The transferable belief model for quantified belief representation.

Philippe Smets[1]
IRIDIA, Université Libre de Bruxelles, Brussels, Belgium.

Abstract. The transferable belief model is a model to represent quantified beliefs based on the use of belief functions, as initially proposed by Shafer. It is developed independently from any underlying related probability model. We summarize our interpretation of the model and present several recent results that characterize the model. We show how rational decision must be made when beliefs are represented by belief functions. We explain the origin of the two Dempster's rules that underlie the dynamic of the model through the concept of specialization and least commitment. We present the canonical decomposition of any belief functions, and discover the concept of 'debt of beliefs'. We also present the generalization of the Bayesian Theorem to belief functions.

1. Introduction.

We present the transferable belief model (TBM), a model for the representation of quantified beliefs. The model aims in representing the same concept as the Bayesian model, i.e., the graded dispositions that guide 'our' behavior. We use the word 'belief' in a broad sense. It could be replaced by quantified credibility, subjective support, strength of opinion... These beliefs are not categorical as in modal logic, but admits degrees as in probability theory. Our approach is normative. The beliefs are held by an idealized rational agent, denoted by You. This 'You' can be a human, but also a robot, a computer program...

In order to construct a model of quantified beliefs, we must be able to represent the static beliefs held by You at a given time and the dynamic of these beliefs when new information is taken in consideration. Decisions are the only observable outcome of the belief process, so we must also explain how decisions will be made. Classically, the probabilistic model, in particular the Bayesian model, is the most popular and best developed model for the representation of quantified belief. This model has nevertheless some limitations. The transferable belief model (Smets and Kennes, 1994) is an alternative model that resolves these weaknesses. It is based on the belief functions that were introduced by Shafer (1976). This paper focused on the specificity of the TBM. It

[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.

should be understood as a more or less self-contained complement to Shafer's initial work (Shafer, 1976) and of Smets' presentation (Smets, 1988).

In Section 2, we present the concept of actual world and the domain on which beliefs are held.

In section 3, we present the TBM, with its two level mental model. We summarize its major mathematical properties. We present the concept of least commitment that would correspond within the TBM to the maximum entropy principle described in probability theory. We present the concept of specialization that leads to the justification of both Dempster's rule of conditioning and Dempster's rule of combination, as well as to a definition of the concept of distinctness. We proceed with the canonical decomposition of a belief function, a still immature problem but that opens the doors to the idea of debts of beliefs that might represent the notion of 'good reasons NOT to believe'.

In section 4, we explain how decisions can be achieved when beliefs are represented by belief functions, and when decisions must be obtained through the use of probability functions in order to avoid Dutch Books, and other irrational behaviors. We show how the TBM resists to the Dutch Books criticisms. It is achieved essentially by not assuming some temporal coherence principle that underlies the Dutch Books argument in the classical probabilistic derivation.

In section 5, we develop the Generalized Bayesian Theorem, a theorem that will play in the transferable belief model the same role as the classical Bayesian Theorem does in probability theory. It will be central to solve diagnostic problems, solving nicely the problem of the too often arbitrary a priori probability functions classically required.

Section 6 discusses the specificity of the TBM versus its major contender, the probabilistic model.

In section 7 we conclude and illustrate the potential use of the model.

Our presentation is based on examples that intuitively characterize the properties of the model. We focus on understanding. Formal presentations and proofs can be found in Shafer (1976) and Smets (1990a, 1990b, 1993b, 1993c, 1995).

2. The credibility domain.

In this section, we present the concept of actual world and the domain on which beliefs are held. We shall use expressions like 'Your belief that a proposition A is true is .7', or as a shortcut, 'Your belief in A is .7'. The meaning of the expression 'Your belief at time t_0 that it will rain the next day is .7' can be either: 'the measure of the belief held by You at time t_0 that the proposition "it will rain tomorrow" is true is .7' or: 'You believe at level

.7 that tomorrow belongs to the set of rainy days'. Beliefs can be equivalently given to 'propositions' or to the subsets of worlds that denote the propositions. We use the possible worlds approach (Carnap, 1962, Ruspini, 1986, Bradley and Swartz, 1979). Beliefs will be given to sets of worlds. In probability theory, these subsets are called 'events'.

2.1. The propositional space.

Let 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 L. Propositions identify the subsets of Ω, and the subsets of Ω denote propositions. For a proposition P, let $[\![P]\!] \subseteq \Omega$ be the set of worlds identified by P (i.e., those worlds where P is true). The worlds of Ω are built in such a way that no two worlds denote the same proposition.

We assume that among the worlds of Ω a particular one, denoted ω_0, corresponds to the actual world. You ignore at time t which world is ω_0. You can only express Your belief at t that the actual world ω_0 belongs or not to various subsets of worlds of Ω. The value bel(A) denotes the degree of belief hold by You at time t that the actual world ω_0 belongs to the set $A \subseteq \Omega$ (or equivalently that the proposition denoted by A is true in the actual world).

Beliefs are not always expressed for every subset of Ω. Thanks to the expressive power of L, it may occur that the worlds of Ω denote very precise propositions, and due to Your limited understanding or interest, You cannot or do not express Your beliefs on such a detailed domain. When You want to assess Your belief about tomorrow weather in Brussels, You will not asses Your belief on the weather at every point of the earth. You will restrict Yourself to Brussels even though L could contains the propositions: 'Brussels weather is fine', 'New York weather is fine', 'Tokyo weather is fine', ... When asked about Your belief about Brussels weather, You build Your belief on a domain $\Re$ built from the two propositions: 'Brussels weather is fine', and 'Brussels weather is not fine'. You will not express Your belief that 'Brussels weather is fine and New York weather is not fine and Tokyo weather is fine', etc... You just do not care about such a refined domain.

So for a given question, at a given time, You build a partition of Ω. We call 'atom' the elements of this partition. Let $\Re$ be the Boolean algebra of subsets of Ω built from these atoms: so $\Re$ contains the unions of the atoms and is closed under complement, union and intersection. Let $At(\Re)$ denote this set of atoms. When $\Re$ is the power set 2^Ω of Ω, the atoms of $\Re$ are the singletons of Ω. Your belief is given to the elements of $\Re$.

We call Ω the frame of discernment (the frame for short). We call the pair $(\Omega, \Re)$ a propositional space. We call the triple $(\Omega, \Re, \text{bel})$ a belief state, where bel is the belief function that represents the beliefs held by You at time t on $\Re$. The full specification of a belief state should also mention the evidential corpus on which it is based (see chapter 1), but we omit it for simplicity's sake.

For the full description of the TBM , we introduce a subset of Ω on which the algebra and the belief are defined. This distinction is not detailed here. Details are given in (Smets, this handbook series, volume 2).

2.2. The belief functions.

Belief functions (Shafer, 1976) are capacities monotone of order ∞ (Choquet, 1953). Let $(\Omega, \Re)$ be a finite propositional space. A belief function is a function bel from $\Re$ to [0, 1] such that :

1) $\text{bel}(\emptyset) = 0$
2) for all $A_1, A_2, \ldots A_n \in \Re$,

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

Usually, $\text{bel}(\Omega) = 1$ is also assumed. It can be ignored. We only require that $\text{bel}(\Omega) \leq 1$.

The basic belief assignment (bba) related to a belief function bel is the function m from $\Re$ to [0,1] such that :

$$m(A) = \sum_{B: B \in \Re, \emptyset \neq B \subseteq A} (-1)^{|A|-|B|} . \text{bel}(B) \qquad \text{for all } A \in \Re, A \neq \emptyset \qquad (2.2)$$

$$m(\emptyset) = 1 - \text{bel}(\Omega)$$

The value $m(A)$ for $A \in \Re$ is called the basic belief mass (bbm) given to A[2]. It may happen that $m(\emptyset) > 0$, what reflects some kind of contradiction in the belief state, but $\text{bel}(\emptyset)$ is always 0.

The m and bel functions are in one-to-one correspondence via:

$$\text{bel}(A) = \sum_{B: B \in \Re, \emptyset \neq B \subseteq A} m(B) \qquad \text{for all } A \in \Re, A \neq \emptyset \qquad (2.3)$$

Related to bel and m, one can also define the plausibility function pl and the commonality function q, both from $\Re$ to [0,1], by:

$$\text{pl}(A) = \text{bel}(\Omega) - \text{bel}(\bar{A}) \qquad \text{for all } A \in \Re \qquad (2.4)$$

---

[2] Shafer speaks of basic *probability* masses and assignment. We avoid the probability label as it induces the idea that there is some underlying probability function, what is not the case in the transferable belief model, in contrast with the model studied by Dempster (1967) and those defended today by Shafer (1992).

$$q(A) = \sum_{B:B\in\Re, A\subseteq B} m(B) \qquad \text{for all } A\in\Re,\ A\neq\emptyset \qquad (2.5)$$

The meaning of these functions will be clarified in the following sections.

Shafer assumes that bel is normalized so that bel(Ω) = 1, or equivalently pl(Ω) = 1 and m(Ø) = 0. We do not require such a normalization (Smets, 1992a). We use the notation bel and pl, whereas Shafer uses Bel and Pl, in order to enhance that our functions are unnormalized.

3. The transferable belief model.

3.1. Two mental levels.

Beliefs manifest themselves at two mental levels: the credal level where beliefs are entertained and the pignistic level where beliefs are used to make decisions[3].

Usually these two levels are not distinguished and probability functions are used to quantify beliefs at both levels. The justification for the use of probability functions is usually linked to "rational" behavior to be held by an ideal agent involved in some betting or decision contexts (Ramsey, 1931, Savage, 1954, DeGroot, 1970). They have shown that if decisions must be "coherent", the uncertainty over the possible outcomes must be represented by a probability function. This result is accepted here, except that such *probability functions quantify the uncertainty only when a decision is really involved.* Therefore uncertainty must be represented by a probability function at the pignistic level. We also accept that this probability function is induced from the beliefs entertained at the credal level. What we reject is the assumption that this probability function represents the uncertainty at the credal level.

We assume that the pignistic and the credal levels are distinct which implies that the justification for using probability functions at the credal level does not hold anymore (Dubois et al., 1995). At the credal level, beliefs are represented by belief functions, at the pignistic level, they induce a probability function that is used to make decision. This probability function should not be understood as representing Your beliefs, it is nothing but the additive measure needed to make decision, i.e., to compute the expected utilities. Of course this probability function is directly induced by the belief function representing Your belief at the credal. The link between the two levels is achieved by the pignistic transformation that transforms a belief function into a probability function. Its nature and justification is detailed in section 4.

[3] Credal and pignistic derive both from the latin words ‘credo’, I believe and ‘pignus’, a wage, a bet (Smith, 1961).

3.2. Basic belief assignment and degree of belief.

The basic belief assignment receives a natural interpretation. For $A \in \Re$, m(A) is that part of Your belief that supports A, i.e., that the actual world ω_0 belongs to A, and that, due to lack of information, does not support any strict subset of A.

