# TOWARDS A BAYESIAN THEORY OF SECOND-ORDER UNCERTAINTY: LESSONS FROM NON-STANDARD LOGICS

HYKEL HOSNI

Abstract. Second-order uncertainty, also known as model uncertainty and Knightian uncertainty, arises when decision-makers can (partly) model the parameters of their decision problems. It is widely believed that subjective probability, and more generally Bayesian theory, are ill-suited to represent a number of interesting second-order uncertainty features, especially "ignorance" and "ambiguity". This failure is sometimes taken as an argument for the rejection of the whole Bayesian approach, triggering a Bayes vs anti-Bayes debate which is in many ways analogous to what the *classical vs non-classical* debate used to be in logic. This paper attempts to unfold this analogy and suggests that the development of non-standard logics offers very useful lessons on the contextualisation of justified norms of rationality. By putting those lessons to work I will flesh out an epistemological framework suitable for *extending* the expressive power of standard Bayesian norms of rationality to second-order uncertainty in a way which is both formally and foundationally conservative.

Contents



*Date*: September 24, 2012.

## 1. An imaginary opinion poll

*What is the right logic*? If logicians worldwide were polled I would imagine the vast majority of answers falling into two categories:

**Type 1:** "There is no right logic – the question is ill-posed".
**Type 2:** "Logic $L$ is obviously the right logic".

As to the distribution of the answers, I would expect mostly Type 1 responses, a non-negligible proportion of which possibly adding "You should rather be asking *Is logic L adequate for context C?* Those are the questions that logicians today find well-posed and indeed very much worth attacking."

The same opinion poll thirty years ago would have very probably resulted in the opposite outcome, with respondents readily picking their favourite $L$.

Needless to say some (substantial) disagreement is likely to emerge on my projections here. Be it as it may, it's hard to question that over the past few decades the logician's way of thinking about logic has changed. It is too early to say, but it is quite plausible that our way of thinking about logic has changed (and keeps changing) as a consequence of our way of doing logic, especially in connection to computer science, artificial intelligence and cognitive science. Those research areas added a whole new stock of pressing problems which largely contributed to revamping a number of long-standing ones arising from philosophy and linguistics. Those questions, in turn, could be very sharply defined and attacked from the vantage point of an unprecedented mathematical understanding of classical logic, leading to a virtuous circle which provided great momentum for the development of non-standard logics.

Let us now consider a slight variation on our imaginary poll, so that the question now becomes *What is the right measure of uncertainty?* Now I suppose the vast majority of uncertain reasoners would give responses of Type 2, with many adding "why are you asking at all?". In many ways, the current debate on the foundations of uncertain reasoning mirrors what the debate about classical vs non-classical logics used to look like during the first half of the nineteenth century, a discussion about who is right and who is hopelessly wrong.

The purpose of this paper is to suggest that uncertain reasoning can greatly benefit from undergoing a change in perspective similar to the one which provided such favourable a context for the development of non-standard logics, some of which today compare in formal depth and philosophical interest to "mathematical logic" as epitomised by Barwise and Keisler (1977). Just as in non-standard logics initial formal development led to a gradual change in foundational perspective, which in turn gave rise to enough formal advance to change our way of thinking about logic, I will make a suggestion to the effect that a similar virtuous circle may pave the way for increasingly more expressive formal models of rationality. This will throw new light on the Bayes vs non-Bayes contrast which continues to be the focus of much theoretical work in uncertain reasoning, but rarely provides enough insight to facilitate the much-needed formal advance that the broad area of rational reasoning and decision under uncertainty so very urgently needs.

The paper is organised as follows. Section 2 presents very briefly the Type 2 position in uncertain reasoning. Section 3 focusses on one specific thread in the development of non-standard logics, which I will refer to as the contextualisation of reasoning. This will provide the background against which I will draw an analogy and a contrast between logic and uncertain reasoning. The resulting suggestion will be for uncertain reasoners to abandon positions of Type 2. The second part of the paper is devoted to outlining an epistemological framework which may enable this transition. Section 4 fleshes out the main rationale of such a framework whilst Section 5 informally illustrates its applicability to produce conservative extensions of standard Bayesian theory which capture interesting aspects of second-order uncertainty. Section 6 concludes.

## 2. A snapshot of the foundations of uncertain reasoning

Bayesian theory abounds with Type 2 positions. Consider for instance de Finetti

> Bayesian standpoint is noways [sic] one among many possible theories but it is an almost self-evident truth, simply and univocally relying on the indisputable coherence rules for probability. (de Finetti, 1973, p.468)

Similar uncautiousness is easily found in connection to Bayesian methods, as opposed to Bayesian theory.[1] Jaynes for instance, put it as follows:

> The superiority of Bayesian methods is now a thoroughly demonstrated fact in a hundred different areas. One can argue with philosophy; it is not so easy to argue with a computer printout, which says to us: "Independently of all your

[1] For a recent appraisal of the distinction see (Gelman, 2011)

> philosophy, here are the facts of actual performance".(Jaynes, 2003, p. xxii)

Interestingly enough, de Finetti and Jaynes endorsed radically different forms of Bayesianism, which prompts a clarification of the use I will make of the adjective Bayesian for the purposes of this work.

Bayesian theory justifies two interdependent norms of individual rationality, namely (i) degrees of belief should be probabilities and (ii) decisions should maximise the subjective expected utility the outcomes. Anti-Bayesian approaches reject the adequacy of such norms and put forward alternative research programmes, which are just too many to mention here. The approach exemplified by Gilboa et al. (2011) shares with Bayesian theory the concern for norms of rationality, but departs from it in that it presupposes the existence of uncertainties which cannot be quantified probabilistically. This line of criticism, which goes back to both Keynes and Knight, albeit in rather distinct forms, is effectively summed up in the opening lines of 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 subjective probability is unfit as a norm of rational belief. 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 grounds that, say I have never observed signs of "unfairness" in my coin, whilst I just don't know anything about the stranger's coin. This familiar argument against the completeness of the revealed belief approach[2] goes hand in hand with a radical challenge to probability as a norm of rational belief:

[2]The so-called Ellsberg paradox, of which the coin tossing problem is the simplest example, appears as far back as in (Keynes, 1921; Knight, 1921). This is often acknowledged in the literature by referring to probabilistically unquantifiable belief as "Knightian uncertainty". The synonym "ambiguity" is due to Ellsberg (1961). Much of the recent revival of the interest in "Knightian decision theory" owes to Bewley (2002).

> The main difficulty with [. . . ] the entire Bayesian approach is in my mind the following: for many problems of interest there is no sufficient information based on which one can define probabilities. Referring to probabilities as subjective rather than objective is another symptom of the problem, not a solution thereof. It is a symptom because, were one capable of reasoning one's way to probabilistic assessments, one could have also convinced others of that reasoning and result in a more objective notion of probability. Subjective probabilities are not a solution to the problem: subjectivity [. . . ] does not give us a reason to choose one probability over another. (Gilboa, 2009, pp.130-1)

The Bayesian failure to represent ignorance leads to a plea for rational modesty:

> It is sometimes more rational to admit that one does not have sufficient information for probabilistic beliefs than to pretend that one does. (Gilboa et al., 2011)

Similar lines of argument had been put forward by Shafer (1986). Colyvan (2008) argues against the suitability of probability by suggesting that the so-called *probabilistic excluded middle* misrepresents ignorance to such an extent that it can hardly be taken as a sound principle of uncertain reasoning.[3]

## 3. An analogy and a contrast

This sort of debate is reminiscent, in style and spirit, of the heated disputations[4] over "the right logic", of which the classical vs intuitionistic logic one is certainly the best-known case. The purpose of this Section is to put forward a suggestion as to why the success of non-standard logics, and hence of Type 1 mentality, is intimately connected with the progressive loss of appeal of quarrels of that sort.

3.1. **Non-classical and non-standard logics: a terminological remark.** In everyday usage the adjective "classical" bears at least two connotations. One refers to something which is traditional. The other refers, in stricter adherence to the Latin etymology, to the class of the "best", the

[3]Roughly speaking, what is objected is that a rational agent must assign probability 1 to any tautology of the form $\theta \vee \neg\theta$, even when the agent knows nothing about $\theta$.

[4]See e.g. (Heyting, 1956).

privileged and the important. The expression "classical logic" retains both meanings. Yet I take the current state of research in logic, broadly construed, to support the claim to the effect that "classical logic" is somewhat a misnomer with respect to both connotations.

The history of logic clarifies that "classical logic" is certainly not the invariant to be found in the millennial development of the subject. On the other hand, the mathematical depth of the results and the wealth of applications, from philosophy to computer science to game theory, of whole families of modal, non-monotonic and many-valued logics, stand as obvious evidence to the effect that classical logic is certainly not the only citizen of the "first-class". In the light of this, *non-standard logics* appears to be far better a terminology for those logics which arise by making more imaginative uses of the concepts and techniques of classical logic, which is therefore not rejected as ill-founded.

