The Normative Representation of Quantified Beliefs by Belief Functions.

Philippe SMETS[1]
I.R.I.D.I.A. Université Libre de Bruxelles

Summary: The use of belief functions has recently been advocated as an alternative to the use of probability functions in order to represent quantified beliefs. Such a proposal lacked justification. We present a set of requirements that justify the use of belief functions. The assessment of the validity of these requirements provides a tool for assessing the relative value of normative models of subjective behaviors.

Keywords: Belief functions, quantified belief, subjective probabilities, requirements for belief, transferable belief model.

1. Introduction.

To build a 'thinking robot' can be seen as the ultimate aim of artificial intelligence. To be 'viable' such a robot should be able to reason and act within an uncertainty-riddled environment. Uncertainty assumes numerous forms (Smithson, 1989, Smets, 1991) and usually induces 'beliefs', i.e. the graded dispositions that guide 'our' behavior. If the robot is to hold such 'beliefs', then a mathematical model representing beliefs is needed. This paper develops such a model. Our approach is normative. The 'robot' - the agent that holds the beliefs - henceforth called You, is an ideal rational subject. We propose requirements that should be satisfied by the beliefs held by such an agent. These requirements are satisfied if beliefs are quantified by belief functions (Shafer, 1976). The derived model is the transferable belief model (Smets and Kennes, 1994).

As far as we know, this is the first axiomatization that justifies the use of belief functions to represent quantified beliefs. Wong et al. (1990) have proposed qualitative axioms for a belief ordering and shown that it can always be represented by belief functions, but they fail to show that only belief functions can represent such an ordering. Of course, using belief functions clashes with the current trend advocated by the Bayesian School that claims that quantified beliefs must be represented by probability functions. What makes our axiomatization interesting is that the analysis of the proposed requirements provides a tool for comparing competing normative models.

---

[1]Research work partly supported by the Action de Recherches Concertées BELON funded by a grant from the Communauté Française de Belgique and the ESPRIT III, Basic Research Action 6156 (DRUMS II) funded by a grant from the Commission of the European Communities.

After summarizing the necessary technical information, we present the proposed requirements together with a few illustrative examples by way of illustration. Proofs are omitted[2]. We focus essentially on the rationality constraints that underpin the requirements and that really justify them.

2. The credibility domain.

The aim of this paper is to develop the mathematical structure of a function Cr, temporarily called a credibility function that quantifies Your beliefs that the actual world belongs to such or such subset of possible worlds. Beliefs can equivalently be allocated to 'propositions', to the subsets of worlds that denote the propositions or to events. Our presentation is based on the possible worlds approach (Bradley and Swartz, 1979, Ruspini, 1986).

The strength of the beliefs entertained by You at time t is defined relatively to a given evidential corpus (EC_t^Y), i.e., the set of evidence in Your mind at time t. This EC_t^Y is equivalent to the 'background knowledge' used by the Bayesians, and intuitively just means 'all You know'. Only one belief holder, You, is considered in this paper, and time t is unique except when belief revision is studied.

2.1. The propositional space.

Let $\mathtt{L}$ be a finite propositional language, supplemented by the tautology and the contradiction, denoted $\top$ and $\perp$, respectively, and closed under the usual Boolean connectives $\neg$, $\vee$ and $\wedge$. Let Ω be the set of worlds that correspond to the interpretations of $\mathtt{L}$ and built so that no two worlds denote logically equivalent propositions. For any proposition X in $\mathtt{L}$, let $[\![X]\!] \subseteq \Omega$ be the set of worlds identified by X (i.e., those worlds where X is true).

Let $\Re$ be a Boolean algebra of subsets of Ω (thus closed under union, intersection, complement, and containing Ω and $\emptyset$). Let $At(\Re)$ be the set of atoms of the algebra $\Re$, i.e., the non-empty elements of $\Re$ whose intersection with any element of $\Re$ is either themselves or the empty set. The atoms of $\Re$ are the elements of a partition of Ω. When $\Re$ is the power set 2^Ω of Ω, the atoms of $\Re$ are the singletons of Ω.

We assume that among the worlds of Ω a particular one, denoted ω_0, corresponds to the actual world. You do not know at t which world is ω_0. You can only express Your beliefs at t that the actual world ω_0 does or does not belong to various subsets of worlds. By definition the actual world ω_0 is an element of Ω.

[2]Proofs are presented in a Technical Report that can be obtained from the author.

The domain of Your beliefs is assumed to be a Boolean algebra of subsets $\Re$ of Ω. Indeed whenever You can express Your belief that ω_0 belongs to a set A and to a set B, You can also express Your belief that ω_0 belongs to their complement (relative to Ω), union and intersection. We do not assume that $\Re$ is 2^Ω. Because of Your limited reasoning power, Your beliefs that result from Your evidential corpus at time t are not necessarily so detailed that every subset of Ω can be assigned a belief. In Smets (1995), we further examine the case where the credibility is defined only on a subalgebra Ω' of Ω, in which case the belief holder is not even sure that ω_0 is an element of Ω'. For simplicity's sake, that particular case will not considered here.

