reading, improvement suggestions and encouraging remarks.

References

D. Aerts, E. D'Onts, and L. Gabora. Why the disjunction in quantum logic is not classical. *Foundations of Physics*, 30:1473–1480, 2000.

T. Augustin, F. P. A. Coolen, G. de Cooman, and M. C. M. Troffaes, editors. *Introduction to Imprecise Probabilities*. Wiley, 2014.

A. Avron and A. Zamansky. Non-deterministic semantics for logical systems. In D.M. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume 16, pages 227–304. Springer Verlag, 2nd edition, 2011.

D.E. Bell, H. Raiffa, and A. Tversky. *Decision making: Descriptive, normative, and prescriptive interactions*. Cambridge University Press, 1988.

J.M. Crawford and D.W. Etherington. A non-deterministic semantics for tractable inference. In *AAAI/IAAI*, pages 286–291, 1998.

M.L. Dalla Chiara, R. Giuntini and R. Greechie. *Reasoning in Quantum Theory*. Springer-Science+Business Media, B.V. 2004.

M. D'Agostino. An informational view of classical logic. *Theoretical Computer Science*, 606:79–97, 2015. doi: http://dx.doi.org/10.1016/j.tcs.2015.06.057.

M. D'Agostino, M. Finger, and D. Gabbay. Semantics and proof-theory of depth bounded Boolean logics. *Theoretical Computer Science*, 480:43–68, 2013.

B. de Finetti. *Theory of Probability, Vol 1.* John Wiley and Sons, 1974.

Thierry Denoeux. 40 years of DempsterShafer theory *International Journal of Approximate Reasoning* 79- 1-6, 2016.

D. Ellsberg. Risk, Ambiguity and the Savage Axioms. *The Quarterly Journal of Economics*, 75(4):643–669, 1961.

T. Flaminio, L. Godo, and H. Hosni. On the logical structure of de Finetti's notion of event. *Journal of Applied Logic*, 12(3):279–301, 2014.

T. Flaminio, L. Godo, and H. Hosni. Coherence in the aggregate: A betting method for belief functions on many-valued events. *International Journal of Approximate Reasoning*, 58:71–86, 2015.

I. Gilboa. *Theory of Decision under Uncertainty.* Cambridge University Press, 2009.

I. Gilboa and M. Marinacci. Ambiguity and the Bayesian Paradigm. In D. Acemoglu, M. Arellano, and E. Dekel, editors, *Advances in Economics and Econometrics: Theory and Applications, Tenth World Congress of the Econometric Society.* Cambridge University Press, 2013.

I. Gilboa, A. Postlewaite, and D. Schmeidler. Rationality of belief or: Why Savage's axioms are neither necessary nor sufficient for rationality. *Synthese*, 187(1):11–31, 2012.

J.M Keynes. *Treatease on Probability.* Harper & Row, 1921.

F.H. Knight. *Risk, uncertainty and profit.* Beard Books Inc, 1921.

D. V. Lindley. *Understanding uncertainty.* John Wiley and Sons, 2006.

M. Oaksford. Normativity, interpretation, and Bayesian models. *Frontiers in Psychology*, 5, 15, 2014.

J.B. Paris. *The uncertain reasoner's companion: A mathematical perspective.* Cambridge University Press, 1994.

G. Parmigiani and L. Inoue. *Decision Theory: Principles and Approaches.* Wiley, 2009.

I. Pitowsky. Betting on the outcomes of measurements: A Bayesian theory of quantum probability. *Studies in History and Philosophy of Science Part B - Studies in History and Philosophy of Modern Physics*, 34(3), 395414, 2003.

E. M. Pothos, and J. R. Busemeyer. Can quantum probability provide a new direction for cognitive modeling? *Behavioral and Brain Sciences*, 36(3), 255274,2013.

W.V.O. Quine. *The Roots of Reference.* Open Court, 1973.

L.J. Savage. Difficulties in the theory of personal probability. *Philosophy of Science*, 34(4): 305–310, 1967.

L.J. Savage. *The Foundations of Statistics.* Dover, 2nd edition, 1972.

D. Schmeidler. Subjective Probability and Expected Utility without Additivity. *Econometrica*, 57(3):571–587, 1989.

G. Shafer. *A mathematical theory of evidence*. Princeton University Press, 1976.