This leads to a general pattern in the construction of non-standard logics. Certain reasoning schemes are not adequately modelled by the classical notion of consequence so that the goal is to define more *realistic* notions of "follows". In this attempt logicians are often guided by contextually clear modelling needs, which are translated into suitable restrictions on the applicability of classically valid principles of inference. A paramount example of this *extension by restriction* pattern is provided by the logic of defeasible reasoning.

3.2. **Context in action I: Defeasible reasoning.** Consider the naive logical modelling setting in which an agent's reasoning is identified with a consequence relation.[5] The repeated use of a consequence relation is interpreted as determining those formulae (or equivalently, for present purposes, sentences) which the agent is forced to accept –on pain of violating the underlying norms of rationality– given that the agent is accepting a (possibly empty) set of formulae which are interpreted as the premisses of the agent's reasoning.

"Accepting a formula" can be formalised in a number of essentially equivalent ways in classical logic. Say that an agent accepts the formula $\theta$ if $v(\theta) = 1$, where $v : \mathcal{L} \to \{0, 1\}$ is the usual notion of valuation which extends uniquely to the set of sentences $\mathcal{SL}$ recursively built from propositional

[5]As the content of this section is purely heuristic, I will not burden the reading with otherwise unnecessary definitions. On the general questions of providing rigorous characterisations of logical systems and context of reasoning, Gabbay (1995) is a slightly dated yet still very valuable reference.

language $\mathcal{L}$. As usual a set $\Gamma \subseteq \mathcal{SL}$ has a model if $v(\gamma) = 1$ for all $\gamma \in \Gamma$. This minimal logical setting leads to the Tarskian analysis of consequence:

> $\theta$ is a logical consequence of a set of assumptions $\Gamma$, written $\Gamma \models \theta$, if every model of $\Gamma$ is also a model of $\theta$.

Under the present naive modelling, $\Gamma \models \theta$ can be interpreted as $\Gamma$ giving the agent reasons to accept $\theta$. Since classical logic is monotonic, an agent whose reasoning is captured by the relation $\models$ must reason monotonically, that is it must satisfy

(MON) $$\frac{\Gamma \models \theta}{\Gamma, \psi \models \theta},$$

where $\psi$ is an arbitrary formula. (MON) can thus be interpreted as saying that $\psi$ does not provide relevant reasons either for or against accepting $\theta$ beyond those already available to the agent who accepts $\Gamma$. Since $\psi$ is arbitrary[6] we can say something rather stronger, namely that *anything* beyond $\Gamma$ is actually irrelevant to the acceptance of $\theta$. Thus $\models$ captures a notion of acceptance which is based on having *sufficient reasons*.

There are many situations in which sufficient reasons are never available to the rational agent, no matter how accurate or thoughtful they may turn out to be. Any minimally realistic scenario will feature exogenous, dynamic aspects which make reasoning according to sufficient reasons largely inapplicable. Indeed, with the notable exception of formal deduction, there hardly seems to be a "real-world" context for which *unconstrained* monotonicity can be taken as generally adequate principle of rational reasoning. *This* goes some way towards vindicating the use of the term "classical logic" to denote the logic of formal deduction.

Virtually all the development of defeasible logics, from the early syntactic and modal approaches, to abstract theory of non-monotonic consequence relations, thus focussed on modelling inference based on reasons which are responsive to potentially invalidating refinements of the agent's currently held information. By the end of the 1980s, suitable constraints on the applicability of (MON) were identified –notably *cautious* and *rational* monotonicity[7]– which led to the model-theoretic analysis of defeasible logic as Tarskian consequence nuanced by some suitable minimisation:

> $\theta$ is a defeasible logical consequence of a set of assumptions $\Gamma$ if all minimal models of $\Gamma$ are also models of $\theta$.

[6] The restriction to single sentences is clearly immaterial here.

[7] The reader who is not familiar with the details may wish to consult Makinson (1994, 2005).

The resulting supraclassical consequence relations extend the expressive power of classical reasoning from "acceptance based on sufficient reasons" to "acceptance based on good reasons" in way which is formally and foundationally conservative.[8]

### 3.3. Context in action II: Rationality, coherence and consistency.

The contextual analogy offers also a way of contrasting uncertain reasoning and non-standard logics on the relation between (in)consistency and (ir)rationality. Whilst rational norms of uncertain reasoning appeal to a variety of distinct and often mutually incompatible intuitions about what "rational" actually means,[9] the formal notion of consistency offers much less room for controversy: A set of sentences $\Gamma$ is *inconsistent* if it logically implies any sentence $\alpha \in \mathcal{SL}$.[10]

According to the naive agent-representation of consequence relations recalled above, the standard relation between logical inconsistency and rationality can be stated as follows: Rational agents cannot accept inconsistent sets of sentences. This view lies at the heart of the belief change approach of which AGM has been a vastly successful example. Inconsistency triggers revision exactly because logical inconsistencies are normatively incompatible with the epistemic state of a rational agent.

Yet, it is easy to imagine situations in which an agent can rationally find themselves accepting an inconsistent set of beliefs. Since the definition of inconsistency is essentially unique, it is again the context of reasoning which must be invoked to justify the rationality –i.e. normative adequacy– of entertaining inconsistent beliefs which however need not trigger a revision.

Interestingly enough one such context was also noted by David Makinson in his *Preface Paradox*. Makinson considers a situation in which the author of a monograph who believes that each statement in the book is true is apologising for the possible persistence of errors. According to Makinson's analysis of the problem, the author

[8]No doubt some meta-mathematical properties are lost in this extension! Supraclassical consequence relations, for instance, need not be closed under substitution (Makinson, 2005).

[9]Daston (1988) emphasises how the virtually unanimous agreement on the "expectations of reasonable men" –one of the trademark of the Enlightment– played an important role in the construction of the theory of probability. The objective Bayesian approach of Paris (1994) takes for granted that our intuitions about rationality are largely intersubjective, an idea which is further developed in (Hosni and Paris, 2005)

[10]A set of sentences is consistent if it is not inconsistent.

> is being rational even though inconsistent. More than this: he is being rational even though he believes each of a certain collection of statements, which he knows are logically incompatible [...] This appears to present a living and everyday example of a situation [in] which it is [...] rational to hold incompatible beliefs. (Makinson, 1965)

The Preface Paradox constitutes an early exploration of the idea that under suitable circumstances it may be rational to violate consistency. This rather subversive intuition paved the way for a number of subsequent investigations which contributed to making inconsistency respectable to logicians.[11] Hence, the relaxation of the tight connection between irrationality and logical inconsistency.

In contrast to this, the wider domain of uncertain reasoning is currently dotted with disagreement on the very definition of consistency, i.e. what should be taken as the formal counterpart of a rational norm of belief. This essentially accounts for the Type 2 majority which I suggest would dominate the imaginary poll of Section 1. The reason is quite simply that there is currently no space for a Preface Paradox-like manouvre in Bayesian uncertain reasoning: Since choice behaviour reveals the epistemic state of an agent, there is clearly no distinction between displaying irrational behaviour and entertaining an inconsistent epistemic state.

3.4. **Methodological lessons.** It useful to single out more explicitly the key methodological lessons that emerge from the informal analysis of the development of non-standard logics sketched above. Since reasoning doesn't occur in a vacuum, logical modelling of rational agents is best relativised to an intended domain of application, or in case our interest lies in "pure" modelling, to some well-specified reasoning *context* which pins down the salient properties of the domain of interest.

I have mentioned some of the motivations for extending Tarskian consequence to capture defeasible reasoning, but similar arguments are easily put forward for a large number of non-standard logics which entered the stage with the dynamic, practical and many-valued turns in logic.[12] Those are just a handful of very familiar examples of how new ways of *doing* logic contributed essentially to changing our way of *thinking about logic*. In turn, this clearly affected our way of doing logic, thus igniting a "doing-thinking-doing'

[11] The expression is borrowed from (Gabbay and Hunter, 1991). See also (Paris, 2004). Carnielli et al. (2007) offer a comprehensive survey of the field.