We call Ω the frame of discernment (the frame for short). We call the pair $(\Omega, \Re)$ a propositional space.

2.2. Doxastic equivalence.

We now introduce the concept of doxastic equivalence, i.e., equivalence relative to the evidential corpus EC_t^Y. As an example, suppose You want to decide whether to go to a movie or stay at home tonight. You have decided to toss a coin, if it is heads, You will go to the movie, and if it is tails, You will stay at home. (These are the pieces of evidence in EC_t^Y). Then 'heads' and 'going to the movie' are doxastically equivalent as they share the same truth status given what You know at t. Of course, they are not logically equivalent (Kyburg, 1987a). Logical equivalence implies doxastic equivalence, not the reverse.

Formally, in propositional logic, two propositions p and q defined on L are logically equivalent iff $[\![p]\!] = [\![q]\!]$. They are doxastically equivalent (for You at t, i.e., given EC_t^Y) iff the sets of worlds $[\![p]\!]$ and $[\![q]\!]$ that denote them share the same worlds among those in $[\![EC_t^Y]\!]$, i.e., $[\![EC_t^Y]\!] \cap [\![p]\!] = [\![EC_t^Y]\!] \cap [\![q]\!]$, where $[\![EC_t^Y]\!]$ denotes the set of worlds where all the propositions deduced on L from EC_t^Y are true.

Doxastic equivalence under EC_t^Y is denoted by: $[\![p]\!] =_{EC_t^Y} [\![q]\!]$.

2.3. Change in the evidential corpus.

A piece of evidence is defined here as a proposition that You will accept as true once You learn it. Adding a piece of evidence Ev to Your evidential corpus EC_t^Y leads to a new evidential corpus, denoted $EC_t^Y \cup \{Ev\}$. The piece of evidence $\neg Ev$ is the negation of the piece of evidence Ev, i.e., the proposition that would characterize Ev is accepted as false. By construction, $[\![Ev]\!] \cap [\![\neg Ev]\!] = \emptyset$, and $EC_t^Y \cup \{Ev \cup \neg Ev\} = EC_t^Y$. The last equality results from the fact that $Ev \cup \neg Ev$ is a tautology and adding a tautology to an evidential corpus will leave it unchanged.

We say that Ev is compatible with EC_t^Y if $[\![EC_t^Y]\!] \cap [\![Ev]\!] \neq \emptyset$.

2.4. Belief functions

We will derive that quantified beliefs are represented by 'belief functions' (Shafer, 1976). Given a propositional space $(\Omega, \mathfrak{R})$, a belief function is a function bel from $\mathfrak{R}$ to [0, 1] such that :

1) $bel(\emptyset) = 0$

2) for all $A_1, A_2, \ldots A_n \in \mathfrak{R}$,

$$bel(A_1 \cup A_2 \cup \ldots A_n) \geq \sum_i bel(A_i) - \sum_{i>j} bel(A_i \cap A_j) \ldots - (-1)^n bel(A_1 \cap A_2 \cap \ldots A_n) \quad (2.1)$$

Usually, $bel(\Omega) = 1$ is also assumed. It can be ignored and we will only require that $bel(\Omega) \leq 1$.

3. Axiomatic justification for the use of belief functions.

Let $\mathcal{F}$ denote the set of functions that could be used to quantify someone's beliefs and let us call 'credibility functions' the elements of $\mathcal{F}$. Initially, credibility functions can be any set-function. We then introduce requirements that we feel any measure of belief should satisfy. Each requirement increasingly limits $\mathcal{F}$, up to the point where $\mathcal{F}$ reduces itself to the set of belief functions. Even though probability functions are special cases of belief functions, the family of probability functions is not expressive enough to satisfy all requirements (see section 3.5).

3.1. The credibility function.

The first requirement states that beliefs are pointwise valued, non negative and monotone for inclusion.

Requirement 1: Let $(\Omega, \mathfrak{R})$ be a propositional space. Your beliefs allocated to the elements of $\mathfrak{R}$ are quantified by a pointwise function $Cr: \mathfrak{R} \rightarrow [0,1]$, where $Cr(\emptyset) = 0$, and Cr is monotone for inclusion, i.e., for all $A, B \in \mathfrak{R}$, if $A \subseteq B$, then $Cr(A) \leq Cr(B)$.

The triple $(\Omega, \mathfrak{R}, Cr)$ is called a credibility space. As Cr is induced by EC_t^Y, the belief state of You at t is fully described by $(\Omega, \mathfrak{R}, Cr, EC_t^Y)$ and we call that quadruple a belief state.

That Cr is a pointwise function is not without consequence as it leads to rejecting belief representations based on interval valued probabilities or sets of probability functions (Kyburg, 1987b, 1993, Voorbraak, 1993, Levi, 1980, Waley, 1991).

We introduce the strong but obvious requirement that propositions doxastically equivalent for You at t receive the same beliefs (Kyburg, 1987a).