Let m: $\Re \rightarrow [0,1]$ with

$$\sum_{A \in \Re} m(A) = 1.$$

In general, the basic belief assignment looks similar to a probability distribution function defined on the power set 2^Ω of the frame of discernment Ω. This analogy led several authors to claim that the transferable belief model is nothing but a probabilistic model on 2^Ω. Such an interpretation does not resist once conditioning is introduced, as far as it does not lead to Dempster's rule of conditioning we derive in section 3 (Smets, 1992b).

Example 1. Let us consider a somehow reliable witness in a murder case who testifies to You that the killer is a male. Let $\alpha = .7$ be the reliability You give to the testimony. Suppose that *a priori* You have an equal belief that the killer is a male or a female. A classical probability analysis would compute the probability P(M) of M = 'the killer is a male' given the witness testimony as:

$$\begin{aligned} P(M) &= P(M|\text{Reliable})\ P(\text{Reliable}) + P(M|\text{Not Reliable})\ P(\text{Not Reliable}) \\ &= 1.\ x\ .7 + .5\ x .3\ = .85 \end{aligned}$$

where Reliable and Not Reliable refer to the witness reliability. The value .85 is the sum of the probability that the witness is reliable (.7) plus the probability of M given the witness is not reliable (.5) weighted by the probability that the witness is not reliable (.3). The transferable belief model analysis will give a belief .7 to M: bel(M) = .7. In P(M) = .7 + .15, the .7 value can be viewed as the *justified* component of the probability given to M (called the belief or the support) whereas the .15 value can be viewed as the *aleatory* component of that probability. It would be relevant to bel(M) only if the murderer had been really selected by a random process from a population where 50% are male. In our example, such a random selection does not apply, so the aleatory component is not considered when building Your beliefs. The transferable belief model deals only with the justified components. (Note: the Evidentiary Value Model (Ekelof, 1982, Gärdenfors et al., 1983, Smets, 1994) describes the same belief component, within a strict probability framework. It differs from the transferable belief model once conditioning is introduced.)

If some further evidence becomes available to You and implies that B is true, then the mass m(A) initially allocated to A is transferred to $A \cap B$.

Continuing with the murder case, suppose there are only two potential male suspects: Phil and Tom. Then You learn that Phil is not the killer. The testimony now supports that the killer is Tom. The reliability .7 You gave to the testimony initially supported 'the killer is Phil or Tom'. The new information about Phil implies that the value .7 now supports 'the killer is Tom'. t

More formally, given a propositional space $(\Omega, \Re)$, the degree of belief bel(A) for $A \in \Re$ quantifies the total amount of *justified specific support* given to A. It is obtained by summing all basic belief masses given to subsets $X \in \Re$ with $X \subseteq A$ (and $X \neq \emptyset$). Indeed a part of belief that supports that the actual world ω_0 is in B also supports that ω_0 is in A whenever $B \subseteq A$. So for all $A \in \Re$,

$$\mathrm{bel}(A) = \sum_{\emptyset \neq X \subseteq A, X \in \Re} m(X). \tag{3.1}$$

We say *justified* because we include in bel(A) *only* the basic belief masses given to subsets of A. For instance, consider two distinct atoms x and y of $\Re$. The basic belief mass $m(\{x,y\})$ given to $\{x,y\}$ could support x if further information indicates this. However given the available information the basic belief mass can only be given to $\{x,y\}$. We say *specific* because the basic belief mass $m(\emptyset)$ is not included in bel(A) as it is given to the subset $\emptyset$ that supports not only A but also $\overline{A}$.

The degree of plausibility pl(A) for $A \in \Re$ quantifies the maximum amount of *potential specific support* that could be given to A. It is obtained by adding all those basic belief masses given to subsets X compatible with A, i.e., such that $X \cap A \neq \emptyset$:

$$\mathrm{pl}(A) = \sum_{X \cap A \neq \emptyset, X \in \Re} m(X) = \mathrm{bel}(\Omega) - \mathrm{bel}(\overline{A}) \tag{3.2}$$

We say *potential* because the basic belief masses included in pl(A) could be transferred to non-empty subsets of A if new information could justify such a transfer. It would be the case if we learn that $\overline{A}$ is impossible.

The plausibility function pl is just another way of presenting the information contained in bel and could be forgotten, except inasmuch as it often provides a mathematically convenient alternate representation of the beliefs.

3.3. Vacuous belief function.

Total ignorance is represented by a vacuous belief function, i.e. a belief function such that $m(\Omega) = 1$, hence $\mathrm{bel}(A) = 0\ \forall A \in \Re$, $A \neq \Omega$, and $\mathrm{bel}(\Omega) = 1$. The origin of this particular quantification for representing a state of total ignorance can be justified. Suppose that there are three propositions labeled A, B and C, and You are in a state of total ignorance about which is true. You only know that one and only one of them is true but even their content is unknown to You. You only know their number and their label. Then You have no reason to believe any one more than any other, hence, Your beliefs about their truth are equal: $bel([\![A]\!]) = bel([\![B]\!]) = bel([\![C]\!]) = \alpha$ for some $\alpha \in [0,1]$. Furthermore, You have no reason to put more (or less) belief in $[\![A]\!] \cup [\![B]\!]$ than in $[\![C]\!]$: $bel([\![A]\!] \cup [\![B]\!]) = bel([\![C]\!]) = \alpha$ (and similarly $bel([\![A]\!] \cup [\![C]\!]) = bel([\![B]\!] \cup [\![C]\!]) = \alpha$). The vacuous belief function is the only belief function that satisfies equalities like: $bel([\![A]\!] \cup [\![B]\!]) = bel([\![A]\!]) = bel([\![B]\!]) = \alpha$. Indeed the inequalities (3.1) imply that $bel([\![A]\!] \cup [\![B]\!]) \geq bel([\![A]\!]) + bel([\![B]\!]) - bel([\![A]\!] \cap [\![B]\!])$. As $[\![A]\!] \cap [\![B]\!] = \emptyset$, $bel([\![A]\!] \cap [\![B]\!]) = 0$. The inequality becomes $\alpha \geq 2\alpha$ where $\alpha \in [0,1]$, hence $\alpha = 0$.

3.4. The Principle of Minimal Commitment.

Example 2. Let $\Omega = \{a, b, c\}$. Suppose You know only that My[4] belief function over Ω is such that $bel_{Me}(\{a\}) = .3$ and $bel_{Me}(\{b,c\}) = .5$, and You do not know the value I give to bel_{Me} for the other subsets of Ω. Suppose You have no other information on Ω and You are ready to adopt My belief as Yours. How to build Your belief given these partial constraints? Many belief functions can satisfy them. If You adopt the principle that the subsets of Ω should not receive more support than justified, then Your belief on Ω will be such that $m_{You}(\{a\})=.3$, $m_{You}(\{b,c\})=.5$ and $m_{You}(\{a,b,c\})=.2$. Among the belief functions compatible with the constraints given by the known values of bel_{Me}, bel_{You} is the one that gives the smallest degree of belief to every subsets of Ω. The principle evoked here is called the Principle of Minimal Commitment. It fits in with the idea that degrees of belief are degrees of 'justified' supports and You should never give more belief than justified. □

With un-normalized belief functions, the principle definition is essentially based on the plausibility function. Suppose pl_1 and pl_2 are two plausibility functions defined on $\Re$ such that:

$$pl_1(A) \leq pl_2(A) \qquad \forall A \in \Re. \tag{3.3}$$

Then we say that pl_2 (bel_2, m_2) is not more committed than pl_1 (bel_1, m_1) (and less committed if there is at least one strict inequality). When expressed with belief functions, the principle becomes:

$$bel_1(A) + m_1(\emptyset) \geq bel_2(A) + m_2(\emptyset) \qquad \forall A \in \Re. \tag{3.4}$$

In particular the vacuous belief function is the least committed belief function among all belief functions on Ω.

The Principle of Minimal Commitment consists in selecting the least committed belief function in a set of equally justified belief functions. This selection procedure does not always lead to a unique solution in which case extra requirements are added. The

[4] 'I' (or 'Me') is an agent different from 'You'.

principle formalizes the idea that one should never give more support than justified to any subset of Ω. It satisfies a form of skepticism, of uncommitment, of conservatism in the allocation of our belief. In its spirit, it is not far from what the probabilists try to achieve with the maximum entropy principle (see Dubois and Prade 1987, Hsia, 1991, Smets 1993b)

3.5. Specializations.

The major rules that describe the dynamic of the beliefs in the TBM are Dempster's rule of conditioning and Dempster's rule of combination. In order to explain their origin, we introduce the concept of specialization, both rules being particular cases of the specialization process.

Let m_0 be the basic belief assignment induced on the propositional space $(\Omega,\Re)$ by You at time t_0. The value $m_0(A)$ is that part of Your belief that supports $A\in\Re$ and does not support any strict subset of A due to lack of information. If further information obtained by You at time t_1 with $t_1 > t_0$ justifies it, the basic belief mass $m_0(A)$ that was supporting $A\in\Re$ at t_0 might support more specific subsets of A. This fits in with the idea that m(A) was not allocated to subsets more specific than A by lack of information. When new information is obtained, m(A) might thus ‘flow’ to subsets of A, and it may not move outside of A as we already knew that it specifically supports A. Therefore, the impact of a new piece of evidence results in a redistribution of $m_0(A)$ among the subsets of A. This redistribution can be characterized by a set of non negative coefficients $s(B,A)\in[0,1]$, $A,B\in\Re$, where s(B,A) is the proportion of $m_0(A)$ that is transferred to $B\in\Re$ once the new piece of evidence is taken into account by You. The s coefficients depends of course on the piece of evidence that initiated the belief revision.

In order to conserve the whole mass $m_0(A)$ after this transfer, the s(B,A) must satisfy:

$$\sum_{B\subseteq A, B\in\Re} s(B,A) = 1 \qquad \forall A\in\Re \tag{3.5}$$

As masses can only flow to subsets, s(B,A) = 0 for all B not included in A. The matrix **S** of such coefficients s(B,A) for $A,B\in\Re$ is called a specialization matrix on $\Re$ (see Yager, 1986, Dubois and Prade, 1986, Kruse and Schwecke, 1990, Delgado and Moral, 1987).

After You learn the new piece of evidence E, Your initial basic belief assignment m_0 is transformed into the new basic belief assignment m_1 such that:

$$m_1(A) = \sum_{X\in\Re} s(A,X)\, m_0(X) \tag{3.6}$$

This formula reflects the idea that the bbm $m_0(X)$ initially allocated to X is distributed among the subsets of X after applying the specialization operator. This down-flow reflects the meaning of $m_0(X)$ as the part of belief that specifically supports X, but might support more specific subsets if further information justifies it.

The basic belief assignment m_1 is called a specialization of m_0. For a bba m, we use the same notation **m** to represent to column vector with elements m(A) for $A\in\Re$. Relation 3.6 can be written as:

$$\mathbf{m}_1 = \mathbf{S}\cdot\mathbf{m}_0 \tag{3.7}$$

Yager (1986) has shown that if the basic belief assignment m_1 is a specialization of the basic belief assignment m_0, then m_1 is at least as committed as m_0.