[12] See, e.g. Gabbay and Woods (2005, 2007); van Ditmarsch et al. (2007). See also Moss (2005).

virtuous circle which largely accounts for the intuition that respondents of Type 1 in our imaginary logic poll outnumber those of Type 2. This virtuous circle can in fact be credited for diluting the excitement over the disputes about "the right logic". Take, as an illustration, intuitionistic logic. The so-called Brower-Heyting-Kolmogorov semantics is solidly interpreted within $S4$, a modal logic which conservatively extends classical logic.[13]

Thanks to the virtuous circle, logicians who are sensitive to the pressing need for modelling rational yet defeasible reasoning, may dispense with (unconstrained) monotonicity without plunging into foundational discomfort. Similarly for many other cornerstones of classical logic, should those turn out to be inappropriate for the specific context of interest. Even consistency, as recalled above in connection to the Preface Paradox, might turn out to be undesirable for a logic which aims at modelling rational norms of reasoning. And we can go on restricting and relaxing even more "entrenched" principles of classical logic, such as compositionality, which is clearly incompatible with a number of features of interest in (qualitative and quantitative) uncertain reasoning.[14]

This delivers a rather unequivocal message. Classical logic provides the starting point for a number of non-standard logics which aim at capturing more realistic aspects of rational reasoning than those formalised by the classical notion of consequence. Once a particular context of reasoning is sufficiently well defined so as to give rise to sufficiently sharp modelling intuitions, we can set out to construct logics which address what then appear as clear limitations of classical logic.[15] Increasingly more realistic contexts of reasoning demand that we use classical logic in increasingly more imaginative ways, yet none of the above-mentioned restrictions and relaxations of classical logic is motivated by the desire to replace classical logic because it is wrong.

[13]The first representation theorem to this effect goes all the way back to 1933 when Gödel introduced what we now call the Gödel translation and then proved that the algebra of open elements of every modal algebra for S4 is a Heyting algebra and, conversely, every Heyting algebra is isomorphic to the algebra of open elements of a suitable algebra for S4 (see, e.g. Blackburn et al., 2007, especially Ch7.9).

[14]Arieli and Zamansky (2009) illustrate the applicability of non-deterministic matrices in modelling non-deterministic structures, where uncertainty is taken to be an intrinsic feature of the world, in addition to being a subjective epistemic state of the reasoning agents. Adams (2005) gives a feel for the *taboo* of abandoning compositionality in connection with Lewis's Triviality results.

[15]This process is mostly, but not necessarily, application-driven. As Banach is often reported to have said, "Good mathematicians see analogies between theorems; great mathematicians see analogies between analogies" (as quoted by Jaynes who quotes from Ulam).

This *contextualisation* strategy lends a precious methodological lesson to uncertain reasoning. Reasoning by analogy naturally suggests the following question: What is the context in which standard Bayesian norms are justified? I will put forward an answer in Section 4 (in the limited case of belief norms) and it can hardly be surprising that such a context will be rather narrow. Hence the comprehensible dissatisfaction with standard Bayesian theory to which –contextualisation suggests– we should *not* reply by denying the relevance of Bayesian theory to normative models of rationality.[16] This would be like saying that logic is irrelevant to rational reasoning because classical logic is inadequate for air-traffic control. The success story of non-standard logics recommends a different, more foundationally conservative reaction: Identify the contexts of interest, and then work hard to extend the expressive power of standard Bayesian norms to capture those specific aspects of rational reasoning and decision under uncertainty. If we set ourselves free from the ill-founded quest for the right measure of uncertainty and start looking for what measures are justified in a certain context, we can hope to ignite a doing-thinking-doing virtuous circle similar to the one which proved so useful in the development of non-standard logics.

## 4. First-order uncertainty and classical Bayesianism

Uncertain reasoning is best understood in terms of a fact and an assumption. The fact, roughly speaking, is that whenever we face a (non-trivial) choice problem we find ourselves in an epistemic state of uncertainty concerning the outcome of our choices. The assumption, on the other hand, is that we can make sense of such an epistemic state. Jacob Bernoulli's *Ars Conjectandi* (1713) is widely recognised as putting forward the first convincing framework based on the idea that making sense of uncertainty goes hand in hand with *measuring* it. As briefly outlined in Section 1 however, the intervening three centuries have brought a rather limited consensus, among the wider community of uncertain reasoners, as to the precise details of how uncertainty should be measured. Our running logical analogy suggests not only that consensus may, in fact, never be achieved, but also and perhaps more radically, that we should stop looking for it and start articulating in some detail the lack of it instead.

This clearly requires an epistemological framework, to which I now turn.

4.1. **Choice roots for Bayesian theory.** As a starting point I will assume that making sense of our epistemic state of uncertainty –measuring

[16] This is precisely what the title of (Gilboa et al., 2011) provocatively recommends.

uncertainty– matters to us insofar as we are faced with suitably defined choice problems. Whilst such a "behavioural" perspective might not account for the whole story, it certainly constitutes a very important part of it as illustrated by the recent analysis of *objective* Bayesian epistemology put forward by Williamson (2010).

Williamson singles out three characterising norms of rational belief which jointly pin down the bulk of the objective Bayesian framework. The *Probability Norm* demands that rational agents' degrees of belief should conform to the laws of probability. The *Calibration Norm* requires that the subjective probabilities licensed by the Probability Norm should be further constrained by known frequencies or, if these exist at all, single-case physical probabilities.[17] Finally, the *Equivocation Norm* further refines the choice of subjective probabilities by excluding extreme probability values unless these are being prescribed by the previous norms, and subject to this requirement, it constrains probabilities to be otherwise minimally prejudiced, or equivalently, maximally equivocal.[18] Objective Bayesianism is the epistemological framework in which Probability, Calibration and Equivocation are endorsed as *norms* of rational belief.

Williamson points out that Probability, Calibration and Equivocation are individually justified[19] by appealing to (formal) variations of essentially the same argument which involves the minimisation of a certain loss function. The gist of the argument can be described as follows. Provided that an individual is faced with a suitably defined choice problem, each of the three above norms can be justified by showing that contravening them would increase the agent's expectation of incurring a loss, possibly in the long run. The basic instantiation of this line of argument will be recalled in some detail in Section 4.4 below, where the construction of de Finetti's betting problem is seen to force a choice of betting odds which prevents the bookmaker from facing a sure loss.

This account of objective Bayesian epistemology lends itself to two considerations. First, the framing of the choice problem which motivates the quantification of uncertainty is fundamental in the definition and justification of the candidate norms of rational belief. This makes a strong case for connecting the context of reasoning to the specific description of the

[17] A familiar instantiation of this Norm is Lewis's *Principal Principle*.

[18] The best-known instantiation of the Equivocation norm is the Maximum entropy principle, which has been extensively discussed over the past three decades, often from heterogeneous points of view. For two comprehensive presentations, see Paris (1994); Jaynes (2003)

[19] See Chapter 3 of (Williamson, 2010) for a detailed analysis of such justifications.

choice problem, as discussed above in comparison with non-standard logics. Hence the choice problem should be centrally involved in the definition of the formal counterpart to the intuitive notion of "irrationality". Second, and related to the first point, there seems to a very general and very basic principle which gives rise, through distinct instantiations, to distinct norms of rationality. I propose to render this principle as follows:

> **Choice Norm**: A rational agent must not choose inadmissible alternatives.

The Choice[20] Norm captures a central aspect of Bayesian epistemology by making explicit how the subjective component of individual uncertainty is intertwined with the objective features of the underlying choice problem. This connects choice, belief and decision in a way which is typical of Bayesian theory. De Finetti, for instance, put it as follows:

> A decision must [...] be based on probabilities: i.e. the posterior probabilities as evaluated on the basis of all information so far available. This is the main point to note. In order to make decisions, we first require a statistical theory which provides conclusions in the form of posterior probabilities. The Bayesian approach does this: other approaches explicitly refuse to do this. (de Finetti, 1974, p.252)

This virtuous circle connecting choice, belief and decision accounts for the ubiquitous synergies of the broad concepts of "rationality" and "uncertainty" in statistics, epistemology, economics and related fields. Rukhin (1995), for example, puts it as follows:

> How should one choose a decision rule whose performance depends upon the unknown state of Nature? Since there is no uniquely recognized optimality principle that would provide a complete ordering of all statistical decision rules, this question is probably unanswerable when posed with such generality. However, it seems clear which procedures should not

[20] Alternative denominations might have included "Dominance", "Pareto" or even of course "Admissibility". As they all have rather specific connotations in distinct areas of the uncertain reasoning literature, from statistics to decision and game theory to social choice theory, it appears that "Choice Norm" sits more comfortably at the desired level of generality whilst avoiding potential confusion.

> be used – the inadmissible ones which can be improved upon no matter what the unknown state of Nature.[21]

The Choice Norm can thus be seen as a good candidate to address the contrast outlined in Section 3.2 above and provide unity in the formalisation of the intuitive notion of "irrationality". To see this note that the Choice Norm is best seen as an attempt to define rationality in terms of *avoiding blatantly irrational* behaviour which is arguably more readily identified than its positive counterpart.[22] This negative characterisation of rationality is ubiquitous in Bayesian epistemology and is tightly connected to the idea of rationality as maximisation. Wald's seminal "idea of associating a loss with an incorrect decision" (de Finetti, 1975, p. 253), for instance, paved the way to the analysis of subjective expected utility as the standard norm of rational decision.