Requirement 2: Doxastic Consistency.
Let $(\Omega, \Re_i, Cr_i)$, i=1,2, be two credibility spaces induced by the same EC_t^Y. Let $A_1 \in \Re_1$, $A_2 \in \Re_2$.

$$\text{If } A_1 =_{EC_t^Y} A_2\text{, then } Cr_1(A_1) = Cr_2(A_2).$$

Requirement 2 implies that those subsets of Ω that belong to both algebras will receive the same belief: indeed the propositions that denote them are doxastically equivalent. Hence the belief given to a subset of Ω does not depend on the structure of the algebra to which the subset belongs.

3.2. Convexity of the set of credibility functions.

We first accept that probability functions belong to $\mathcal{F}$. This could be either directly assumed, or equivalently derived from the Hacking frequency principle that claims:

If Chance(X) = p, then Belief(X) = p.

where Chance corresponds to Objective Probability, and Belief to Cr.

Requirement 3: Probability functions are credibility functions.

We then proceed by showing that the linear combination of two credibility functions is a credibility function.

Example 1: The Horse Race. Suppose a horse race involving three horses: Allan, Blues and Carol. Tomorrow at 7 AM, it will be decided depending on the outcome of a coin tossing experiment, if the race will be run at 10 AM or 4 PM. Let α be the probability that the race is run at 10 AM. The time of the race influences Your beliefs about which horse will win. Let Cr_1 and Cr_2 be the credibility functions that describe Your beliefs about which horse will win if the race is run at 10 AM or at 4 PM, respectively. You must buy a ticket now. Let Cr_{12} be the credibility function that describes Your beliefs held by now about the winner. We essentially assume that $Cr_{12}(A)$ for $A \subseteq$ {Allan, Blues, Carol} depends only on $Cr_1(A)$, $Cr_2(A)$ and α. The next requirement formalizes this constraint. □

Requirement 4: Let an evidential corpus EC_t^Y and the pieces of evidence Ev_1 and Ev_2 where $Ev_2 = \neg Ev_1$, and where both are compatible with EC_t^Y. Let α be the probability that You learn Ev_1 and $1-\alpha$ the probability that You learn Ev_2. Let Cr_1, Cr_2 and Cr_{12} be the credibility functions that represent Your beliefs on a propositional space $(\Omega, \Re)$ as induced by $EC_t^Y \cup \{Ev_1\}$, $EC_t^Y \cup \{Ev_2\}$, and EC_t^Y, respectively.
Then there exists a function $F : [0,1]^3 \rightarrow [0,1]$ such that for all $A \in \Re$,

$$Cr_{12}(A) = F_\alpha(Cr_1(A), Cr_2(A))$$

where $F_\alpha(x,y)$ is continuous in $(x,y)\in[0,1]^2$, strictly monotone for both components and idempotent ($F_\alpha(x,x) = x$).

Under requirements 1, 3 and 4, we prove that:

$$Cr_{12}(A) = \alpha\, Cr_1(A) + (1-\alpha)\, Cr_2(A) \qquad \text{for all } A\in\mathfrak{R} \tag{3.1}$$

Thus F is a convex set, a property shared by subjective probability functions and belief functions, but neither by the set of possibility nor by the set of necessity measures. Hence requirements 3 and 4 eliminate possibility and necessity measures for representing beliefs.

3.3. Uninformative changes of $\mathfrak{R}$.

Example 2: The Killer's Nationality. Suppose a person was murdered. Let Cr_0 represent Your beliefs that the killer (k) is English, German, French or Italian. Cr_0 is defined on the subsets of {E, G, F, I}. We consider how Cr_0 will be adapted when the domain of Your belief is changed. Two transformations are considered: coarsening and refinement. In the first case, suppose French and Italian are grouped into the set ‘Mediterranean’. The new space {E, G, M} is a coarsening of the initial space. In the second case, suppose the set ‘French’ is partitioned into two subsets, the sets ‘FrenchTuc’ (FT) and ‘FrenchPic’ (FP). The new space {E, G, FT, FP, I} is a refinement of the initial space. These transformations of the frames on which Your beliefs are defined are said to be 'uninformative' inasmuch as Your evidential corpus EC_t^Y is unchanged for what concerns Your beliefs about the killer's nationality. To change the granularity of the frames does not modify Your beliefs for those propositions that are doxastically equivalent.

Let Cr_1 and Cr_2 represent Your belief on {E, G, M} and {E, G, FT, FP, I}, respectively. By Doxastic Consistency, $Cr_1(E) = Cr_0(E)$, $Cr_1(M) = Cr_0(F\cup I)$, etc... and in fact Cr_1 is entirely defined from Cr_0. Identically, $Cr_2(E) = Cr_0(E)$, $Cr_2(FT\cup FP) = Cr_0(F)$, ... but some values of Cr_2 are not derivable from Cr_0 by Doxastic Consistency: this is the case for $Cr_2(FT)$, $Cr_2(FT\cup E)$... Hence extra requirements will be introduced. □

Formally, we have the next definitions.