Example 3.1. An example of specialization.
Let $\Omega = \{a, b, c\}$. Suppose a bba m_0 defined on 2^Ω. Let **S** be a specialization matrix, and let $\mathbf{m}_1 = \mathbf{S}.\mathbf{m}_0$. Table 1 presents the values of **S**, $\mathbf{m}_0$ and $\mathbf{m}_1$. We use the iterated order illustrated in table 1 to list the elements of the vectors **m** and the specialization matrix **S**.

	m_1	=	S									m_0
			Ø	{a}	{b}	{a,b}	{c}	{a,c}	{b,c}	{a,b,c}		
Ø	.13		1.	.3	.5	.2	.0	.4	.0	.2		.0
{a}	.04			.7		.5		.1		.1		.0
{b}	.13				.5	.0			.4	.0		.1
{a,b}	.00	=				.3				.0	•	.0
{c}	.50						1.	.3	.0	.5		.3
{a,c}	.00							.2		.0		.0
{b,c}	.12								.6	.0		.2
{a,b,c}	.08									.2		.4

Table 1. values of the specialization matrix **S**, and of the bba $\mathbf{m}_0$ and $\mathbf{m}_1$ with $\mathbf{m}_1 = \mathbf{S}.\mathbf{m}_0$. The blanks in the **S** matrix indicate those values of **S** that must be null so that **S** is a specialization matrix.

3.5.1. Specialization and Dempster's rule of conditioning.
Suppose Your beliefs over $\Re$ at time t_0 are represented by a basic belief assignment m_0. At time t_1, You learn the piece of evidence Ev_A that says that 'the actual world ω_0 is not in $\overline{A}$' for $A\in\Re$. Let m_A be the basic belief assignment obtained after conditioning m_0 on A. The bbm $m_0(X)$ that was specifically supporting X now supports $X\cap A$, so after conditioning the bbm $m_0(X)$ is transferred to $m_A(X\cap A)$. This transfer explains the name of the model. The resulting bba and its related functions are given by:

$$m_A(B) = \sum_{X: X\in\Re, X\subseteq\overline{A}} m_0(X)$$

$$\begin{array}{lcll} bel_A(B) & = & bel_0(B\cup\overline{A}) - bel_0(\overline{A}) & \text{for } B\in\Re \\ pl_A(B) & = & pl_0(A\cap B) & \text{for } B\in\Re \\ q_A(B) & = & q_0(B) & \text{if } B\subseteq A, \\ & & 0 & \text{otherwise.} \end{array}$$

These rules are called the Dempster's rule of conditioning (except for the normalization factor that is neglected here).

The same results are obtained by a specialization of $\mathbf{m}_0$ where the coefficients of the specialization matrix $\mathbf{S}_A$ are:

$$s_A(C,B) = \begin{cases} 1 & \text{if } C = B \cap A \\ 0 & \text{otherwise.} \end{cases} \tag{3.8}$$

Then $\mathbf{m}_A = \mathbf{S}_A \cdot \mathbf{m}_0$.

We have shown (Klawonn and Smets, 1992) that $\mathbf{m}_A$ is in fact the least committed specialization of $\mathbf{m}_0$ such that $pl_A(\overline{A}) = 0$. Hence the rule for conditioning is derived from the concept of specialization and least commitment.

Example 3.2. Specialization and Dempster's rule of conditioning.
In table 2, we present the specialization matrix $\mathbf{S}_A$ that performs a conditioning on A = {a,c} when Ω = {a,b,c}, the bba $\mathbf{m}_0$ and $\mathbf{m}_A = \mathbf{S}_A.\mathbf{m}_0$. □

	$\mathbf{m}_A$	=				$\mathbf{S}_A$						$\mathbf{m}$
			Ø	{a}	{b}	{a,b}	{c}	{a,c}	{b,c}	{a,b,c}		
Ø	.1		1.	.0	1.	.0	.0	.0	.0	.0		.0
{a}	.0			1.		1.		.0		.0		.0
{b}	.0				.0	.0			.0	.0		.1
{a,b}	.0	=				.0				.0	·	.0
{c}	.5						1.	.0	1.	.0		.3
{a,c}	.4							1.		1.		.0
{b,c}	.0								.0	.0		.2
{a,b,c}	.0									.0		.4

Table 2. Values of the specialization matrix $\mathbf{S}_A$ that represents a conditioning on A = {a,c}, and of the bba $\mathbf{m}_0$ and $\mathbf{m}_A$ with $\mathbf{m}_A = \mathbf{S}_A.\mathbf{m}_0$.

The origin of the unnormalization and the positive bbm $m_A(\emptyset)$ can be understood by the study of Dempster's rule of conditioning. Suppose m_0 is a normalized belief function, i.e., $m_0(\emptyset) = 0$. Then You learn that the actual world does not belong to $\overline{A}$. Suppose You had given a non null belief to $\overline{A}$: $bel_0(\overline{A}) > 0$. The new piece of evidence enters in conflict with Your initial beliefs. You had some reason to believe that the actual world ω_0 belonged to $\overline{A}$, You learn that ω_0 does not belong to $\overline{A}$, hence the conflict. The intensity of the conflict is exactly $bel_0(\overline{A})$, and our conditioning rule transfers it to $m_A(\emptyset)$. So the value of $m_A(\emptyset)$ can be understood as the amount of conflict that results from the accumulation of the pieces of evidence that lead to Your present state of belief. In Smets (1992a) we also consider that a positive $m_A(\emptyset)$ could result from an incomplete specification of the frame of discernment on which bel_0 is built.

The meaning of the commonality function can also be understood when studying Dempster's rule of conditioning. It results from the relation:

$q_0(A) = m_A(A)$ for all $A \in \Re$.

In general, the value $m(\Omega)$ has a natural interpretation: it represents the part of belief completely free as $m(\Omega)$ could be given to any subset of Ω if further information justifies it. We called it the non-dogmatic component (Smets, 1983). When $\Re$ is based on two atoms, the term $m(\Omega)$ is also equal to pl(A) - bel(A), and often it has been understood as the amount of 'ignorance' included in the belief function bel. The term $m_A(A)$ computed by conditioning m_0 on A (thus $q_0(A)$) corresponds to the non dogmatic component of m_A, i.e., the part of belief completely free after having conditioned bel_0 on A. Besides the $q_0(A)$ values are also the eigenvalues of the operator that characterizes Dempster's rule of combination (see next section).

3.5.2. Specialization and Dempster's rule of combination

A more general form of specialization is considered. Suppose Your beliefs over $\Re$ are represented by the basic belief assignment m_0, hence m_0 results from the evidential corpus EC_0. Suppose two pieces of evidence Ev_1 and Ev_2 such that their conjunction is compatible with EC_0, i.e., their conjunction is not contradictory.

After the expansion of EC_0 by Ev_1, the initial basic belief assignment m_0 is transformed into a new basic belief assignment m_{01}. Let $\mathbf{S}_{01}$ be the specialization matrix that transforms $\mathbf{m}_0$ into $\mathbf{m}_{01}$. The impact of the second piece of evidence Ev_2 on EC_0 already expanded by Ev_1 results in the transformation of $\mathbf{m}_{01}$ into $\mathbf{m}_{012}$. Let $\mathbf{S}_{012}$ be the specialization matrix that transforms $\mathbf{m}_{01}$ into $\mathbf{m}_{012}$. In an identical way, we could consider expanding EC_0 first by Ev_2, and the result by Ev_1, with $\mathbf{m}_{02}$ and $\mathbf{m}_{021}$, $\mathbf{S}_{02}$, $\mathbf{S}_{021}$ the resulting basic belief assignment and specialization matrices.

So: $\mathbf{m}_{012} = \mathbf{S}_{012} \cdot \mathbf{m}_{01} = \mathbf{S}_{012} \cdot \mathbf{S}_{01} \cdot \mathbf{m}_0$

and $\mathbf{m}_{021} = \mathbf{S}_{021} \cdot \mathbf{m}_{02} = \mathbf{S}_{021} \cdot \mathbf{S}_{02} \cdot \mathbf{m}_0$

The resulting $\mathbf{m}_{012}$ and $\mathbf{m}_{021}$ must be equal as they both quantify the belief induced by the expansion of EC_0 by the same two pieces of evidence Ev_1 and Ev_2, and the order is assumed irrelevant (Gärdenfors, 1988, page 50)

Assume the specialization matrices that represent the impact of Ev_1 and of Ev_2 do not depend on the basic belief assignment on which they are applied. This translates the idea that the impacts of each of Ev_1 and Ev_2 are context independent. The requirement is thus:

$$\mathbf{S}_{01} = \mathbf{S}_{021} = \mathbf{S}_1 \text{ and } \mathbf{S}_{02} = \mathbf{S}_{012} = \mathbf{S}_2. \quad (3.9)$$

In that case, the equality $\mathbf{m}_{012} = \mathbf{m}_{021}$ becomes:

$$\mathbf{S}_2 \cdot \mathbf{S}_1 \cdot m_0 = \mathbf{S}_1 \cdot \mathbf{S}_2 \cdot m_0. \quad (3.10)$$

This constraint must be satisfied whatever $\mathbf{m}_0$. Hence the specialization matrices commute: $\mathbf{S}_1 \cdot \mathbf{S}_2 = \mathbf{S}_2 \cdot \mathbf{S}_1$.

Let Σ_D be the largest family of specialization matrices that contains the specialization matrices characterizing the conditioning process (see 3.5.1) and whose elements commute among themselves. The elements of Σ_D are called the dempsterian specialization matrices. In Klawonn and Smets (1992), we have shown that every dempsterian specialization matrices **S** is in one to one correspondence with a unique basic belief assignment m on Ω such that the coefficients of **S** satisfy:

$$s(B,A) = m_A(B) \qquad \text{for all } A,B\subseteq\Omega.$$

where the $m_A(B)$'s are obtained from m by Dempster's rule of conditioning.

Example 3.3. A dempsterian specialization matrix.
Let the bba m on 2^Ω with $\Omega = \{a,b,c\}$. The dempsterian specialization matrix $\mathbf{S}_m$ characterized by m is presented in table 3. The coefficients of a column of $\mathbf{S}_m$ are obtained by conditioning the bba m on the label of the column.

	m	$\mathbf{S}_m$: Ø	{a}	{b}	{a,b}	{c}	{a,c}	{b,c}	{a,b,c}
Ø	.05	1.0	.60	.44	.25	.20	.15	.09	.05
{a}	.04		.40		.19		.05		.04
{b}	.10			.56	.35			.11	.10
{a,b}	.01				.21				.01
{c}	.20					.80	.45	.35	.20
{a,c}	.15						.35		.15
{b,c}	.25							.45	.25
{a,b,c}	.20								.20

Table 3. Values of the dempsterian specialization matrix $\mathbf{S}_m$ built from the bba **m** given on the left of the table.

The meaning of that basic belief assignment m that characterizes the dempsterian specialization matrix **S** becomes clear once **S** is applied to a vacuous basic belief assignment, i.e., the basic belief assignment induced by an evidential corpus ECØ that leaves You in a state of total ignorance about the value of the actual world ω_0 in Ω. Let **S** be the dempsterian specialization matrix that characterizes the impact of a piece of evidence Ev. Expanding a ‘vacuous’ ECØ with Ev results in an evidential corpus that contains only Ev. The result of the specialization matrix **S** on the vacuous basic belief assignment is the bba m that characterizes **S**. Hence m is the basic belief assignment induced by the evidential corpus that contains only Ev (for what concerns Ω).