The Choice Norm captures a notion of rationality which is general enough to be applicable even when no uncertainty enters directly the picture. In fact it builds on a very weak, yet non-empty, characterisation of "purposeful behaviour".[23] It is non-empty because the occurrences of "rational" and "inadmissible" appearing in the above phrasing of the Choice Norm have distinct epistemological statuses. Whilst the former is intended intuitively, the latter is taken in a precise, technical sense. So, if "rationality" means anything at all, it cannot be rational to behave in a way which blatantly contradicts the purpose of our own behaviour. It is the ordinary parlance meaning of "irrational" that is being used here, i.e. "stupid", "against commonsense", "illogical", etc. On the other hand the term "inadmissible" which occurs in the Choice Norm aims at capturing what we intuitively regard as being self-defeating –as opposed to purposeful behaviour– in the context of a formally specified choice problem.

What is perhaps the simplest example goes as follows. Let $X$ be a set of feasible alternatives, $R_i \subseteq X^2$ a binary relation such that $xR_iy$ interpreted as "$i$ doesn't prefer $y$ to $x$", and $\emptyset \neq C_i(X) \subseteq X$ be $i$'s *choice set* from $X$ (i.e. the non-empty subset of feasible alternatives selected by $i$). We say

[21]See (Levi, 1986; Bossert and Suzumura, 2011) for an appraisal of similar ideas in epistemology and economics, respectively.

[22]In this respect, the concept of rationality seems to be analogous with that of democracy: there is usually more disagreement on what democracy should be than on what counts as a violation of a democratic society.

[23]I am adapting the terminology from Bossert and Suzumura (2010) who use it in connection to their notion of *Suzumura consistency*, arguably the weakest requirement in the formal theory of rational choice.

that $i$'s choices are *inadmissible* if and only if there exists $y \in X$ such that $yR_iz$, for all $z \in X$ but $y \notin C_i(X)$.

Going back to the comparison of the previous Section, "irrationality" is to "inadmissibility" what "incoherence" is to logical "inconsistency". This analogy suggests that a stupid choice is clearly *always* inadmissible. To see that an inadmissible choice is intuitively irrational, suppose $i$'s preferences are modelled by the binary relation $R_i$. Under this assumption, an inadmissible choice captures the idea of self-inconsistency –blatant irrationality– so that the Choice Norm specifies the conditions under which an informal problem can be modelled as a *choice problem*, i.e. a formalised situation in which inadmissible alternatives cannot be rationally chosen.[24] Put the other way round, if an agent is normatively justified in selecting inadmissible choices, the situation at hand falls short of being a choice problem. But if there is no well-posed choice problem, we have no reason to worry about making irrational choices, a situation in which the problem of quantifying uncertainty need not arise at all.

This suggests that Bayesian theory –in the rendering which I am emphasising here– takes rationality as a function of two arguments: a choice problem and an individual facing it. As a consequence, whether a certain norm of rationality is justified or not, will have to be discussed in relation to the assumptions we make on the choice problem, and those we make on the agent who is confronting it. The remainder of this section is devoted to illustrating how, under this interpretation, standard Bayesianism can be fruitfully seen as the solution to the following problem:

> **First-order uncertainty:** How should a maximally idealised agent behave when facing a maximally abstract choice problem?

Maximum idealisation will be shown to go hand in hand with the requirement that our model be normative. The role of maximal abstraction will be, as in all modelling, essential to set off mathematical formalisation.

4.2. **Idealisation.** Bayesian theory is normative at root, a feature which is abundantly emphasised by its proponents. In a number of early presentations of his ideas, de Finetti likens probability to the "logic of the uncertain", a view which is fully articulated in (de Finetti, 1972, Chapter 2). Savage (1954) refers to his postulates as "logic-like criteria". More recently, the

[24] This is one way of interpreting the axiom of the Independence of irrelevant alternatives, according to which a dominated alternative should not be chosen from any superset of the original set of feasible alternatives Sen (1970).

logical characterisation of "common sense principles" has been the focus of the objective Bayesian approach of Paris and Vencovská (1990); Paris (1994, 1998).[25]

According to the received view, normative models of rational behaviour should not take into account the potential cognitive limitations of the agents which are being modelled. This is usually motivated by the fact that building cognitive limitations into a normative model deprives it from its role in correcting mistakes, arguably one of the key reasons for developing normative models in the first place.[26] So, whilst the received view is not unchallenged (see, e.g. Gabbay and Woods, 2003), I think it offers a very useful starting point and I will therefore largely conform to it. Not completely though, because idealisation is usually thought as a binary, all-or-nothing, property of agents, whereas it is much more natural, in the view I am articulating here, to think of idealisation as coming in degrees.

The idea is as follows. Students recruited by economists for their experiments, certainly qualify as minimally idealised agents, as do all of us individual, real, agents. A Turing machine, with its indefinitely extensible tape, is certainly an example of a maximally idealised agent, one whose memory limitations are of no consequence for the model of computation it defines. Arguably all agents of interest sit somewhere in between the spectrum delimited by those two examples. It is indeed tempting to go on and define an ordering relation on the minimum-maximum abstraction interval, which could be interpreted as "is less subject to cognitive limitations than" and investigate the consequences of this for uncertain reasoning modelling.[27] For present purposes, however, I will restrict the attention to maximally idealised agents only and refer to all other agents as "non-idealised".

4.3. **Abstraction.** All modelling requires abstraction, a fact that certainly contributes to making all models wrong, in Box's notorious dictum. Abstraction is the inevitable price that we must pay to grant ourselves the

[25]The consensus on granting logic a normative status is far from being unanimous, and indeed the question might turn out to be ill-posed. A discussion of this point would take us too far, but the interested reader might wish to consult, among others Gabbay and Woods (2003); van Benthem (2008); Wheeler (2008).

[26]As a well-known story goes, Savage was initially tricked into Allais's "paradox", but once Allais pointed that out, Savage acknowledged his mistake and corrected his answer accordingly. A similar reaction to the descriptive failures of Bayesian theory is condensed in the one-page paper (de Finetti, 1979).

[27]Gabbay and Woods (2005) go some way towards developing the idea of a hierarchy of agents based on a similar relation.

privilege of quantitative thinking[28]. The specific aspect of abstraction that will be of direct interest for present purposes concerns the features of the choice problem which, as recalled above, motivates the need for the quantification of the agent's uncertainty, or as I will simply say from now on, *the choice problem*.

De Finetti's betting problem, to be discussed shortly, and Savage's decision matrix are two very familiar examples of maximally abstract choice problems which intuitively can be associated with the formal and complete description of a "real-world" problem. I am stressing *intuitively* here because a precise characterisation of the abstraction of the choice problem is exactly what the framework under construction aims at achieving.

The key idea is that a maximally abstract choice problem includes all and only the information which is relevant to the decision-maker who is facing the choice problem. Put otherwise, when an agent is facing a maximally abstract choice problem it is *that* problem that the agent is facing, and not some other problem, however related. This is a fundamental yet too often overlooked modelling principle which has been given various names in the uncertain reasoning literature, including *the Watts assumption*[29] by Paris (1994), where the idea is presented as follows:

> The [linear constraints on a probability distribution are] not simply the shadow or description of the expert's knowledge but [...] *(essentially) all the expert's relevant knowledge.* [...] If we make this assumption then our task of giving a value to $Bel(\theta)$ given $K$ is exactly the task that the expert himself carries out. (p.67)[30]

In this spirit, we can naturally think of an agent facing a maximally abstract choice problem as a *decision-maker with no modelling privileges.* The distinction between decision-makers and decision-modellers can be elusive and this goes some way towards explaining why so much foundational confusion arises on this point, some of which underlies the discussion recapped

[28]This need not pertain only to applied mathematics. In a number of widely known mathematical expositions George Polyá insists that abstraction is intrinsic to mathematical reasoning for the solution to hard mathematical problems is often best achieved by solving simpler problems from which the general idea can be extrapolated.

[29]This assumption is clearly related to Carnap's "Principle of Total Evidence" and Keynes' "Bernoulli's maxim". The fact that experimental subject systematically violate this principle motivates the introduction of the "editing phase" in Prospect Theory.

[30]Here $Bel()$ denotes the expert's belief function, $\theta$ is a sentence and $K$ is a finite set of expressions of the form $Bel(\theta_i) = \beta_i$, where all the $\beta_i \in [0, 1]$.

in Section 2. For definiteness, recall Schmeidler's "two coins" example,[31] in which we are told that an agent's preference for tossing their own coin is normatively rational on the grounds that it has never given any sign of being unfair, whereas nothing is known about the stranger's coin. But what is precisely the choice problem here? We might certainly be justified in distrusting the stranger and their coin, but the Watts assumption requires this to be explicitly represented in the choice problem. If it is, no paradoxical situation arises, for the preference for one's own coin is so to speak, tautological. If the extra information concerning one's beliefs about the stranger and their coin are *not* represented in the choice problem, then there is no reason to prefer one's own coin.