Coarsening: Let $(\Omega, \mathfrak{R})$ be a propositional space. A coarsening C is a mapping from $\mathfrak{R}$ to $\mathfrak{R}''$, where $\mathfrak{R}''$ is an algebra also defined on Ω, such that one or several atoms of $\mathfrak{R}$ are mapped into one atom of $\mathfrak{R}''$ and each atom of $\mathfrak{R}$ is mapped into one and only one atom of $\mathfrak{R}''$.

Refinement: Let (Ω, $\Re$) be a propositional space. A refinement. R is a mapping from $\Re$ to $\Re'$ where $\Re'$ is an algebra on Ω' such that each atom of $\Re$ is mapped into one or several atoms of $\Re'$ and each atom of $\Re'$ is derived from one and only one atom of $\Re$. Let R(A) be the image of $A \in \Re$ in $\Re'$, and let $R(\emptyset)=\emptyset$.

The nature of the frames Ω and Ω' is irrelevant to our presentation. The only important components are the algebras. In practice, we can always redefine Ω and Ω' such that the resulting frames are equal.

Coarsenings and refinements are called uninformative if they do not modify the evidential corpus EC_t^Y held by You at t. Uninformative change fits in with the idea that only the structure of the algebra on which beliefs are held is modified; no further information is added to the evidential corpus.

The uninformative nature of the changes is formalized in the next requirement that states that the credibility function induced by such mappings from an initial credibility function Cr depends only on Cr and the mapping.

Requirement 5: Let (Ω, $\Re$, Cr, EC_t^Y) be a belief state. Let R be a refinement from (Ω, $\Re$) to (Ω, $\Re'$) and let C be a coarsening from (Ω, $\Re$) to (Ω, $\Re''$). Let the belief states (Ω, $\Re'$, Cr', EC_t^Y) and (Ω, $\Re''$, Cr'', EC_t^Y). Then Cr' and Cr'' are completely determined by Cr and by R and C, respectively. So there are g and h functions such that:

$$Cr' = g(Cr, R) \qquad \text{and} \qquad Cr'' = h(Cr, C).$$

3.3.1. Uninformative Coarsening.

The derivation of the nature of the h transformation is immediate as illustrated in example 2. The only difference between Cr'' and Cr is that Cr provides more detailed information on Ω than Cr''. Indeed Cr describes a belief over an algebra $\Re$ whose granularity is finer. The next theorem is proved by the direct application of the Doxastic Consistency Requirement.

Theorem 1. Let (Ω, $\Re''$, Cr'', EC_t^Y) be the belief state derived from the belief state (Ω, $\Re$, Cr, EC_t^Y) by the uninformative coarsening C from (Ω, $\Re$) to (Ω, $\Re''$). Given Requirements 1, 2 and 5,

$$Cr''(A) = Cr(C^{-1}(A)) \qquad \text{for all } A \in \Re'' \qquad (3.2)$$

where $C^{-1}(A)$ denote the union of the atoms of $\Re$ which are mapped by C into an atom of A.

3.3.2. Uninformative Refinement.

In requirement 5, some of the values of Cr' are derived by the direct application of the Doxastic Consistency Requirement. But this does not work for the elements of $\Re$' that are not the image of some elements of $\Re$ under R. They will be deduced once we study the conditioning process to which we now turn our attention. We only need one requirement: it states that if an atom of $\Re$ is refined into a very large number of new atoms in $\Re$' by the refinement R, then the credibility given to this new atoms should be very small.

Requirement 6: Let $(\Omega, \Re_n, Cr_n, EC_t^Y)$ be the belief state derived from the belief state $(\Omega, \Re, Cr, EC_t^Y)$ by the uninformative refinement R_n from $(\Omega, \Re)$ to $(\Omega, \Re_n)$, where R_n is so defined that it refines a given atom ω of $\Re$ into n atoms of $\Re_n$. Let ω_n be one of the atoms of $\Re_n$ that belongs to $R_n(\omega)$. Then:

$$\lim_{n\to\infty} Cr_n(\omega_n) = 0.$$

3.4. Informative changes of $\Re$: Conditioning.

Let EC_t^Y be the evidential corpus held by You at time t and let $(\Omega, \Re, Cr, EC_t^Y)$ be Your belief state. Suppose You revise EC_t^Y by adding to it the piece of evidence Ev_A where Ev_A is the proposition: 'all worlds in $\overline{A}\subseteq \Omega$ are impossible'. How do You revise Your beliefs, hence Cr, after adding Ev_A to EC_t^Y? For simplicity's sake, we assume that Ev_A is compatible with EC_t^Y. Generalization for Ev_A not compatible with EC_t^Y is possible but useless here. The fact that Ev_A is compatible with EC_t^Y implies that we are restricting ourselves to the expansion process (Gärdernfors, 1988), i.e., to the conditioning process usually described in probability theory.

Let Cr_A be the credibility function (qualified as conditional) that results from the adjunction of Ev_A to EC_t^Y. It is postulated that Cr_A is completely determined by the credibility function Cr based on EC_t^Y and by A.