Let $\mathbf{S}_m$ denote the dempsterian specialization matrix built from the bba m. The result of the application of the dempsterian specialization matrix $\mathbf{S}_{m_2}$ on Ω to a basic belief assignment m_1 on Ω can be written without making $\mathbf{S}_{m_2}$ explicit. Let the basic belief assignment $m_{12} = \mathbf{S}_{m_2}\cdot m_1$. The basic belief mass $m_{12}(A)$ can be written as:

$$m_{12}(A) = \sum_{B,C\in\Re,\ B\cap C=A} m_1(B)\ m_2(C) \qquad \text{for } A\in\Re. \tag{3.11}$$

It can be shown that $\mathbf{m}_{12} = \mathbf{S}_{m_2}\cdot\mathbf{m}_1 = \mathbf{S}_{m_1}\cdot\mathbf{m}_2$, and $\mathbf{S}_{m_1}\cdot\mathbf{S}_{m_2} = \mathbf{S}_{m_2}\cdot\mathbf{S}_{m_1}$. The relation (3.11) corresponds to the so-called Dempster's rule of combination (except for the normalization factor).

3.5.3. Commonalties are eigenvalues.

The eigenvalues of the dempsterian specialization matrix $\mathbf{S}_m$ are the values of the commonality function q derived from m. The eigenvectors of $\mathbf{S}_m$ do not depend of m. Their matrix is the matrix that transforms a basic belief assignment m into its related commonality function q (Klawonn and Smets, 1992). These properties explain the mathematical importance of the commonality functions.

3.5.4. The concept of distinctness.

The combination rule (3.11) is known as Dempster's rule of combination. It was initially proposed by Shafer (1976) as the rule to be applied to combine the basic belief assignments m_1 and m_2 induced by two distinct pieces of evidence Ev_1 and Ev_2, respectively. Historically, distinctness was not precisely defined[5]. So far, distinctness means in fact that the specialization matrix that transforms m_1 into m_{12} does not depend on m_1. In Smets (1992c), we analyze the meaning of distinctness. It is based on the agent You, two extra agents, denoted You_1 and You_2, and two pieces of evidence Ev_1 and Ev_2. You_i knows only the piece of evidence Ev_i. Let $B_\Re$ be the set of basic belief assignments defined over $\Re$. In order to represent his beliefs on $\Re$, each You_i selects one basic belief assignment in $B_\Re$. You_i is free to represent his beliefs by selecting any basic belief assignment in $B_\Re$. You will collect the basic belief assignments m_1 and m_2, and combine them as You feel that both agents are perfectly reliable.

Suppose You know that if You_1 selects the basic belief assignment m_1, then You_2 can only select m_2 in a strict subset of $B_\Re$. In that case the pieces of evidence Ev_1 and Ev_2 are not distinct as there is something in Ev_1 that induces You_2 to restrict the domain in which he can select the basic belief assignment that represents the beliefs on $\Re$ induced by Ev_2.

Two pieces of evidence will be said to be distinct, if there is no basic belief assignment in $B_\Re$ such that if You learn that it has been selected by You_1 to represent his beliefs on $\Re$, than You know that You_2 can only select his basic belief assignment in a strict sub domain of $B_\Re$ (and vice versa).

---

[5]except for the Dempsterian model (Shafer and Tversky, 1985) but the Dempsterian model is different from the transferable belief model presented here.

When the two pieces of evidence are distinct, You will apply Dempster's rule of combination to combine the belief functions they induce.

How is that property related to the Dempsterian nature of the specialization matrices. Suppose a vacuous evidential corpus ECØ and two pieces of evidence Ev_1 and Ev_2. Let m_1 and m_2 be the basic belief assignments induced by You on $\mathfrak{R}$ when Your evidential corpus contains only Ev_1 or Ev_2, respectively ('only' is relative for what concerns Ω). Let $\mathbf{S}_{12}$ and $\mathbf{S}_{21}$ be the specialization matrices that transform m_1 into m_{12} and m_2 into m_{21}, respectively. One must have $m_{12} = m_{21}$. To say that the $\mathbf{S}_{12}$ does not depend on m_1 means that $\mathbf{S}_{12}$ depends only on Ev_2, and not on Ev_1 (and analogously for $\mathbf{S}_{21}$). When $\mathbf{S}_{12}$ and $\mathbf{S}_{21}$ depend only on Ev_2 and Ev_1, respectively, then we showed that the equality $m_{12} = m_{21}$ implies that $\mathbf{S}_{12}$ and $\mathbf{S}_{21}$ are the dempsterian specialization matrices $\mathbf{S}_{m_2}$ and $\mathbf{S}_{m_1}$, thus completely characterized by m_2 and m_1, respectively. To say that $\mathbf{S}_{m_1}$ is independent of m_2 means that the value of $\mathbf{S}_{m_1}$, hence m_1, is not constrained by the knowledge of the value of m_2. Hence the pieces of evidence Ev_1 and Ev_2 inducing m_1 and m_2 are distinct where distinctness could be understood as 'epistemically independent'. Therefore distinctness implies that the specialization matrices are Dempsterian, hence that Dempster's rule of combination is the appropriate rule to combine them.

The distinctness concept can be understood as the fact that the impact of Ev_1 is fully characterized by a specialization matrix which values depend only on m_1 and not on m_2, (and similarly for the impact of Ev_2). A first generalization to the case where the pieces of evidence are not distinct is tackled in Kennes (1991).

3.6. The canonical decomposition of a belief function

3.6.1. Simple support function.
A particular but highly useful belief function is the so-called simple support function. It is a belief function that has at most two non-null basic belief masses, one being given to Ω, the other being given to a subset of Ω called the focus. Note that the empty set Ø can be a focus. So for $A \subseteq \Omega$, let A^x denote the simple support function with focus A and basic belief masses $m(A)=1-x$, $m(\Omega)=x$. It represents the belief function induced by a piece of evidence that supports A (at a level 1-x) and leaves the remaining belief uncommitted ($m(\Omega)=x$).

3.6.2. Absorbing Beliefs.
A strange state of belief can be described that we call a state of 'absorbing' belief. The simple support function A^x represents a state of belief that translates the idea that "You have some reason to believe that the actual world is in A (and nothing more)". The 1-x is the weight corresponding to "some reasons". Suppose another state of belief that would translate the idea that "You have some reason *not* to believe that the actual world is in A". This cannot be represented by a belief function over Ω and it seems there is no way to represent it by a meta-belief over the set of belief functions over the propositional space

$(\Omega,\mathfrak{R})$. Suppose that You are in a situation where You have simultaneous some reason to believe A and some reason *not* to believe A. It might occur that the weights of both 'some reasons' are exactly counter-balancing each other. In that case, it seems reasonable to assert that You end up in a state of total ignorance, hence Your belief over Ω is represented by a vacuous belief function. The first state of belief is represented by a simple support function A^x. So the second state of belief must be represented by 'something' whose combination with A^x leads to a vacuous belief function. But there is no belief function whose combination with another belief function by Dempster's rule of combination would result in a vacuous belief function. The state of belief encountered when there are some reasons *not* to believe A is called a state of absorbing belief as it is a state of belief that will absorb A^x. It looks like a state of belief where You have a 'debt of belief' as the accumulation of new pieces of evidence could lead You to a classical state of belief. The representation of such a state of absorbing belief cannot be achieved by a single belief function.

3.6.3. Latent beliefs.

A way to solve this strange state consists in creating a structure of latent beliefs and a structure of apparent beliefs. A latent belief structure is represented by a pair of belief functions (C, D) where $C,D\in \mathtt{B}$ and $\mathtt{B}$ is the set of belief functions over the propositional space $(\Omega,\mathfrak{R})$. C and D are respectively quantifying a confidence and a diffidence component of the latent belief structure. So C is called the confidence component and D is called the diffidence component of (C,D). Let T represent the vacuous belief function. (C,T) describes a state of belief where You have only a confidence component. It is the classical state of belief considered so far. (T,D) describes a pure state of absorbing belief where there is only a diffidence component. For example, the state of belief induced by "You have some reasons to believe A" is represented by (A^x,T) and the state "You have some reasons *not* to believe A" is represented by (T,A^y) (for $x,y\in[0,1]$, and where x and y are the complements of the weights corresponding to the 'some reasons').

Two latent belief structures are combined by Dempster's rule of combination (denoted by $\oplus$) applied to both the confidence and diffidence components.

$$(X,Y) \oplus (U,V) = (X\oplus U,Y\oplus V).$$

In particular,

$$(X,T) \oplus (T,X) = (X,X)$$

A latent belief structure can induce an apparent belief structure represented by an element of $\mathtt{B}$. Let Λ be an operator that transforms a latent belief structure into an apparent belief structure: $\Lambda:\mathtt{B}\mathrm{x}\mathtt{B}\rightarrow\mathtt{B}$. We require that if there is only a non vacuous confidence component then the apparent belief structure is equal to the confidence component: $\Lambda(C,T) = C$. We also want that if the confidence and the diffidence components are equal, they counter-balance each other and the resulting apparent belief structure is vacuous: $\Lambda(X,X) = T$ and in particular $\Lambda(A^x,A^x)=T$ as already developed in 3.6.2.

Consider now the equalities:

$$(X\oplus Y,X) = (X,X)\oplus(Y,T) = (T,T)\oplus(Y,T) = (Y,T)$$

Thus that for every $X,Y,Z\in \mathrm{B}$, $(X,Y) = (X\oplus Z,Y\oplus Z)$: $\oplus$–combining the same belief function Z with both the confident and the diffident part leaves the latent belief structure unchanged.

Let us introduce the operator Θ. It is defined as the inverse of the $\oplus$ operator in that sense that: $(X\oplus Y)\Theta Y = X$ for all $X,Y\in \mathrm{B}$

and $Y\Theta Y = T$.

So if q_X and q_Y are the commonality functions related to the belief functions X and Y, then the commonality function related to $X\Theta Y$ is defined as:

$$q_{X\Theta Y}(A) = \frac{q_X(A)}{q_Y(A)} \qquad \text{for all } A\in \Re,$$

whereas $q_{X\oplus Y}(A) = q_X(A)\, q_Y(A)$ for all $A\in \Re$. This last relation is just Dempster's rule of combination. Whatever X and Y, the function $q_{X\oplus Y}$ so defined is always related to a belief function, but this is not always true for $q_{X\Theta Y}$.

It happens that whenever $m_{12}(\Omega)>0$, one can always recover m_2 from the knowledge of m_1 and $m_{12} = m_1\oplus m_2$. We can then write:

If $X\Theta Y\in \mathrm{B}$, then $(X,Y) = (X\,\Theta Y,T)$ and $\Lambda(X,Y) = X\Theta Y$ (3.12)

If $X\Theta Y\notin \mathrm{B}$, then $\Lambda(X,Y)$ is undefined.

So Λ is not defined on the whole space BxB, but only on those elements (X,Y) where $X\Theta Y$ is a belief function in B. We could have hoped that such a state of belief would not occur. Unfortunately we already encountered a counter example when we introduced the latent belief structure (T,A^x) that characterizes the case where all You know is that You have good reasons not to believe A. This means that the apparent belief structures are not rich enough to characterize every belief state. Some state of belief can only be represented by their latent belief structure.