More generally, the main source of confusion here lies in the fact that in real life we very often play *both* the role of decision-makers and that of decision-modellers. Take an agent who is about to do their shopping. This real-world problem can be given a maximally abstract representation as a consumer's choice problem (see, e.g. Rubinstein, 2006). When doing our shopping, however, we make some choices as modellers, say whether to consider items without the fair-trade certification as feasible alternatives, and some as decision-makers, e.g preferring pomegranate juice to orange juice. Standard Bayesian theory relies, albeit this point is sometimes only implicit, on the rigid distinction between decision-making and decision-modelling. This is in fact central to the very idea of revealing consistent beliefs and preferences through a formally defined elicitation mechanism. Whenever such an elicitation device is assumed to capture all the relevant features of the real-world problem for which uncertainty needs to be quantified, we can think of the agent whose degrees of belief are being elicited as facing a maximally abstract choice problem.

This allows us to reframe de Finetti's Dutch book argument as the result of instantiating the Choice Norm with a maximally abstract betting problem. The purpose of this reframing is to make explicit the context for which standard Bayesian norms are justified. This will then constitute our starting point for the extension of standard Bayesian norms to second-order uncertainty.

4.4. **The abstraction of de Finetti's betting problem.** In a nutshell[32] de Finetti's betting problem is one in which a bookmaker is asked to write

[31] Which in turn is a variant of a problem known to both Keynes and Knight before it was revamped by Ellsberg.

[32] Since this is one of the most intensely studied aspect of Bayesian epistemology, I will take many details for granted and focus on the specific aspects which are directly relevant to the present discussion. Readers who are not familiar with the argument are

their betting odds for a set of events of interest, otherwise known as a *book*. In this context it is natural to take as blatantly irrational a choice of odds that exposes the bookmaker to potential loss independent of the outcome of the events in the book. Admissibility, which de Finetti calls *coherence*[33] can thus be rigorously characterised as "avoiding sure loss" (otherwise known as a *Dutch book*). As Ramsey independently anticipated and de Finetti proved, a Dutch book is avoided exactly if the betting odds conform to the laws of probability.

The formalisation of admissibility as avoiding sure loss rests on a rather convoluted elicitation framework which has caused much discussion over the past eight decades or so and which certainly justifies the present choice of considering de Finetti's betting problem as maximally abstract. Central to this is a complete contract[34] which is introduced to regulate the exchange of money between gamblers and bookmakers. The contract includes the following clauses:

**Completeness:** The bookmaker's choice is forced for (boolean) combinations of bets and, after the book has been published, the bookmaker is forced to accept all of a potentially infinite number of bets.

**Swapping:** After reading the published book, the gambler bets by paying to the bookmaker a real-valued stake of her choice. Since the gambler can choose *negative* stakes (betting negative money), she can unilaterally impose a payoff-matrix swap to the bookmaker.

**Rigidity:** Stakes involved in the betting problem correspond to actual money (in some currency).

Completeness is justified by de Finetti (1931) on the grounds that it provides the following modelling constraints. Were the bookmaker allowed to refuse selling certain bets, the bookmaker's betting odds could not be claimed to reveal his sincere degrees of belief on the relevant events, and as a consequence, the betting problem would fail its fundamental purpose of connecting a rational agent's degrees of belief to their willingness to bet. As he would retrospectively notice, the betting problem is a "device to force the

urged to consult the original (de Finetti, 1931, 1974). Paris (2001) offers a very general proof whilst Williamson (2010) provides ample background.

[33]The term "consistency" is also frequently used in English translations.

[34]In economics and political science a contract is said to be complete if it contemplates all possible contingencies. Despite being blatantly unrealistic, it is a widely used assumption in those areas (see, e.g. Tirole, 1999).

individual to make conscious choices, releasing him from inertia, preserving him from whim" (de Finetti, 1974, p.76).[35]

In the presence of Completeness, Swapping entails that the bookmaker's degrees of belief should be *fair* betting odds. For suppose the bookmaker were to publish a book with non-zero expectation. Then he could be forced into sure loss by a gambler who put a negative stake on the book. Note that the abstraction leading to fair betting odds is justified only if the agents involved are maximally idealised. This amounts to saying that the bookmaker interacts with gamblers who will exploit any logical possibility of making a Dutch book against him, no matter how computationally demanding this might be. In game theoretic language, publishing fair betting odds constitutes the bookmaker's best response against rational gamblers under the usual common knowledge assumptions. Thus the idealisation of the gamblers is part and parcel of the abstraction of de Finetti's betting problem, which as recalled above, aims at eliciting the *bookmaker*'s degrees of belief.

Rigidity is an immaterial abstraction of de Finetti's betting problem and is motivated by de Finetti's reluctance to appeal to the mathematical theory of utility (see de Finetti, 1969, especially Chapter 4). In order to avoid the potential complications arising from the diminishing marginal utility of money, de Finetti assumes that stakes should be *small*, an assumption to which he refers to as the *rigidity hypothesis* (de Finetti, 1974, p.77-78).

The abstraction of de Finetti's betting problem is fundamental to provide an instantiation of the Choice Norm which formalises admissibility as the avoidance of sure loss. De Finetti's (formal) notion of coherence is therefore context-dependent and it is justified only for the specific context defined by the betting problem. Under those restrictions only, the Choice Norm does entail probability as a first-order uncertainty norm of rational belief. De Finetti is explicit about the fact that no second-order uncertainty can be accommodated in his framework:

> Among the answers that do not make sense, and cannot be admitted are the following: "I do not know", "I am ignorant of what the probability is", "in my opinion the probability does not exist". Probability (or prevision) is not something which in itself can be known or not known: *it exists in that it serves to express, in a precise fashion, for each individual, his*

[35]The question as to how suitable the betting problem is as an elicitation device is raised by de Finetti in his later work on *proper scoring rules* and especially Brier's (see, especially de Finetti, 1962, 1969, 1972). I will postpone the analysis of admissibility as "miniminum expected loss under Brier's score" to further research.

> *choice in his given state of ignorance.* To imagine a greater degree of ignorance which would justify the refusal to answer would be rather like thinking that in a statistical survey it makes sense to indicate, in addition to those whose sex is unknown, those for whom one does not even know "whether the sex is unknown or not". (de Finetti, 1974, p.82, my emphasis)

As the quotation clearly suggests, his radically subjectivist position brings de Finetti to deny that there is anything to be modelled outside first-order uncertainty. Whilst de Finetti's radical stance certainly played a role in giving the subjective approach to probability full mathematical citizenship, mostly as a consequence of the much celebrated Representation theorem (see de Finetti, 1974), his highly idiosyncratic style contributed to hiding the applicability the corresponding epistemological framework to modelling interesting and relevant aspects of second-order uncertainty. It is interesting to note that at approximately the same time –de Finetti's monograph was published in Italian in 1970– Savage added to the second edition of his Savage (1954) the following footnote:

> One tempting representation of the unsure is to replace the person's single probability measure $P$ by a set of such measures, especially a convex set[. . . ]. (p.58)

I.J. Good is one prominent Bayesian who has not resisted the temptation (see Good, 1983).

4.5. **A two-dimensional framework.** Our running logical analogy now triggers the obvious question: If the Probability Norm —the requirement that uncertainty should be measured by probability— is fully justified for first-order uncertainty only, what happens outside its rather narrow borders?

By identifying first-order uncertainty with the epistemic state of a maximally idealised agent who is facing a maximally abstract choice problem, we can tackle the question by either relaxing the abstraction of the choice problem or (inclusively) the idealisation of the agent. Figure 1 illustrates the logical space of modelling which arises in this abstraction-idealisation coordinate system.