Requirement 7: Markovian Axiom. Let the belief state $(\Omega, \Re, Cr_0, EC_0)$. For $A\in\Re$, let Ev_A be a piece of evidence compatible with EC_0. Let $EC_A = EC_t^Y\cup\{Ev_A\}$ be the evidential corpus obtained by adding Ev_A to EC_0. Let $(\Omega, \Re, Cr_A, EC_A)$ be the belief state after Ev_A has been added to EC_0. It is assumed that Cr_A is completely determined by Cr_0 and A.

To derive the conditioning process, we use the idea of iterated conditioning. Let $A,B\subseteq \Omega$, and the three pieces of evidence Ev_A, Ev_B and $Ev_{A\cap B}$. Suppose You learns 1) Ev_A and then Ev_B, or 2) Ev_B and then Ev_A, or 3) directly $Ev_{A\cap B}$. In order to satisfy Doxastic Consistency, the final conditional credibility functions must be the same in the three cases. This is obtained by accepting that the order with which pieces of evidence are taken into consideration is irrelevant. Furthermore, the Doxastic Consistency Requirement allows us also to prove that the conditional credibility function Cr_A depends only on a few terms of Cr.

Theorem 2: Let the belief state $(\Omega, \Re, Cr, EC_t^Y)$. For $A \in \Re$, let Ev_A be a piece of evidence compatible with EC. Let $(\Omega, \Re, Cr_A, EC_t^Y \cup \{Ev_A\})$ be the new belief state obtained after conditioning the previous belief state on Ev_A. Then Cr_A satisfies:

1: $Cr_A(B) = 0$ for all $B \subseteq \overline{A}$, $B \in \Re$

2: $Cr_A(B) = Cr_A(B \cap A)$ for all $B \in \Re$

and there is an f function such that for all $B \in \Re$

3: $Cr_A(B) = f\,(Cr(B \cap A), Cr(\overline{B} \cap A), Cr(\overline{A}), Cr(A), Cr(B \cup \overline{A}), Cr(\overline{B} \cup \overline{A}), Cr(\Omega))$.

We introduce the principle of doxastic stability through example 2.

Example 2 Continued.

Consider the refinement of example 2. Let Cr_0 and Cr_2 denote the credibility function defined on {E, G, F, I} and {E, G, FT, FP, I}, respectively. Two types of conditionings, called generic and factual conditionings, can be considered (Dubois and Prade, 1994). The first results from an information relative to a set of worlds to which the actual world belongs, the second from an information relative to the actual world itself. For the generic conditioning, suppose You learn that there was no FrenchTuc (FT) at the place where the killing occurred. For the factual conditioning, suppose You had a perfect witness who can only recognize if someone is FT or not. The witness saw the killer, and states that the killer is not FT. Are these two types of conditioning equivalent? As far as You are concerned, they are. Both state that the killer is not FT. For the factual conditioning, the situation would have been different if the killer had been randomly selected from a population and You had learned that the killer was not FT, a good reason for such an event being that the killer is not French, and the resulting probabilistic analysis would be appropriate. Here we are not concerned with a randomly selected killer, but with one killer. And the two pieces of conditioning information are equivalent as far as Your beliefs about the killer's nationality are concerned.

Once You know that the killer is not a FrenchTuc, 'the killer is French' and 'the killer is FrenchPic' are doxastically equivalent. Similarly, Your belief that the killer is German was not affected by the refinement of the French into FP and FT: neither was it affected by the knowledge that the killer was not FT. So: $Cr_{notFT}(G) = Cr_2(G) = Cr_0(G)$.

Similarly, that the killer is 'German or FrenchPic' is doxastically equivalent to the fact that the killer is 'German or French' (as French and FrenchPic are doxastically equivalent once You know the killer is not FrenchTuc). So: $Cr_{notFT}(\{G,FP\}) = Cr_2(\{G, F\}) = Cr_0(\{G, F\})$.

These equalities are natural but must nevertheless be assumed. Formally, we have the next requirement.

Requirement 8: Doxastic Stability.
Let the belief state (Ω, $\Re$, Cr, EC_0). Let R be an uninformative refinement from $\Re$ to $\Re'$. Let ω be an atom of $\Re$, and $R(\omega) = A \cup B$ where $A \cap B = \emptyset$, $A \neq \emptyset$, $B \neq \emptyset$. Let $Ev_{\overline{\mathbf{B}}}$ be the piece of evidence that states that all atoms in B are impossible and let $EC_1 = EC_0 \cup \{Ev_{\overline{\mathbf{B}}}\}$, so $R(\overline{\omega}) \cup A =_{EC_1} \Omega$. Then under EC_1, $R(X) \cap \overline{B}$ and $R(X)$ are doxastically equivalent for every X in $\Re$: $R(X) \cap \overline{B} =_{EC_1} R(X)$.

Gärdenfors (1988) suggests two compelling properties for probabilistic revision functions: homomorphism and preservation, whose meanings are illustrated hereafter.