What should be an appropriate apparent belief structure when $X\Theta Y\notin \mathrm{B}$ is not clear. What is the apparent belief structure in the case (T,A^x)? We could claim that $\Lambda(T,A^x) = T$, but then the apparent vacuous belief structure T could correspond to many non equivalent latent belief structures. How to solve the general case? We could propose that $\Lambda(X,Y)$ is the belief function 'closest' from $X\Theta Y$. Unfortunately such a concept of 'closeness' is not yet available. The specialization concept can be used to create a partial order on the set of belief functions. Pointwise measures of the information contained in a belief function have been proposed (Pal and Bezdek, 1992), but none seems really convincing as THE appropriate measure to define 'closeness', and furthermore they should be adapted to functions like $X\Theta Y$ when they are not belief functions.

3.6.4. The anatomy of a belief function.

We show now how to build the latent belief structure from a given belief function. Suppose two set functions X and Y that map $\Re$ onto the real. Let q_X and q_Y be their related 'commonality' functions (the relations between q_X and X and q_X and Y are the same as those between a commonality function and a belief function). Then we define the ⊕-combination of X and Y such that the resulting related 'commonality' function satisfies:

$$q_{X\oplus Y}(A) = q_X(A)\, q_Y(A) \qquad \text{for all } A\in\Re.$$

It can be shown that for every belief function bel in $\mathbb{B}$ with $m(\Omega)>0$, there is an unique mathematical decomposition of bel into a set of generalized simple support functions:

$$bel = \oplus_{A\subseteq\Omega} A^{x_A}$$

where $x_A \geq 0$ [6] and where Dempster's rule of combination is generalized to any set functions. We say 'generalized' simple support function as x_A can be larger than one. Let

$$\mathbb{A}+ = \{A: x_A<1\} \text{ and } \mathbb{A}- = \{A: x_A>1\}$$

then $bel = \oplus_{A\in\mathbb{A}+} A^{x_A} \;\Theta\; \oplus_{A\in\mathbb{A}-} A^{1/x_A}$

This decomposition corresponds in fact to (3.12). So the latent belief structure underlying bel is given by $(\oplus_{A\in\mathbb{A}+} A^{x_A}, \oplus_{A\in\mathbb{A}-} A^{1/x_A})$.

The case $m(\Omega) = 0$ can be solved by considering that $m(\Omega) = 0$ is only the limit of $m(\Omega) = \varepsilon$ for $\varepsilon\to 0$. Then every computation is done with ε and limits are taken as the last operation.

The meaning of the decomposition is illustrated in the next example.
Let $\Omega = \{a,b,c]$ and the basic belief masses associated to bel be $m(\{a,b\}) = m(\{a,c\}) = m(\{a,b,c\}) = 1/3$. Then: $\quad bel = \{a,b\}^{1/2}\oplus\{a,c\}^{1/2}\oplus\{a\}^{4/3}$,
in which case the latent belief structure is $(\{a,b\}^{1/2}\oplus\{a,c\}^{1/2},\{a\}^{3/4})$. Such a latent belief structure describes the situation where:

You have some reasons to believe {a,b} (weight 1/2),
You have some reasons to believe {a,c} (weight 1/2),
You have some reasons *not* to believe {a} (weight 3/4).

Whenever Your belief is represented by a belief function, the latent belief structure can always be determined. The apparent belief structure can always be recovered from that latent belief structure as it should. Formally:

if $bel = \oplus_{A\in\mathbb{A}+} A^{x_A} \;\Theta\; \oplus_{A\in\mathbb{A}-} A^{1/x_A}$ then $bel = \Lambda(\oplus_{A\in\mathbb{A}+} A^{x_A}, \oplus_{A\in\mathbb{A}-} A^{1/x_A})$.

The only belief states that cannot be represented by an apparent belief structure are those states where the diffidence components 'dominate' the confidence components (when Λ

---

[6] Shafer (1976, page 94) describes the algorithm to compute the weights x_A. Their logartithms are linear functions of the logarithms of the commonality function, and the Fast Möbius Transform can be used for their computation (Kennes, 1992).

in undefined). Such states of belief require the use of the latent structure in order to obtain a full representation.

Example 4: The Newspaper Bias.

You visit a foreign country and You read in the local Journal that the economic situation in the region X is good. You never heard about that Journal nor of the existence of region X. So You had no a priori whatsoever about the economic status of region X, and now after having read the Journal, You might have some reasons to believe that the economic status is good. The 'some reasons' reflects the strength of the trust You put in the Journal's information. Then a friend in which You have some confidence mention to You that the Journal is completely under control by the local propaganda organization, therefore You have some reasons *not* to believe the Journal when it describes the good economic status of the country, it might just be propaganda.

The reasons to believe that the economic status is good that results from the information presented in the Journal, and the reasons *not* to believe it that results from what Your friend said could counterbalance each other, in which case You end up in a state of total ignorance about the economic status in region X. The diffidence component that results from what Your friend said about the Journal is balanced by the confidence component that results form the information in the Journal .

It might be that the confidence component is stronger than the diffidence component, then You will end up with a slight belief that the economic status is good (but the belief is not as strong as if You had not heard what Your friend said). If the diffidence component is still stronger than the confidence component, then You are still in a state of 'debt of belief', in the sense that You will need a further confidence component (some extra information that supports that the economic status is good) in order to balance the remaining diffidence component. In such a case, if You are asked to express Your opinion about the economic status, You might express it under the form: 'So far, I have no reason to believe that the economic status is good, and I need some extra reasons before I start to believe it'.

These arguments might look like a discounting of what the Journal claims. But it is not the case as You will believe what the Journal says when the claims are not favorable to the local regime (a statement like 'the economic status is not good' would be such an example). Discounting would mean that You reject whatever the Journal claims. What we face here is a kind of focused and weighted discounting, and we solve such a situation by introducing the latent belief structure.

4. Decision Making and Dutch Books.

4.1. The pignistic probability function for decision making.

Suppose a credibility space (Ω, $\Re$, bel) where bel quantifies Your beliefs at the credal level. When a decision must be made that depends on ω_0, You construct a probability function on $\Re$ in order to make the optimal decision, i.e., the one that maximizes the expected utility (Savage, 1954, DeGroot, 1970). We assume that the probability function defined on $\Re$ is a function of the belief function bel. It translates the saying that beliefs guide our actions. Hence one must transform bel into a probability function that will be used for selecting the best decision. Let BetP denoted this probability function. The transformation is called the pignistic transformation and is denoted by $\Gamma_\Re$. The $\Re$ index in $\Gamma_\Re$ mentions the Boolean algebra $\Re$ on which bel and BetP are defined: so BetP = $\Gamma_\Re$(bel) where bel and BetP map $\Re$ into [0, 1]. We call BetP a pignistic probability to insist on the fact that it is a probability measure used to make decisions (Bet is for betting). Of course BetP is a classical probability measure.

The structure of the pignistic transformation is derived from the rationality requirement that underlies the following scenario.

Example 5: Buying Your friend's drink. Suppose You have two friends, G and J. You know they will toss a fair coin and the winner will visit You tonight. You want to buy the drink Your friend would like to have tonight: coke, wine or beer. You can only buy one drink. Let D = {coke, wine, beer} and $\Re = 2^D$.

Let $\mathrm{bel}_G(d)$, for all $d \subseteq D$, quantify Your belief about the drink G is liable to ask for. Given bel_G, You build the pignistic probability BetP_G about the drink G will ask by applying the (still to be defined) pignistic transformation. You build in identically the same way the pignistic probability BetP_J based on bel_J, Your belief about the drink J is liable to ask for. The two pignistic probability distributions BetP_G and BetP_J are the conditional probability distributions about the drink that will be asked for given G or J comes. The pignistic probability distributions BetP_{GJ} about the drink that Your visitor will ask for is then:

$$\mathrm{BetP}_{GJ}(d) = .5\ \mathrm{BetP}_G(d) + .5\ \mathrm{BetP}_J(d) \qquad \text{for all } d \in D$$

You will use these pignistic probabilities $\mathrm{BetP}_{GJ}(d)$ to decide which drink to buy.

But You might as well reconsider the whole problem and first compute Your belief bel_V about the drink Your visitor (V) would like to have. We have shown (Smets, 1995) that bel_V is given by:

$$\mathrm{bel}_V(d) = .5\ \mathrm{bel}_G(d) + .5\ \mathrm{bel}_J(d) \qquad \text{for all } d \subseteq D$$

Given bel_V, You could then build the pignistic probability BetP_V You should use to decide which drink to buy. It seems reasonable to assume that BetP_V and BetP_{GJ} must be equal. In such a case, the pignistic transformation is uniquely defined. t

Formally, we have assumed:

Linearity Assumption: Let bel_1 and bel_2 be two belief functions on the propositional space $(\Omega, \Re)$. Let $\Gamma_\Re$ be the pignistic transformation that transforms a belief function over $\Re$ into a probability function BetP over $\Re$. Then $\Gamma_\Re$ satisfies, for any $\alpha \in [0,1]$,

$$\Gamma_\Re(\alpha bel_1 + (1-\alpha) bel_2) = \alpha\, \Gamma_\Re(bel_1) + (1-\alpha)\, \Gamma_\Re(bel_2)$$

Two technical assumptions must be added that are hardly arguable:

Anonymity Assumption: An anonymity property for BetP by which the pignistic probability given to the image of $A \in \Re$ after permutation of the atoms of $\Re$ is the same as the pignistic probability given to A before applying the permutation.

Impossible Event Assumption The pignistic probability of an impossible event is zero.

Under these assumptions, it is possible to derive uniquely $\Gamma_\Re$ (Smets, 1990b)

Pignistic Transformation Theorem: Let $(\Omega, \Re, bel)$ be a credibility space, with m the bba related to bel. Let $BetP = \Gamma_\Re(bel)$. The only solution that satisfies the Linearity, Anonymity and Impossible Event Assumptions is:

$$\mathrm{BetP}(\omega) = \sum_{A:\omega \subseteq A \in \Re} \frac{m(A)}{|A|\ (1-m(\emptyset))} \quad \text{for any atom } \omega \text{ of } \Re \tag{4.1}$$

where $|A|$ is the number of atoms of $\Re$ in A,

and $$\mathrm{BetP}(A) = \sum_{\omega:\ \omega \in At(\Re),\ \omega \subseteq A \in \Re} \mathrm{BetP}(\omega) \quad \text{for all } A \in \Re.$$

Historical note. In a context similar to ours, Shapley (1953) derived the same relation (4.1). The model he derived was later called the 'transferable utility model' whereas, unaware of it, we called our model the 'transferable belief model'.

It is easy to show that the function BetP obtained from (4.1) is a probability function and the pignistic transformation of a probability function is the probability function itself.