Starting from first-order Bayesian theory, we can relax completely the idealisation of the agent and consider abstract choice problems faced by experimental subjects. This identifies the domain of behavioural decision theory for which Prospect Theory offers the best-known framework (Kahneman

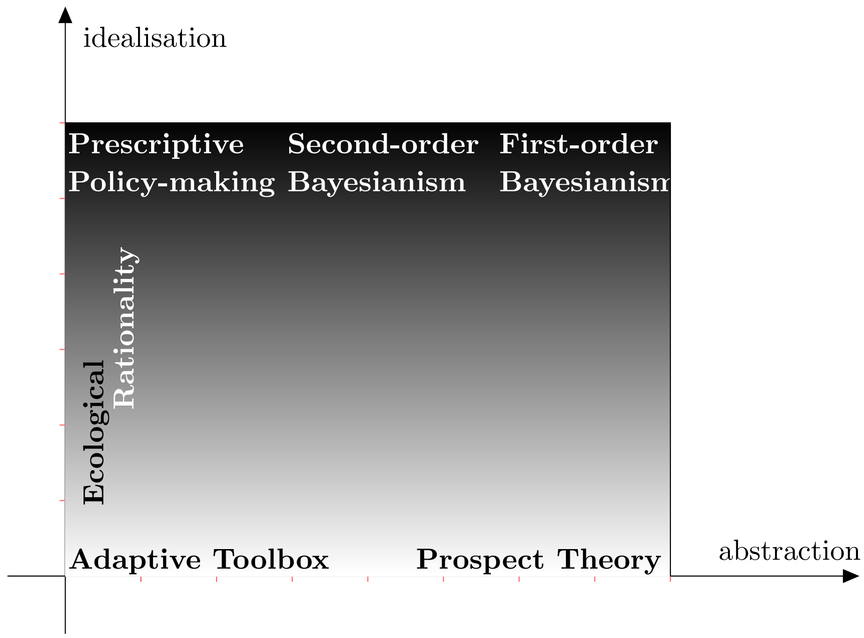


FIGURE 1. A stylised two-dimensional space of models determined by the parameters intervening in the Choice Norm, namely the idealisation of the agent and the abstraction of the choice problem.

and Tversky, 1979; Wakker, 2010). From there we can fully relax also the abstraction of the choice problem, ending up in the domain of the "adaptive toolbox" investigated in (Gigerenzer, 1999), where real people use a whole stock of heuristics to make decisions in the real world. Which heuristics should be used for which particular choice problem is a question of clear normative import which is addressed by the ecological approach to rationality (see Gigerenzer, 2012, especially Chapter 19). The top-left corner of the rectangle depicted in Figure 1 finally identifies the context in which idealised agents face real-world problems, of obvious interest in prescriptive policy making.

As recalled above, the present study is only concerned with normative models of rationality so the focus is on the interval delimited by maximal and minimal abstraction, corresponding to the darker area of the rectangle of Figure 1. The central intuition of the present proposal is that second-order uncertainty can be given a solid foundation by instantiating the Choice Norm on increasingly less abstract choice problems, or equivalently, by granting the decision-maker an increasingly large number of modelling privileges.

The next Section illustrates this familiar "extension by restriction" strategy (see Section 3.2 above) with an example in which the Choice Norm is instantiated with a second-order uncertainty problem featuring ignorance or ambiguity – as referred to the elicitation of subjective degrees of belief.

Before moving to that, however, it is important to stress that the rectangle in Figure 1 should not be interpreted as providing a 'grand unified theory of uncertain reasoning and rational decision', as it were. Whilst it accommodates a number of currently popular albeit hardly mutually consistent approaches to the topic, it is not meant to suggest that all models of rationality are or should be commensurable in some particularly precise way. Again in the spirit of our central logical analogy, Figure 1 is best interpreted as highlighting the variety which is determined by the context-dependence of rational norms of belief.

## 5. Towards second-order uncertainty

Let us go back to the (Gilboa et al., 2011) plea for rational modesty recalled in Section 2 above. In real life, where time and information are scarce, it might well be wiser to hold our judgment and to postpone our decisions until all the facts are in, as it were. However, the completeness of the revealed belief approach normatively requires that subjective degrees be linearly ordered forcing the Bayesian agent to have a probability for all elementary events of interest. As there is no *ought* without a *can*, many anti-Bayesians conclude that the Probability Norm is not necessary for rational belief. This goes hand in hand with a foundational perspective which can be traced back to Keynes (1921) and Knight (1921) who insisted, albeit with rather distinct arguments, that not all (economic) problems admit of a probabilistic quantification of uncertainty. This point of view has been supported over the past four decades by vast yet not uncontroversial experimental evidence.[36] As a consequence, a number of anti-Bayesian proposals currently take issue with probability and the maximisation of expected utility as adequate norms of rationality under uncertainty. The main line of the argument, as recalled in Section 2 above, is this: since it fails to reflect the agent's ignorance, the Probability Norm fails as a measure of rational belief.

[36] Binmore et al. (2012) report on a recent experiment which casts substantial doubts on the alleged universality of the "ambiguity aversion" phenomenon in Ellsberg-type problems.

To fix ideas consider a policy maker who is faced with a problem whose uncertainty is intuitively felt to be "hard to quantify". Take, for example, the event "Greece will exit the eurozone by the end of May 2013", or GREXIT as it is referred to in the financial lingo.[37] There is certainly no obvious state space that captures the relevant scenarios connected to GREXIT. Since nobody can really come up with such a state space, anti-Bayesians suggest that GREXIT relates to the kind of uncertainty which cannot be quantified probabilistically. Putting to one side an analysis of the slippery concept of 'non-probabilistic uncertainty' I would like to focus on how the two-dimensional characterisation of rationality of Figure 1 helps us to see that whilst intuitively plausible, the anti-Bayesian answer to the issues raised by GREXIT-like problems fails to be normatively persuasive.

Let us begin by noticing that Figure 1 gives us two main modelling options. The first such option is to take GREXIT as a first-order uncertainty problem, i.e. maximally specified so that no modelling options are available to the decision-maker who is facing it. In this case a maximally idealised policy maker *must* come up with a probability value representing their belief in GREXIT. On a first-order uncertainty reading, it doesn't matter at all if the process of producing an admissible (i.e. satisfying the Choice Norm) quantification of the policy maker's degree of belief for GREXIT is "hard" or difficult, in any sense. Under these assumptions the idealised policy maker is normatively forced to attach GREXIT a unique point in the real unit interval.

Our second modelling option is to frame GREXIT as a second-order uncertainty problem. In this case, we might want to capture the fact that the scenarios which are relevant to GREXIT are so many and currently so poorly understood that too much information is lost by summarising our uncertainty about GREXIT with one real number. This clearly means allowing some modelling privileges to the decision-maker facing the GREXIT problem. Hence, in a second-order uncertainty framing, the agent might rationally (at the second-order) decide that in the present state of information an interval representation of uncertainty is preferable, or even that it is best not to give a (public) answer in order to prevent the attack of financial speculators and so on, depending on the degree of modelling privileges that we are willing to grant to the decision-maker.

It goes without saying that I will not offer a formal solution to the problem of, say, measuring uncertainty in problems as complex as GREXIT.

---

[37] The term appears to have been coined by Willem Buiter of Citigroup, in his 6 February 2012 report.

The purpose of this example is rather that of suggesting a natural framing for the following question: What norms of rationality are adequate if we give decision-makers some modelling privileges, or equivalently, relax the abstraction of the choice problem? The answer must clearly come in two steps. First we need to provide an account of what a principled relaxation of abstraction might be, i.e. the extent to which we are willing to grant modelling privileges to decision-makers. Having done this, we can instantiate the Choice Norm approach illustrated above and provide a justified notion of admissibility for the specific class of choice problems at hand.

As an example of the applicability of this general framework, I will illustrate how imprecise probabilities can be justified as the rational norm of belief in a for-profit betting problem –a more realistic version of de Finetti's problem which relaxes the Swapping condition.

5.1. **Betting for profit.** Recall from Section 4.4 above that de Finetti's betting problem is so formulated as to force the (idealised) bookmaker to choose fair betting quotients. Publishing books with zero expectation is necessary and sufficient to protect the bookmaker from being forced into sure loss by rational gamblers, possibly through Swapping.

Fedel et al. (2011) investigate the relaxation of this first-order uncertainty modelling feature by considering a betting scenario in which bookmakers are motivated by making profit in a market-like environment. This clearly imposes the relaxation of Swapping so that gamblers can no longer choose the sign of the stake for their bets. Analogously, bookmakers are allowed to differentiate between buying and selling prices, thus giving rise to the notions of lower and upper probabilities which are well familiar from the theory of imprecise probabilities.[38] How can the Choice Norm be instantiated for this (second-order uncertainty) *for-profit betting problem*? A simultaneous extension of classical Bayesianism to imprecise and fuzzy probabilities is carried out in (Fedel et al., 2011) by constructing the analytic framework of imprecise probabilities on top of a many-valued algebraic semantics. For present purposes I will limit myself to an informal discussion of how the notion of admissibility is arrived at and refer the interested reader to the original paper for precise mathematical details and for further motivation concerning the extension to *fuzzy* events.

The key idea, as anticipated above, is that the bookmaker publishes his book by assigning a pair of real numbers $\alpha_i, \beta_i$, intuitively interpreted as

[38]The standard reference for the field is (Walley, 1991), which also provides detailed historical background. Miranda (2008) offers an outline of the most recent developments.

the sup of the buying price and the inf of the selling price, respectively, to all events in $E_i$ in the book. So a book (or a system of bets) is defined by a set of pairs $(E_1, [\alpha_1, \beta_1]), \ldots, (E_n, [\alpha_n, \beta_n])$ such that $0 \leq \alpha_i \leq \beta_i \leq 1$ for $i = 1, \ldots n$. In standard market models, the inequality between buying and selling prices is assumed to be strict, thus defining the so-called *bid-ask spread* (see e.g. Hasbrouck, 2007). The assumption is variously motivated: from potential asymmetric information to the bookmaker's need to cover fixed transaction costs. So, in the perspective of reducing the abstraction of the choice problem, it is very natural to consider bookmakers who are set a positive spread. Note however that de Finetti's framework is mathematically fully recovered when $\alpha_i = \beta_i$ for all $E_i$ in the book.