Example 1 Continued. In the horse race example, suppose that You learn that Carol is a sure loser. You can derive the conditional credibility function either directly from the combined credibility function Cr_{12} or from the linear combination of the individual credibility functions Cr_1 and Cr_2. This requirement would have been satisfied in probability theory if probabilities had not been normalized, i.e., if the axiom $P(\Omega) = 1$ were abandoned, and the Bayesian conditioning rule were $P(A|B) = P(A \cap B)$ instead of $P(A \cap B)/P(B)$. □

Requirement 9: Homomorphism:

If $Cr = pCr' + (1-p)\, Cr''$, $p \in [0,1]$, then $Cr_A = pCr'_A + (1-p)\, Cr''_A$ for any $A \in \Re$.

The Preservation Requirement asserts essentially that a proposition as much believed as a tautology will be as believed as the conditioning proposition after conditioning.

Example 1 Continued. Consider the horse race example involving four horses: Allan, Blues, Carol and Daisy. Suppose You learn that Daisy is a sure loser. Then {Allan, Blues, Carol} and{Allan, Blues, Carol, Daisy} are Doxastically Equivalent, hence Cr({Allan, Blues, Carol}) = Cr({Allan, Blues, Carol, Daisy}). Then if You also learn that Carol is a sure loser, then {Allan, Blue, Carol}, {Allan, Blue} and {Allan, Blue, Daisy} are Doxastically Equivalent, hence $Cr_{\text{not-Carol}}$({Allan, Blues, Carol}) = $Cr_{\text{not-Carol}}$({Allan, Blues, Daisy}). Furthermore, once I know that Daisy is a sure loser, no new information about other sure losers could change this knowledge, hence Cr(Daisy) was nul and cannot become positive by learning that Carol is a sure loser. □

Requirement 10: Preservation:

If $Cr(B) = Cr(\Omega)$, then $Cr_A(B) = Cr_A(A)$ for any $A,B \in \Re$,

and if $Cr(B) = Cr(\Omega)$ and $Cr(\overline{B}) = 0$, then $Cr_A(\overline{B}) = 0$ for any $A,B \in \Re$.

Given requirements 1 to 10, we can establish the exact mathematical relations that represent the impact of both the conditioning and the coarsening processes.

Theorem 3. Let the belief state $(\Omega, \mathfrak{R}, Cr, EC)$. Let R be an uninformative refinement from $(\Omega, \mathfrak{R})$ to $(\Omega, \mathfrak{R}')$. Let Cr' be the credibility function derived from Cr on $\mathfrak{R}'$ by R. For $A\in\mathfrak{R}$, let Cr_A be the conditional credibility function induced from Cr by the evidence Ev_A. The only solutions for the coarsening and conditioning processes that satisfy Requirements 1 to 10 are respectively:

$$Cr'(X) = \max_{Y:R(Y)\subseteq X} Cr(Y) \qquad \text{for all X in } \mathfrak{R}'.$$

$$Cr_A(B) = Cr(B\cup\overline{A}) - Cr(\overline{A}) \qquad \text{for } A,B\in\mathfrak{R} \tag{3.3}$$

3.5. Why probability functions and plausibility functions are inadequate?

Before going on to prove that all credibility functions are belief functions, we consider some of the consequences of requirements 1 to 10, and in particular why probability functions and plausibility functions are inadequate to represent quantified beliefs.

To show that probability functions are not adequate, we consider the problem of iterated uninformative refinements. As an illustrative example, take $\Omega_0 = \{a, b\}$, $\Omega_1 = \{a, b_1, b_2\}$, and $\Omega_2 = \{a, b_1, b_{21}, b_{22}\}$. Let R_1 be a refinement from $(\Omega_0, 2^{\Omega_0})$ to $(\Omega_1, 2^{\Omega_1})$ such that $R_1(a) = \{a\}$, and $R_1(b) = \{b_1, b_2\}$. Let R_2 be a refinement from $(\Omega_1, 2^{\Omega_1})$ to $(\Omega_2, 2^{\Omega_2})$ such that $R_2(a) = \{a\}$, $R_2(b_1) = \{b_1\}$ and $R_2(b_2) = \{b_{21}, b_{22}\}$.

Let the belief state $(\Omega_0, 2^{\Omega_0}, Cr_0, EC_0)$. Let Cr_1 (Cr_2) be the credibility function induced from Cr_0 (Cr_1) on 2^{Ω_1} (2^{Ω_2}) by the uninformative refinement R_1 (R_2).

Consider the refinement R_{12} from $(\Omega, 2^{\Omega})$ to $(\Omega_2, 2^{\Omega_2})$ such that $R_{12}(a) = \{a\}$, $R_{12}(b) = \{b_1, b_{21}, b_{22}\}$, and let Cr_{12} be the credibility function induced from Cr_0 on 2^{Ω_2} by the uninformative refinement R_{12}. R_{12} is nothing but the result of combining R_1 with R_2. By the Doxastic Consistency Requirement, $Cr_2 = Cr_{12}$.