Betting frame. The pignistic transformation depends on the structure of the frame on which the decision must be made. One must first define the 'betting frame' $\Re$ on Ω, i.e., the set of atoms on which stakes will be allocated. The granularity of this frame $\Re$ is defined so that a stake could be given to each atom of $\Re$ independently of the stakes given to the other atoms of $\Re$. Suppose one starts with a belief function on a frame $\Re_0$. If the stakes given to atoms A and B of $\Re_0$ must necessarily be always equal, both A and B belong to the same granule of the betting frame $\Re$. The betting frame $\Re$ is organized so that the granules are the atoms of $\Re$. $\Re$ results from the application of a sequence of coarsenings and/or refinements on $\Re_0$. The pignistic probability BetP is then built from the belief function so derived on $\Re$. Thus BetP is a function from $\Re$ to [0,1].

Betting under total ignorance. To show the potency of our approach, let us consider one of those disturbing examples based on total ignorance.

Example 6. Betting and total ignorance.
Consider a guard in a huge power plant. On the emergency panel, alarms A_1 and A_2 are both on. The guard never heard about these two alarms, they were hidden in a remote place. He takes the instruction book and discovers that alarm A_1 is on iff circuit C is in state C_1 or C_2 and that alarm A_2 is on iff circuit D is in state D_1, D_2 or D_3. He never heard about these C and D circuits. Therefore, his beliefs on the C circuit will be characterized by a vacuous belief function on space $\Omega_C = \{C_1, C_2\}$. By the application of (4.1) his pignistic probability will be given by $BetP_C(C_1) = BetP_C(C_2) = 1/2$. Similarly for the D circuit, the guard's belief on space $\Omega_D = \{D_1, D_2, D_3\}$ will be vacuous and the pignistic probability are $BetP_D(D_1) = BetP_D(D2) = BetP_D(D3) = 1/3$. Now, by reading the next page on the manual, the guard discovers that circuits C and D are so made that whenever circuit C is in state C_1, circuit D is in state D_1 and vice-versa. So he learns that C_1 and D_1 are equivalent (given what the guard knows) and that C_2 and (D_2 or D_3) are also equivalent as C is either C_1 or C_2 and D is either D_1 or D_2 or D_3. This information does neither modify his belief nor his pignistic probability about which circuit is broken.

If the guard had been a trained Bayesian, he would have assigned value for $P_C(C_1)$ and $P_D(D_1)$ (given the lack of any information, they would probably be 1/2 and 1/3, but any value could be used). Once he learns about the equivalence between C_1 and D_1, he must adapt his probabilities as they must give the same probabilities to C_1 and D_1. Which set of probabilities is he going to update: P_C or P_D, and why?, especially since it must be remembered that he has no knowledge whatsoever about what the circuits are. In a probabilistic approach, the difficulty raised by this type of example results from the requirement that equivalent propositions should receive identical beliefs, and therefore identical probabilities.

Within the transferable belief model, the only requirement is that equivalent propositions should receive equal beliefs (it is satisfied as $bel_C(C_1) = bel_D(D_1) = 0$). Pignistic probabilities depend not only on these beliefs but also on the structure of the betting frame. The difference between $BetP_C(C_1)$ and $BetP_D(D_1)$ reflects the difference between the two betting frames. t

The fact the TBM can cope easily with such a state of ignorance results from the dissociation between the credal and the pignistic levels. Bayesians do not consider such a distinction and therefore work in a much limited framework, hence the difficulty they encounter in the present situation.

We consider now the problem where the betting frame is ill defined. Suppose bel is a belief function on a frame $\mathfrak{R}$, and let $BetP_{\mathfrak{R}}$ be the pignistic probability obtained by applying the pignistic transformation $\Gamma_{\Re}$ to bel. So for $A\in\Re$, $BetP_{\Re}(A) = \Gamma_{\Re}(bel)(A)$. The index $\Re$ of $BetP_{\Re}$ indicates the betting frame. Suppose one changes the betting frame $\Re$ into a new betting frame $\Re^*_A$ by a sequence of uninformative refinement/coarsening applied to Ω, and such that A is still an element of $\Re^*_A$. Let bel* be the belief function induced from bel on $\Re^*_A$ by the same sequence of uninformative refinement/coarsening. In that new betting frame $\Re^*_A$, one can compute $BetP_{\Re^*}(A) = \Gamma_{\Re^*}(bel^*)(A)$. Suppose now the set of possible betting frames $\Re^*_A$ that can be built from $\Re$ and such that $A\in\Re^*_A$, and the set of belief functions bel* induced from bel on $\Re^*_A$. Consider the set of values $BetP_{\Re^*_A}(A)$ obtained by considering all these betting frames $\Re^*_A$. Wilson (1993) shows that, for all $A\in\Re$, the minimum of $BetP_{\Re^*_A}(A)$ taken over the $\Re^*_A$'s is equal to bel(A). So the set of pignistic probabilities $BetP_{\Re^*}$ that can be obtained from bel by varying the betting frame $\Re$ is directly related to the set P(bel) of probability functions 'compatible' with bel and its associated plausibility function pl, i.e., P(bel) is the set of probability functions P on $\Re$ such that $bel(A) \le P(A) \le pl(A)$ $\forall A\in\Re$. So whatever the betting frame $\Re^*_A$, $BetP_{\Re^*_A}(A) \ge bel(A)$ $\forall A\in\Re$.

Suppose You ignore what is the appropriate betting frame, You nevertheless know that, $\forall A\in\Re$, the lowest bound of BetP(A) is bel(A). Therefore bel(A) can then be understood as the lowest pignistic probability one could give to A when the betting frame is not fixed (Giles, 1982).

This set P(bel) of probability functions compatible with a belief function bel gets a meaning from this result. It is the set of pignistic probability functions define on $\Re$ that could be induced by bel when varying the betting frame. Its definition follows from bel, not the reverse as assumed by the authors who understand bel as the lower envelop of some class of probability functions. In the TBM, we get P(bel) from bel, not bel from P(bel).

4.2. The impact of the two-level model.

In order to show that the introduction of the two-level mental model is not innocuous, we present an example where the results will be different if one takes the two-level approach as advocated in the transferable belief model or a one-level model like in probability theory.

Example 7: The Peter, Paul and Mary Saga.
Big Boss has decided that Mr. Jones must be murdered by one of the three people present in his waiting room and whose names are Peter, Paul and Mary. Big Boss has decided that the killer on duty will be selected by a throw of a dice: if it is an even number, the killer will be female, if it is an odd number, the killer will be male. You, the judge, know that Mr. Jones has been murdered and who was in the waiting room. You know about the dice throwing, but You do not know what the outcome was and who was actually selected. You are also ignorant as to how Big Boss would have decided between Peter and Paul in the case of an odd number being observed. Given the available information at time t_0, Your odds for betting on the sex of the killer would be 1 to 1 for male versus female.

At time $t_1 > t_0$, You learn that if Big Boss had not selected Peter, then Peter would necessarily have gone to the police station at the time of the killing in order to have a perfect alibi. Peter indeed went to the police station, so he is not the killer. The question is how You would bet now on male versus female: should Your odds be 1 to 1 (as in the transferable belief model) or 1 to 2 (as in the most natural Bayesian model).

Note that the alibi evidence makes 'Peter is not the killer' and 'Peter has a perfect alibi' equivalent. The more classical evidence 'Peter has a perfect alibi' would only imply P('Peter is not the killer' | 'Peter has a perfect alibi') = 1. But P('Peter has a perfect alibi' | 'Peter is not the killer') would be undefined and would then give rise to further discussion, which would be useless for our purpose. In this presentation, the latter probability is also 1.

The transferable belief model solution.
Let k be the killer. The information about the waiting room and the dice throwing pattern induces the following basic belief assignment m_0:

$$k \in \Omega = \{\text{Peter, Paul, Mary}\}$$

$$m_0(\ \{\text{Mary}\}\ ) = .5 \qquad m_0(\ \{\text{Peter, Paul}\}\ ) = .5$$

The bbm .5 given to {Peter, Paul} corresponds to that part of belief that supports "Peter or Paul", could possibly support each of them, but given the lack of further information, cannot be divided more specifically between Peter and Paul.

Let $BetP_0$ be the pignistic probability obtained by applying the pignistic transformation to m_0 on the betting frame which set of atoms is {{Peter}, {Paul}, {Mary}}. By relation (4.1), we get:

$$BetP_0(\{\text{Peter}\}) = .25 \qquad BetP_0(\{\text{Paul}\}) = .25 \qquad BetP_0(\{\text{Mary}\}) = .50$$

Given the information available at time t_0, the bet on the killer's sex (male versus female) is held at odds 1 to 1.

Peter's alibi induces an updating of m_0 into m_2 be Dempster's rule of conditioning:

$m_2(\{Mary\}) = m_2(\{Paul\}) = .5$

The basic belief mass that was given to "Peter or Paul" is transferred to Paul.

Let $BetP_2$ be the pignistic probability obtained by applying the pignistic transformation to m_2 on the betting frame whose set of atoms is $\{\{Paul\}, \{Mary\}\}$.

$BetP_2(\{Paul\}) = .50 \qquad BetP_2(\{Mary\}) = .50$

Your odds for betting on male versus female would still be 1 to 1.

The probabilistic solution:

The probabilistic solution is not obvious as one data is missing: the value α for the probability that Big Boss selects Peter if he must select a male killer. Any value could be accepted for α, but given the total ignorance in which we are about this value, let us assume that α =.5, the most natural solution (any value could be used without changing the problem we raise). Then the odds on male versus female before learning about Peter's alibi is 1 to 1, and after learning about Peter's alibi, it becomes 1 to 2. The probabilities are then:

$P_2(\{Paul\}) = 0.33 \qquad P_2(\{Mary\}) = 0.66.$

The 1 to 1 odds of the transferable belief model solution can only be obtained in a probabilistic approach if $\alpha = 0$. Some critics would claim that the transferable belief model solution is valid as it fits with $\alpha = 0$. The only trouble with this answer is that if the alibi story had applied to Paul, than we would still bet at 1 to 1 odds within the TBM approach. Instead the probabilistic solution with $\alpha = 0$ would lead to a 0 to 1 bet, as the probabilities are:

$P_2(\{Peter\}) = 0.0 \qquad P_2(\{Mary\}) = 1.$

So the classical probabilistic analysis does not lead to the transferable belief model solution.

We are facing two solutions for the bet on male versus female after learning about Peter's alibi: the 1 to 1 or at 1 to 2 odds? Which solution is 'good' is not decidable, as it would require the definition of 'good'. Computer simulations have been suggested for solving the dilemma, but they are impossible. Indeed when the killer is a male, we do not know how to choose between Peter and Paul. If we introduce a probability α equal to the probability that Peter is selected when the killer is a male, then the problem is no more the one we had consider in the initial story. If such an α were known, then it would been included in the TBM analysis, and in that case it happens that the TBM and the Bayesian solutions become identical, as it should. So in order to compare the TBM and the Bayesian solution of the initial saga, we are only left over with a subjective comparison of the two solutions... or an in depth comparison of the theoretical foundations that led to these solutions.

4.3. The assessment of the values of bel.

The pignistic transformation can be used in order to assess degrees of belief thanks to the ability to construct several betting frames. The method is essentially identical to the one described to assess subjective probabilities. The numerical value of the credibility function is obtained through some exchangeable bets schema.