The Choice Norm demands that we define an appropriate notion of admissibility that must be satisfied by any justified norm of belief for this refined class of for-profit betting problems. Let us begin by noting that de Finetti's notion of coherence will not do for, when publishing his book, the bookmaker knows that the (maximally idealised) gamblers have a choice between:

(1) paying $\beta\lambda$ for the right to receive $\lambda v(E)$, i.e. betting on $E$
(2) receiving the payment of $\alpha\lambda$ to pay back $\lambda v(E)$, i.e. betting on $E^c$

where $\alpha, \beta \in [0, 1]$, $\lambda > 0$, $v(E)$ maps $E$ to $[0, 1]$ and $E^c$ denotes the complement of $E$. It is immediate to see that de Finetti's coherence is sufficient to avoid blatant irrationality but is not necessary, in this less abstract choice problem. Thus admissibility cannot be defined as "avoiding sure loss". Consider the simple book $B = \{(E, [0, 1]), (E^c, [.5, 1])\}$. Whilst the bookmaker who published $B$ would certainly protect himself against sure loss, there is a clear sense in which $B$ is blatantly irrational in *for-profit betting problems* as the bookmaker is setting too wide a spread. In the market-like scenario which is motivating the construction of this specific choice problem, by setting too wide a spread between buying and selling prices, the bookmaker is effectively encouraging gamblers to trade with those bookmakers who publish more attractive odds. As this openly violates the notion of purposeful behaviour which the Choice Norm intends to capture, an *inadmissible* choice can be defined as a choice of betting intervals which could be refined without leading to sure loss.[39] In other words, a book is inadmissible if the bookmaker writes odds which are unnecessarily conservative.

One central result of (Fedel et al., 2011) can be stated in the terminology of the present note as follows: inadmissibility is avoided if and only if the bookmaker's odds can be extended to an upper prevision over a suitable

[39] Inadmissibility is referred to as the *bad bet criterion* in the terminology of (Fedel et al., 2011).

algebra of events. Dual results follow from lower previsions and probabilities. Thus, instantiated with for-profit betting problems, the Choice Norm justifies imprecise probability as a norm of rational belief.

There is an increasingly wide consensus on the fact that interval-valued probability captures relevant features of second-order uncertainty, especially ignorance. This was partly acknowledged by Savage in the above-quoted remark (see p. 21) to the effect that the expressive power provided by imprecise probability, and in particular by convex sets of probability measures, would be necessary to model increasingly more interesting classes of decision problems under uncertainty. De Finetti, as one would expect, rejected any such extension.[40] In terms of our opening question, clearly de Finetti thinks of his version of Bayesian theory not as one of the many possible alternatives, but as the only possibility worth considering. Yet, the general framework outlined here suggests, as Savage seems to have clearly anticipated, that a Type 1 attitude can lead to foundationally conservative formal advance just by supplying the Choice Norm with context-dependent formalisations of admissibility.

5.2. **Second-order uncertainty vs second-order probability.** As a slightly unfortunate consequence of the terminology which I am using, one might be led to believe that second-order uncertainty is to be measured by second-order probabilities. The above two-dimensional "agent-problem" characterisation of rationality however, suggests that this is not, in general, the case. Whilst the precise details will have to be addressed separately, the intuitive argument as to why second-order uncertainty, as presently characterised, need not be measured by second-order *subjective* probabilities, goes as follows.

Recall the space of models depicted in Figure 1 above. The standard Probability Norm is justified, in such a framework, for first-order uncertainty only. By relaxing the abstraction of de Finetti's betting problem we have introduced some second-order uncertainty features in the choice problem, which in turn led to a justification of imprecise probabilities as a second-order norm for rational belief. The relaxation of abstraction can be interpreted as granting the decision-maker some modelling privileges, such as discouraging potential gamblers by setting very wide intervals as a consequence of the bookmakers' "ignorance" about the events in question. Exactly how wide such a spread between buying and selling price should be, is now a choice which the bookmaker makes as modeller, rather than decision-maker.

[40]See (de Finetti, 1975, Appendix 19.3).

But there is intuitively no reason as to why norms of rational decision-making (i.e. the Probability Norm) should apply "one level up" to rational decision-modelling – say choosing those events $E_i$ of the book on which the bookmaker effectively does not want anyone to bet, a fact which can naturally be expressed by setting the odds for such $E_i$s to $[0, 1]$. One central feature of the present proposal is that norms for decision-making and norms for decision-modelling, albeit contiguous, are distinct problems.

Note that the general strategy of justifying distinct belief norms for *distinct* choice problems, clearly prevents the resulting second-order norms of rational belief from entering an infinite regress, a concern which typically hovers over second-order (subjective) probability. D. Hume is usually credited with spelling out an early version of this argument, whose contemporary version appears in (Savage, 1954, p.58). Interestingly enough this is the very page in which Savage added the second-edition footnote, recalled above, to the effect that the attitude of being "unsure" could be represented by allowing convex sets of probabilities.

Clearly, no infinite regress can arise when first-order uncertainty is represented by *objective* probabilities as in (Hansson, 2009). However, from the point of view of Bayesian theory, which the present framework aims at extending, the assumption that objective probability represents the agent's uncertainty must be made with substantial care. Whether objective probability and possibly second-order probability can be accommodated within the framework outlined in this paper remains an open question.

## 6. Conclusion

In making the comparison which constitutes the leitmotif of this paper, I have implicitly assumed that non-standard logics and uncertain reasoning are distinct research areas. Whilst this tends to be largely the case if we take the quantitative vs qualitative divide very seriously,[41] it might rightly be objected the opposite is true: Mathematical and philosophical overlaps abound between probability and logic in both the sub-areas of quantitative and qualitative representations of uncertainty.

Whilst logic is in itself central to probability and uncertain reasoning, my present aim was not to discuss the fruitful *formal* interactions between logic and uncertain reasoning. This is the object of a number of thorough investigations including Paris (1994); Howson (2009); Haenni et al. (2011);

[41] Suffice it to mention that much of the popularity enjoyed by non-monotonic logics during the 1980s was more or less directly linked to the idea that logic would better serve the (computational) needs of artificial intellingence than probability.

Makinson (2010). My present goal was rather to suggest that the way many non-standard logics developed as extensions of the classical (propositional) one offers a potentially very fruitful methodological example of how the limited expressive power of a formal model can be addressed by clarifying its intended domain of application. By comparing uncertain reasoning with very familiar developments in non-standard logics I suggested that the much-needed formal advance of Bayesian theory, especially with regards to second-order uncertainty, may be greatly facilitated if uncertain reasoners take seriously the lessons offered by non-standard logics, especially contextualisation. Since uncertain reasoning does not occur in a vacuum, questions about rational norms are best relativised to an agent and a choice problem.

This simple observation sheds new light on the "big picture" and provides evident prospects for foundational and formal advance which will hopefully give rise to a *doing-thinking-doing* virtuous circle in uncertain reasoning. From the foundational point of view, the two-dimensional contextualisation presented above clearly illustrates how some popular anti-Bayesian criticisms are intuitively appealing, yet normatively inconclusive. For the arguments based on "probability is hard to quantify" to have normative force, one must clearly put forward a classification of choice problems and prove that for some suitable class, probability offers an incomplete formalisation. The analysis of Ellsberg-type problems recalled above doesn't offer particularly useful insights in this direction.

The strategy of starting with maximally abstract choice problems and then refining them to model increasingly more realistic features of second-order uncertainty appears to be foundationally more transparent than that of postulating the existence of putatively distinct notions of uncertainty, some of which are probabilistically quantifiable, some of which are not. It is hoped that this problem-based approach to measuring uncertainty may prove useful in developing norms of rational belief and decision which are continuous with the real-world applications which eventually motivate much of our interest in this subject. Real-world problems, from climate change to economic uncertainty to biomedical risk, are crying out to be the uncertain reasoner's Tweety.

## Acknowledgments

I have presented the main ideas of this paper at Kent's Centre for Reasoning and at the LSE Choice Group in London. I would like to thank Jon Williamson and Richard Bradley for inviting me to speak at those seminars and both audiences for their very valuable feedback. I am very grateful to Gregory Wheeler for his comments on an earlier draft and for many stimulating discussions on the topics covered in this paper. Thanks also to two referees, whose thorough reviews helped me to improve the chapter in many ways.

Finally, readers familiar with David Makinson's work, will certainly have spotted a number of terms and expressions which are easily associated with, and sometimes directly coming from, David's papers and books. I realised this only when the first draft of this chapter was completed. As I started fetching all the originals to give credit where credit was due, it occurred to me that leaving those paraphrases uncredited would probably be the most direct way to express how influential David's way of doing logic is to my way of thinking about logic. So I stopped.