In order to achieve such an equality in probability theory, we need to know how $Cr_0(b)$ is distributed among b_1 and b_2, and how $Cr_1(b_2)$ is distributed among b_{21} and b_{22}. That knowledge contradicts the Markovian Requirement that states that Cr_1 should depend only on Cr_0 and R_1, not on some extra information like the distributions of $Cr_0(b)$ between b_1 and b_2. The Markovian Requirement can only be satisfied if $Cr_0(b)$ is equally distributed between b_1 and b_2, in which case $Cr_1(b_2)$ should also be equally distributed between b_{21} and b_{22}. Thus $Cr_2(b_{21})$ would be equal to $Cr_0(b)/4$. The same rule applied to Cr_{12}, using R_{12}, implies that $Cr_{12}(b_{21}) = Cr_0(b)/3$, hence $Cr_{12} \neq Cr_2$, an inequality that contradicts the Doxastic Consistency Requirement. Hence equi-repartition cannot de defended. So probability functions are not adequate to represent beliefs once iterated uninformative refinements are applied.

The Preservation Requirement is not satisfied by plausibility functions, the dual of the belief functions. This rejection seems adequate because we feel that Cr should behave like the modality used to represent categorical beliefs, i.e., the ‘box’ operator encountered in doxastic logic. Using plausibility functions to represent quantified beliefs would be equivalent to representing categorical beliefs by the ‘diamond’ operator. Of course, such an interpretation of categorical ‘belief’ could be defended. The question is in defining what is meant by beliefs: we follow the classical interpretation described in doxastic logic (Hintikka, 1962).

In conclusion, probability functions are not expressive enough to satisfy our requirements, and plausibility functions do not cover our interpretation of the belief modality.

3.6. Credibility functions are belief functions.

That belief functions satisfy all requirements 1 to 10 is immediate. The problem is to prove the reverse. We prove it by studying the concept of deconditionalization, i.e., the inverse of the conditioning process, and adding a final requirement. Suppose You had some initial credibility function Cr defined on $\mathfrak{R}$ and You had conditioned it on Ev_X for $X \in \mathfrak{R}$, which resulted in the credibility function Cr_X. Then You learn that the conditioning on Ev_X was inappropriate, i.e., that all the reasons that lead You to condition on Ev_X were unjustified. You want to erase the impact of Ev_X from Cr_X and rebuild a credibility function Cr from which Cr_X could have been obtained by its conditioning on Ev_X.

Formally, let $(\Omega, \mathfrak{R})$ be a propositional space. Let $\mathtt{Cr}$ be the set of credibility functions defined on $\mathfrak{R}$. For $X \in \mathfrak{R}$, let $\mathtt{Cr}_X$ be the set of conditional credibility functions obtained by conditioning the elements of $\mathtt{Cr}$ on Ev_X by (3.3). The impact of conditioning the elements of $\mathtt{Cr}$ on X can be described by an operator $S_X : \mathtt{Cr} \to \mathtt{Cr}_X$ such that:

$$Cr_X = S_X \circ Cr \quad \text{for all } Cr \in \mathtt{Cr}$$

S_X is a linear operator and is uniquely represented by a matrix operator (Klawonn and Smets, 1992)

Consider now the deconditioning operators. The matrix S_X is a singular matrix, so it admits only generalized inverses. Let S^{-}_X be such an operator. S^{-}_X is a generalized inverse of S_X and satisfies:

$$S_X \circ S^{-}_X \circ S_X = S_X. \tag{3.4}$$

The relation translates the idea that re-conditioning after deconditioning annihilates the effect of the deconditioning. Besides S^{-}_X is also idempotent:

$$S^{-}_X \circ S^{-}_X = S^{-}_X \tag{3.5}$$

Indeed, deconditioning twice has the same impact as deconditioning once.

Given S_X, there are many operators S^{-}_X satisfying (3.4) and (3.5). Let $\mathfrak{S}^{-}_X$ be the set of such deconditioning operators.

Example 3: In order to explain the origin of the next requirement, suppose that Cr_X quantifies Your beliefs over $\Re$ based on an evidential corpus EC_t^Y that contains the conditioning evidence Ev_X for $X \in \Re$. You then learn that the evidence Ev_X was unjustified and its impact must be erased. Which operator S^{-}_X should You use? Suppose another agent You* has some opinion about which operator $S^{-}_X \in \mathfrak{S}^{-}_X$ is to be used by You. The opinion of You* is represented by a credibility over $\mathfrak{S}^{-}_X$. Suppose You* is sure about which $S^{-}_X \in \mathfrak{S}^{-}_X$ should be used by You to decondition Cr_X. Suppose You had no a priori about which operator is appropriate and You trust in You*. So You accept the opinion of You* that the appropriate operator is indeed $S^{-}_X{}^*$ and You use $S^{-}_X{}^*$ to decondition Cr_X. Of course, the result must be a credibility function over $\Re$. We want You* to be able to choose $S^{-}_X{}^*$ independently of the value Cr_X representing Your belief over $\Re$. Thus, for every Cr_X and every $S^{-}_X \in \mathfrak{S}^{-}_X$, $S^{-}_X \circ Cr_X$ must be a credibility function. This requirement is sufficient to prove that the credibility functions are belief functions. If Cr is not a belief function, then it is always possible to find a S^{-}_X so that $S^{-}_X \circ Cr_X$ allocates negative beliefs to some elements of $\Re$. □

The next requirement just formalizes the requirement detailed in example 3.