Example 8: Suppose $\Omega = \{a, b\}$ where {a} = 'Circuit X is broken' and {b} = 'Circuit X is not broken'. Consider the betting frame $\Re$ with atoms {a} and {b}. Suppose Your pignistic probabilities on that frame $\Re$ are:

$\mathrm{BetP}(\{a\}) = 4/9 \qquad \mathrm{BetP}(\{b\}) = 5/9.$

Suppose ψ and $\overline{\psi}$ are two complementary but otherwise unknown propositions that state that circuit C whose properties are completely unknown to You is broken or not broken, respectively. $\{a\}\cap\psi$ will occur if circuits X and C are broken. $\{a\}\cap\overline{\psi}$ will occur if circuit X is broken and circuit C is not broken. Let us consider the betting frame $\Re'$ with atoms $\{a\}\cap\psi$, $\{a\}\cap\overline{\psi}$, {b}, and suppose Your pignistic probabilities on that new frame are:

$\mathrm{BetP}'(\{a\}\cap\psi) = \mathrm{BetP}'(\{a\}\cap\overline{\psi}) = 7/27 \qquad \mathrm{BetP}'(\{b\}) = 13/27.$

Then the unique solution for m is: m({a}) = 2/9, m({b}) = 3/9 and m({a,b}) = 4/9.

Let m* be the bba induces by m on $\Re$' by the uninformative refinement:

$m^*(\{\{a\}\cap\psi) = m^*(\{\{a\}\cap\overline{\psi}) = 0,$

$m^*((\{a\}\cap\psi)\cup(\{a\}\cap\overline{\psi})) = m(\{a\}),$

$m^*((\{a\}\cap\psi)\cup\{b\}) = m^*((\{a\}\cap\overline{\psi})\cup\{b\}) = 0$

$m^*(\{b\}) = m(\{b\})$

$m^*((\{a\}\cap\psi)\cup(\{a\}\cap\overline{\psi})\cup\{b\}) = m(\{a, b\}).$

The solution for m must solve two linear equations derived from (4.1):

$$4/9 = m(\{a\}) + m(\{a,b\})/2$$

$$\begin{aligned} 7/27 &= m^*(\{\{a\}\cap\psi) + m^*((\{a\}\cap\psi)\cup(\{a\}\cap\overline{\psi}))/2 + m^*((\{a\}\cap\psi)\cup\{b\})/2 \\ &\quad + m^*((\{a\}\cap\psi)\cup(\{a\}\cap\overline{\psi})\cup\{b\})/3 \\ &= m(\{a\})/2 + m(\{a,b\})/3. \end{aligned}$$

Hence the values of m.

It might seem odd that {b} receives pignistic probabilities of 5/9 and 13/27 according to the betting context. It reflects the fact that a large amount (4/9) of Your initial belief was left unassigned (i.e., given to {a,b}). This example corresponds to a state in which You have very weak support for {a} and for {b}. You are not totally ignorant as in example 6, but still in a state of 'strong' ignorance. Part of BetP({b}) = 5/9 is due to justified beliefs (3/9) but the remainder results from a completely unassigned part of belief that You distribute equally through the pignistic transformation among the alternatives of Your betting frame. t

4.4. Dutch Books.

A classical criticism against any non probabilistic model for quantified beliefs is based on Dutch Books arguments, i.e. on the possibility to build a set of forced bets such that the player will lose for sure. Dutch Books are avoided only if beliefs, when used for decision making, are quantified by probability measures. The transferable belief model with its pignistic transformation resists to such a criticism. Static (synchronic) Dutch Books are of course avoided inasmuch as bets are based on pignistic probabilities. The real delicate point is to resist diachronic Dutch Books, i.e., those built when changes in beliefs are considered and bets can be reconsidered after new information has been collected by the players (Teller, 1973, Jeffrey, 1988).

They show that the impact of the new information should be represented by the classical conditioning rule described in probability theory. The argument is based on the acceptance of the temporal coherence principle that we do not assume in the TBM. The temporal coherence principle requires that Your belief that the actual world ω_0 belongs to $A\in\Re$ once You know that ω_0 belongs to $B\in\Re$ (a factual conditioning) should be the same as Your belief that $\omega_0\in A$ under the hypothesis that You come to know that $\omega_0\in B$ (an hypothetical conditioning).

In Smets (1993a), we show why the temporal coherence principle is not necessary, hence the transferable belief model resists diachronic Dutch Books criticism, and how the appropriate pignistic probabilities can be built up when the player knows that intermediate experiments will be run whose outcomes could affect the bets involved. The originality of the models we obtain in that way is that the player will use different probabilities depending on whether he knows about the experiments to be run or not. The transferable belief model analysis is rich enough to quantify the impact of the knowledge that some relevant intermediate experiments will be run, which is not the case within probability theory. Full details about the construction of the pignistic probabilities in a context of diachronic Dutch Book are presented in Smets (1993a).

5. The Generalized Bayesian Theorem.

5.1. The Theorem.

Bayesian Theorem is a key element of the use of probability theory for diagnosis process. Suppose two spaces, the X space of symptoms and the Θ space of diseases. Given the conditional probability $P_X(x|\theta_i)$ of observing $x\subseteq X$ in each disease class $\theta_i\in\Theta$, and the a priori probability P_Θ over Θ, compute the a posteriori probability $P_\Theta(\theta|x)$ over Θ that the patient belongs to a disease class in θ given the symptom $x\subseteq X$ has been observed. Indices indicate the domain of the functions. By Bayes Theorem, one has:

$$P_\Theta(\theta_i|x) = \frac{P_X(x|\theta_i)\ P_\Theta(\theta_i)}{\sum_j P_X(x|\theta_j)\ P_\Theta(\theta_j)} \qquad \text{for } \theta_i\in\Theta.$$

We have generalized the theorem in the context of the transferable belief model (Smets, 1978, 1981, 1993b). We assume that for each disease class $\theta_i \in \Theta$, there is a belief function $bel_X(.|\theta_i)$ over X that represents Your belief about which symptom can be observed if the patient belongs to the disease class θ_i. Let bel_Θ be Your a priori belief about the disease class to which the patient belongs. Suppose bel_Θ is a vacuous belief function that reflects that You have no a priori about the disease Your patient could present. Then we have shown among others that the a posteriori plausibility function $pl_\Theta(.|x)$ over Θ given $x \subseteq X$ is:

$$pl_\Theta(\theta|x) = 1 - \prod_{\theta_i \in \theta} (1 - pl_X(x|\theta_i)). \qquad (5.1)$$

When bel_Θ is not vacuous because You have some a priori about the disease Your patient could present, bel_Θ and $bel_\Theta(.|x)$ are combined be Dempster's rule of combination.

This theorem can be derived by assuming the Least Commitment Principle and either the distinctness of the pieces of evidence that induce in You the conditional belief functions over X given θ_i or a Generalized Likelihood Principle that states that:

$$\forall \theta \subseteq \Theta, \forall x \subseteq X, pl_X(x|\theta) \text{ depends only on } \{pl_X(x|\theta_i) , pl_X(\bar{x}|\theta_i) : \theta_i \in \theta\}, \qquad (5.2)$$

where $pl_X(x|\theta)$ is the conditional plausibility function on X when all we know about Θ is that θ holds. (5.2) states essentially that this last plausibility should only depends on the plausibility $pl_X(.|\theta_i)$ for $\theta_i \in \theta$.

5.2. The Disjunctive Rule of Combination.

Simultaneously we derive a disjunctive rule of combination. Dempster's rule of combination describes in section 3.5.2 concerns the case where two pieces of evidence are combined conjunctively. Suppose bel_1 and bel_2 are the belief functions on $\Re$ induced by two distinct pieces of evidence Ev_1 and Ev_2. The conjunctive rule of combination allows the computation of $bel(.|Ev_1 \wedge Ev_2)$ from bel_1 and bel_2 (by Dempster's rule of combination). The disjunctive rule of combination allows the computation of $bel(.|Ev_1 \vee Ev_2)$ from bel_1 and bel_2. It provides the belief function when You only know that either Ev_1 or Ev_2 hold (whereas in the conjunctive case You know that both hold). One has:

$$bel(\omega|Ev_1 \vee Ev_2) + m(\varnothing|Ev_1 \vee Ev_2) = (bel_1(\omega) + m_2(\varnothing)) (bel_2(\omega) + m_2(\varnothing)) \qquad \text{for } \omega \in \Re.$$

and

$$pl(\omega|Ev_1 \vee Ev_2) = 1 - (1 - pl_1(\omega))(1 - pl_2(\omega)). \qquad (5.3)$$

These disjunctive combination rules are rarely needed. Nevertheless they are quite useful to reduce computation time and computer memory requirements in belief networks (Xu and Smets, 1994).

5.3. Properties of the Generalized Bayesian Theorem.

An important property satisfied by the Generalized Bayesian Theorem concerns the case where two independent observations are collected. Suppose two symptoms spaces, X and Y. Let bel_X and bel_Y be Your beliefs on X and on Y. We assume that the symptoms are independent within each disease class $\theta_i \in \Theta$. The independence assumption means that if You knew which disease holds the observation of one of the symptoms would not change Your belief about the status of the other symptom. This independence property means that the conditional joint belief $bel_{XxY}(.|\theta_i)$ over the space XxY given θ_i is:

$$pl_{XxY}(x \cap y|\theta_i) = pl_X(x|\theta_i)\, pl_Y(y|\theta_i).$$

Suppose Your a priori belief over Θ is vacuous. Given You observe the symptoms $x \subseteq X$ and $y \subseteq Y$, You can build $bel_\Theta(.|x)$ and $bel_\Theta(.|y)$ by the Generalized Bayesian Theorem, and then combine these two belief functions by Dempster's rule of combination in order to derive Your belief $bel_\Theta(.|x,y)$ about Θ given both symptoms x and y. But You could as well apply the Generalized Bayesian Theorem directly to $bel_{XxY}(.|\theta_i)$ in order to derive $bel_\Theta(.|x,y)$. Both approaches lead to the same result, as it should. Furthermore, we proved that the Generalized Bayesian Theorem is essentially the only solution that satisfies that property (Smets, 1993b).

A nice property of the Generalized Bayesian Theorem is that it allows to extend the disease domain Θ with an extra class, the set of still unknown diseases. In that class, Your belief about the symptoms is of course vacuous. How could You have any belief about which symptom prevails for patients that belong in a disease class You never heard about? You can then compute Your a posteriori belief that the patient belongs to that new class given the observed symptom. It means You can compute Your belief that You have made a ‘discovery’. Such a computation is not possible within the probabilistic framework as we cannot represent in probability theory the state of total ignorance we need to describe the beliefs over the symptoms in the new class. The next example illustrates among other this property.

Example 9. Diagnostic process.

In order to illustrate the use of the Generalized Bayesian Theorem and the disjunctive rule of combination, we consider an example of a medical diagnosis process. Let $\Theta = \{\theta_1, \theta_2, \theta_\omega\}$ be a set of disease with three mutually exclusive and exhaustive diseases. θ_1 and θ_2 are two ‘well known’ diseases, i.e., we have some beliefs on what symptoms could hold when θ_1 holds or when θ_2 holds. θ_ω corresponds to the complement of $\{\theta_1, \theta_2\}$ relative to all possible diseases. θ_ω represents not only all the ‘other’ diseases but also those not yet known. In such a context, our belief on the symptoms can only be vacuous. What do we know about the symptoms caused by a still unknown disease? Nothing of course, hence the vacuous belief function.