## References

E.W. Adams. What is at stake in the controversy over conditionals. In G. Kern-Isberner, W. Roedder, and F. Kulman, editors, *WCII 2002*, pages 1–11. Lecture Notes in Artificial Intelligence, Springer, 2005.

O. Arieli and A. Zamansky. Distance-based non-deterministic semantics for reasoning with uncertainty. *Logic Journal of IGPL*, 17(4):325–350, June 2009.

J. Barwise and H. J. Keisler, editors. *Handbook of Mathematical Logic*. North-Holland, 1977.

T.F Bewley. Knightian decision theory. Part I (1986). *Decisions in Economics and Finance*, 25:79–110, 2002.

K. Binmore, L. Stewart, and A. Voorhoeve. An experiment on the Ellsberg paradox. *Working paper*, 2012.

P. Blackburn, J. Van Benthem, and F. Wolter. *Handbook of modal logic*. Elsevier, 2007.

W. Bossert and K. Suzumura. *Consistency, Choice, and Rationality*. Harvard University Press, 2010.

W. Bossert and K. Suzumura. Revealed preference and choice under uncertainty. *SERIEs*, pages 1–6, February 2011.

W. Carnielli, M.E. Coniglio, and J. Marcos. Logics of formal inconsistency. In D.M. Gabbay and F Guenthner, editors, *Handbook of Philosophical Logic 2nd Edition, Vol 14*, volume 14, pages 1– 93. Springer, 2007.

M. Colyvan. Is probability the only coherent approach to uncertainty? *Risk Analysis*, 28(3):645–52, June 2008.

L. Daston. *Classical Probability in the Enlightenment*. Princeton University Press, 1988.

B. de Finetti. Sul significato soggettivo della probabilità. *Fundamenta Mathematicae*, 17:289–329, 1931.

B. de Finetti. Does it make sense to speak of good probability appraisers? In I.J. Good, editor, *The Scientist Speculates*, pages 357–364. Heinemann, 1962.

B. de Finetti. *Un Matematico e L'economia*. Franco Angeli, Milano, 1969.

B. de Finetti. *Probability, Induction and Statistics*. Wiley and Sons, 1972.

B. de Finetti. Bayesianism: its unifying role for both the foundations and the applications of statistics. *Bulletin of the International Statistical Institute*, 39:349–368, 1973.

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

B. de Finetti. *Theory of Probability, Vol 2*. John Wiley and Sons, 1975.

B. de Finetti. A short confirmation of my standpoint. In M. Allais and O. Hagen, editors, *Expected Utility Hypotheses and the Allais Paradox*, page 161. D. Reidel, 1979.

B. de Finetti. Philosophical Lectures on Probability. Springer Verlag, 2008.

D. Ellsberg. Risk, ambiguity and the savage axioms. *The Quarterly Journal of Economics*, 75(4):643–669, 1961.

M. Fedel, H. Hosni, and F. Montagna. A logical characterization of coherence for imprecise probabilities. *International Journal of Approximate Reasoning*, 52:1147–1170, 2011.

D.M. Gabbay, editor. *What Is a Logical System?* Oxford University Press, 1995.

D.M. Gabbay and A. Hunter. Making inconsistency respectable 1 : A logical framework for inconsistency in reasoning. In *Fundamentals of artificial intelligence research.*, pages 19–32. Lecture notes in Computer Science, Springer, 1991.

D.M. Gabbay and J. Woods. Normative models of rational agency: the theoretical disutility of certain approaches. *Logic Journal of IGPL*, 11(6): 597–613, 2003.

D.M. Gabbay and J. Woods, editors. *Handbook of the History of Logic Vol 8. The Many Valued and Non Monotonic Turn in Logic*. North-Holland, 2007.

D.M. Gabbay and John Woods. The practical turn in logic. In D.M. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic, 2nd Edition, Vol 13*, volume 13, pages 15–122. Springer, 2005.

A. Gelman. Induction and deduction in Bayesian data analysis. *Rationality, Markets and Morals*, 2:67–78, 2011.
(1999).

G. Gigerenzer, P.M. Todd, ABC Research Group. *Simple Heuristics That Make Us Smart*. Oxford University Press, 1999

G. Gigerenzer. *Rationality for Mortals: How People Cope with Uncertainty.* Oxford University Press, 2008.

G. Gigerenzer, P.M. Todd, ABC Research Group. *Ecological rationality: Intelligence in the world.* Oxford University Press, 2012

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

I. Gilboa, A. Postlewaite, and D. Schmeidler. Rationality of belief or: why Savage's axioms are neither necessary nor sufficient for rationality. *Synthese Online First*, 19 October 2011.

I.J. Good. *Good thinking: The Foundations of Probability and Its Applications.* The University of Minnesota Press, 1983.

R. Haenni, J.W. Romeijn, G. Wheeler, and J. Williamson. *Probabilistic Logics and Probabilistic Networks.* Springer, 2011.

S. O. Hansson. Measuring uncertainty. *Studia Logica*, 93(1):21–40, 2009.

J. Hasbrouck. *Empirical Market Microstructure: The Institutions, Economics and Econometrics of Securities Trading.* Oxford University Press, 2007.

A. Heyting. *Intuitionism: An Introduction.* North-Holland, 1956.

H. Hosni and J.B. Paris. Rationality as conformity. *Synthese*, 144(2):249–285, 2005.

C. Howson. Can logic be combined with probability? Probably. *Journal of Applied Logic*, 7(2):177–187, 2009.

E.T. Jaynes. *Probability Theory: The Logic of Science.* Cambridge University Press, 2003.

D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.

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

F.H. Knight. *Risk, Uncertainty and Profit.* Beard Books Inc, 1921.

I. Levi. The paradoxes of Allais and Ellsberg. *Economics and Philosophy*, 2:23–53, 1986.

D. Makinson. The paradox of the preface. *Analysis*, 25(6):205, 1965.

D. Makinson. General patterns in nonmonotonic reasoning. In *Robinson, J. (eds.): Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3.*, pages 35–110. Oxford University Press, 1994.

D. Makinson. *Bridges from classical to nonmonotonic logic.* College Publications, London, 2005.

D. Makinson. Logical questions behind the lottery and preface paradoxes: Lossy rules for uncertain inference. *Synthese Online First*, 30 August 2011.

E Miranda. A survey of the theory of coherent lower previsions. *International Journal of Approximate Reasoning*, 48(2):628–658, June 2008.

L.S. Moss. Applied logic: A manifesto. In D.M. Gabbay, S. Gonchargov, and M Zakharyaschev, editors, *Mathematical Problems from Applied Logic I*, pages 317–343. Springer, 2005.

J.B. Paris. *The Uncertain Reasoner's Companion: A Mathematical Perspective*. Cambridge University Press, 1994.

J.B. Paris. Common sense and maximum entropy. *Synthese*, 117(1):75–93, 1998.

J.B. Paris. A Note on the Dutch Book Method. In T. Seidenfeld, G. De Cooman, T. Fine (eds), *ISIPTA '01, Proceedings of the Second International Symposium on Imprecise Probabilities and Their Applications, Ithaca, NY, USA*. Shaler, 2001.

J.B. Paris. Deriving information from inconsistent knowledge bases: A completeness theorem. *Logic Journal of IGPL*, 12:345 – 353, 2004.

J.B. Paris and A. Vencovská. A note on the inevitability of maximum entropy. *International Journal of Approximate Reasoning*, 4(3):183–223, 1990.

A. Rubinstein. *Lecture Notes in Microeconomic Theory.* Princeton University Press, 2006.

A.L. Rukhin. Admissibility: Survey of a concept in progress. *International Statistical Review*, 63(1):95–115, 1995.

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

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

A.K. Sen. Collective Choice and Social Welfare. 1970.

G. Shafer. Savage revisited. *Statistical Science*, 1(4):463–485, 1986.

J. Tirole. Incomplete contracts: Where do we stand? *Econometrica*, 67(4): 741–781, 1999.

J. van Benthem. Logic and reasoning: Do the facts matter? *Studia Logica*, 88:67–84, 2008.

H. van Ditmarsch, W. van Der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Springer, 2007.

P.P. Wakker. *Prospect Theory For Risk and Ambiguity.* Cambridge University Press, 2010.

P. Walley. *Statistical Reasoning with Imprecise Probabilities.* Wiley, 1991.

G. Wheeler. Applied logic without psychologism. *Studia Logica*, 88(1):137–156, February 2008.

J. Williamson. *In Defence of Objective Bayesianism.* Oxford University Press, 2010.

Scuola Normale Superiore, Pisa and Centre for Philosophy of Natural and Social Science, London School of Economics, London

*E-mail address*: h.hosni@gmail.com