Requirement 11. Let $(\Omega, \Re)$ be a propositional space. Let $\mathtt{Cr}$ be the set of credibility functions defined on $\Re$. For $X \in \Re$, let $\mathtt{Cr}_X$ be the set of conditional credibility functions defined on $\Re$ after conditioning the credibility functions in $\mathtt{Cr}$ on the evidence Ev_X. Let $\mathfrak{S}^{-}_X$ be the set of operators deconditioning the elements of $\mathtt{Cr}_X$ on Ev_X. For every S^{-}_X in $\mathfrak{S}^{-}_X$ and every Cr_X in $\mathtt{Cr}_X$, one has:

$$S^{-}_X \circ Cr_X \in \mathtt{Cr}.$$

Requirements 1 to 11 imply that credibility functions are belief functions.

Theorem 4: Every function that satisfies requirements 1 to 11 is a belief function.

This concludes our task.

4. Conclusions.

In conclusion, we have justified the use of belief functions to represent quantified beliefs. The model we obtain corresponds to the transferable belief model, i.e., a model for the representation of beliefs based on belief functions and developed independently of any probabilistic assumption. This is to be contrasted with Dempster's model (Dempster, 1967) that is also based on belief functions, but they are strongly linked to some underlying probability function. Indeed, the belief function derived within Dempster’s model on some given space Y results from a one-to-many mapping between a space X and the space Y, and the existence of a probability measure on X. This probability measure imposes constraints that we have not included in our modelization.

Similar reasons hold for the random sets interpretation of belief functions. The model we develop does not require any idea of subjective probability. It is derived directly from general rationality principles unrelated to some underlying probability function, as is the case with the transferable belief model.

The value of the model we derive for representing quantified beliefs can be assessed by analyzing the validity of each requirement and assessing their adequacy. No objective test seems to exist to evaluate normative models for quantified beliefs. Hence the interest of the axiomatic justification we propose.

Acknowledgments.

The author is indebted to M. Daniel, R. Gilles, P. Hajek, R. Kennes, J. Kohlas, I. Kramoisil, H. Kyburg, B. Marchal, V. Poznanski, A. Saffiotti, T. Seidenfelt, F. Voorbraak and two anonymous referees for their remarks and comments.

Bibliography.

BRADLEY R. and SWARTZ N. (1979) Possible worlds. Basil Blackwell, Oxford, UK.

DEMPSTER A.P. (1967) Upper and lower probabilities induced by a multplevalued mapping. Ann. Math. Statistics 38: 325-339.

DUBOIS D. and PRADE H. (1994) A survey of belief revision and updating rules in various uncertainty models. Inter. J. Intelligent Systems, 9:61-100.

GÄRDENFORS P. (1988) Knowledge in flux. Modelling the dynamics of epistemic states. MIT Press, Cambridge, Mass.

HINTIKKA J. (1962) Knowledge and belief. Cornell Univ. Press, Ithaca, NY.

KLAWONN F. and SMETS Ph. (1992) The dynamic of belief in the transferable belief model and specialization-generalization matrices. in Dubois D., Wellman M.P., d’Ambrosio B. and Smets P. Uncertainty in AI 92. Morgan Kaufmann, San Mateo, Ca, USA, 1992, pg.130-137.

KYBURG H. (1987a) Objective probabilities. IJCAI-87, 902-904.

KYBURG H.E.Jr. (1987b) Bayesian and non-Bayesian evidential updating. Artificial Intelligence, 31:271-294.

KYBURG H.E.Jr. (1995) Set-Based Bayesianism. forthcoming in IEEE Tranmsactions on Systems, Man, and Cybernetics.
LEVI I. (1980) The enterprise of knowledge. MIT Press, Cambridge, Mass.
RUSPINI E.H. (1986) The logical foundations of evidential reasoning. Technical note 408, SRI International, Menlo Park, Ca.
SHAFER G. (1976) A mathematical theory of evidence. Princeton Univ. Press. Princeton, NJ.
SMETS Ph. (1991) Varieties of ignorance. Information Sciences. 57-58:135-144.
SMETS Ph. and KENNES R. (1994) The transferable belief model. Artificial Intelligence 66:191-234.
SMETS Ph; (1995) The axiomatic justication of the transferable belief model. Technical Report, IRIDIA-TR-95:8.
SMITHSON M. (1989) Ignorance and Uncertainty: Emerging Paradigms. Springer-Verlag, New York.
VOORBRAAK F. (1993) As Far as I Know: Epistemic Logic and Uncerttaiunty. Dissertation, Utrecht University.
WALLEY P. (1991) Statistical reasoning with imprecise probabilities. Chapman and Hall, London.
WONG S.K.M., YAO Y.Y., BOLLMANN P. and BÜRGER H.C. (1990) Axiomatization of qualitative belief structure. IEEE Trans. SMC, 21:726-734.