We consider the set X of symptoms with X = $\{x_1, x_2, x_3\}$. Table 4 presents the beliefs over X within each diseases class. It also shows the beliefs over symptom X when You only know that either θ_1 or θ_2 holds, computed by (5.3). The beliefs translate essentially the facts that θ_1 'causes' (supports) x_3, and θ2 'causes' x_1 or x_2 (without preference).

	$\{\theta_1\}$		$\{\theta_2\}$		$\{\theta_\omega\}$		$\{\theta_1, \theta_2\}$	
X	m	bel	m	bel	m	bel	m	bel
$\{x_1\}$	.0	.0	.0	.0	.0	.0	.00	.00
$\{x_2\}$	.0	.0	.0	.0	.0	.0	.00	.00
$\{x_3\}$	.5	.5	.2	.2	.0	.0	.10	.10
$\{x_1, x_2\}$	.2	.2	.6	.6	.0	.0	.12	.12
$\{x_1, x_3\}$	.0	.5	.1	.3	.0	.0	.05	.15
$\{x_2, x_3\}$	.0	.5	.1	.3	.0	.0	.05	.15
$\{x_1, x_2, x_3\}$	.3	1.0	.0	1.0	1.0	1.0	.68	1.00

Table 4: Conditional basic belief masses (m) and beliefs (bel) on the symptoms $x \subseteq X$ within each of the mutually exclusive and exhaustive disease classes θ_1, θ_2 and $\theta_\omega \in \Theta$. The right hand side of the table presents the beliefs (and basic belief masses) on X given the disease is either θ_1 or θ_2.

Table 5 presents the beliefs induced on Θ by the observation of symptom x_3. The beliefs are computed by the application of (5.1). The symptom supports essentially $\{\theta_1, \theta_\omega\}$.

The meaning of $bel(\theta_\omega|x_3) = 0.12$ merits some consideration. It quantifies Your belief that the symptom x_3 is neither 'caused' by θ_1 nor by θ_2. It supports the fact that the observation is 'caused' by another disease or by some still unknown disease. A large value for $bel(\theta_\omega|x_3)$ somehow supports the fact that You might be facing a new disease. In any case a large value should induce You in looking for other potential causes to explain the observation. t

Θ	$\vert x_3$ m	bel
$\{\theta_1\}$	.00	.00
$\{\theta_2\}$	.00	.00
$\{\theta_\omega\}$	.12	.12
$\{\theta_1, \theta_2\}$	.00	.00
$\{\theta_1, \theta_\omega\}$	.48	.50
$\{\theta_2, \theta_\omega\}$	.08	.20
$\{\theta_1, \theta_2, \theta_\omega\}$	.32	.80

Table 5: The basic belief masses (m) and belief function bel induced on Θ by the observation of symptom x_3.

6. The justification for the use of belief functions.

We present several sets of requirements that justify the use of belief functions for representing quantified beliefs.

Shafer (1976) assumes that any measure of belief Cr on an algebra $\Re$ should satisfy the following inequalities:

$\forall n\geq 1,\ A_1, A_2, \ldots A_n \in \Re,$

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

These inequalities are hardly convincing as the inequalities for $n > 2$ do not have any obvious natural interpretation.

In the TBM (Smets and Kennes, 1994), we start from the concept of parts of beliefs that support a proposition without supporting more specific propositions (Smets and Kennes, 1994). These parts of belief are in fact the values of the bbm related to the belief function.

Both Shafer and the TBM approaches are strictly equivalent. We introduced the second in response to the criticism that the inequalities of Shafer were too artificial and difficult to accept as natural requirements for a measure of belief, hoping ours would be more 'palatable'.

Wong et al. (1990) have presented an axiomatic justification based on the representation of a belief-order relation $\geq$ ($>$) where $B\geq C$ ($B>C$) means 'B is not less believed than C' ('B is more believed than C'). They replace the disjoint union requirement assumed in order to derive probability theory (Koopman 1940, Fine, 1973):

$$A \cap (B \cup C) = \emptyset \Rightarrow (B \geq C \Leftrightarrow A \cup B \geq A \cup C)$$

by a less restrictive requirement:

$C \subseteq B,\ A \cap B = \emptyset \Rightarrow (B > C \Rightarrow A \cup B \geq A \cup C).$

Under this last requirement, the ≥ belief-ordering can always be represented by a belief function. Unfortunately, other functions like the convex capacities can also represent the ≥ ordering.

In Smets (1993c, 1995), we develop a full axiomatization based on rationality requirements. We assume that quantified beliefs must be represented by a pointwise measure, bounded and monotone for inclusion, a function that we call a credibility function. We show that the set of credibility functions defined on the algebra $\mathfrak{R}$ is convex. We derive how these credibility functions behave when the granularity of the algebra on which they are defined is modified, either by splitting the atoms or by regrouping them. The impact of a conditioning event that states that the actual world does not belong to some subset of Ω is derived. Whatever credibility functions are, we show that the impacts of the refinement, coarsening or conditioning processes are those described in the TBM where the credibility functions are belief functions. By introducing the concept of deconditionalization, i.e., eliminating the impact of an abusive conditioning, we prove that only belief functions are fitted for representing quantified beliefs. We have thus produce a set of rationality assumptions that justify the use of belief functions to represent quantified beliefs, and we show also that the set of probability functions is not rich enough to achieve that task.

7. The meaning of ‘belief’.

We present here several comments to show where the TBM departs from the classical probabilistic approach.

Consider the **medical diagnostic process**. Frequentists assume that the patient has been selected at random from the population of patients presenting the observed symptoms, an assumption usually void of any reality: the patient’s presence does not result from any random selection. Bayesians claim that probabilities appear because they describe the clinician a priori opinion about the disease his patient could be suffering from. From that a priori probability, other probabilities result after appropriate conditioning. This is the solution we would obtain in the TBM if such a priori probability was adequately representing the clinician’s a priori opinions. But this is exactly the point we are not accepting. We claim that a priori opinions are usually not adequately represented by probability functions, arguing belief functions are more adequate, even though the idea of ‘family of probability functions’ might be another alternative (Walley, 1991, Voorbraak, 1993). The fact that the patient comes from a population where there are 999 cases with disease A and one without does not mean this proportion is relevant to the clinician’s a priori belief about the fact his patient presents disease A. It would if the clinician knew the patient had indeed been selected at random in such a population. But we are studying the case where such a selection has not been used (or at least is not known by us to have been used). The belief functions we develop are quantifying the beliefs obtained in such general cases.

The measure of belief we study is analogous to the one encountered in judiciary context when culpability has to be assessed. Consider **the rodeo paradox** where out of 1000 persons who attend it, only one paid the entrance fee, the others having forced the gate. Police does not know who paid. Police arrests one person who attended the rodeo for not paying. I am the judge to whom the policeman brings the arrested person who claims - of course - he is the one who paid. If I had to bet on his culpability, I surely would bet with high probability on it, but this does not mean I believe that he is a culprit. I would bet he did not pay (because almost nobody paid) but I have no reason whatsoever to believe that this person did pay or not (because no evidence is brought forward that would justify such a belief). This difference between betting and belief parallels the difference we introduce between the pignistic and the credal levels. The quantification we focus at represents the strength of ‘good reasons’ in the expression ‘I have good reasons to believe’. In the TBM, we accordingly define bel(A) as the amount of ‘justified specific support’ given to A (Smets and Kennes, 1994).

The belief we study is not unsimilar to **the concept of provability**, and it has even been suggested that the degree of belief that a proposition is true represents the probability of proving its truth (Pearl, 1988), except the revision processes are more subtle than the one considered here (Smets, 1991). Indeed the underlying probability measure introduces extra constraints that must be handled appropriately.

In the hints model (Kohlas and Monney, 1995), the authors defends a similar understanding for the degree of support. Their approach is close to the TBM except that they still keep some links with the classical probability theory.

8. Conclusions.

We conclude this paper by pointing to potential applications of the transferable belief model.

The transferable belief model is a general model for quantified beliefs. The kind of application for which the transferable belief model is especially well suited covers diagnostic applications and data fusion problems. Indeed, there is no need to provide a probability to every atom of $\Re$ as in probability theory. Only the known information is fed into the model. No abusive probabilization is required. The transferable belief model is well adapted to represent state of partial or total ignorance that probability theory can hardly represent. There is no real counter part of the vacuous belief function in probability theory.

In Xu et al. (1993), we study a case of radioactivity leakage where there are several potential leaking sites and several locations where radioactivity tests can be performed. The input data included the belief that such and such test will be positive given that there is a leakage at this or that site. We even introduce an extra site whose existence and location are not even known (a secret site). Of course, for each test, the belief about the answers of a test given the leakage took place at the secret site is represented by a vacuous belief function. We also input the cost of performing each test and the test of deciding that the leakage took place at a given site given it really occurs at another site, for every pair of sites. We can then compute the belief about the location of the leakage and decide to clean such or such site. We would also establish the optimal strategy by assessing which is the best test to perform first. Then which test to perform next given the answer observed after the first test has been performed, etc.... All these computations simulate exactly what is commonly done in probability theory but they were based on the use of belief functions and of the pignistic transformation. They are based on the really available information and do not require the assessment of all these probabilities required by the classical probability models, probabilities that are often purely artificial. As an example of such artificial probabilities, in a well known medical diagnosis probabilistic program, the user is required to give a number to the probability that the patients suffers from “none of the considered symptoms” given he belongs to “none of the considered disease classes”. Does any value for such a probability really make sense?

Data fusion and expert opinions pooling is also an excellent domain of application thanks to the rule of combination and the concept of discounting (Smets, 1992d). Applications for data bases are presented in Smets and Kruse (1993).

Application for radar or sonar detection and recognition of mobile vehicles and for business investment decisions have been described. Computer software has been developed for this propagation of belief into belief networks and for optimal decision making (Xu, 1992, Xu et al., 1993). In all cases, the advantage of the transferable belief model is that it requires to feed into the system only what is really known. In the worst case, the computational complexity of the transferable belief model is higher than the one encountered with program based on probability theory. But in practice the complexity is the same or even smaller thanks to the fact that the input information in the transferable belief model is often much simpler. With the transferable belief model, the complexity is proportional to the information really available. With the probabilistic models, the complexity is proportional to the size of the frame, whatever information is really available. The last is often smaller than the former, in which case the TBM beats the probabilistic approach for what concerns the computational complexity.

In conclusion, the transferable belief model is a model to represent quantified beliefs. We hope this model might be useful for the understanding of the human thinking process and could be used as a normative model. We do not claim that this is the way humans do or should behave. Our approach is normative. It is neither descriptive nor prescriptive. It is but an idealized representation whose value can only be assessed by a critical examination of its axiomatic foundation. That it might be implemented in some 'thinking' robot able to cope with uncertainty and belief is not unthinkable.

Bibliography.

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

CARNAP R. (1962) Logical Foundations of Probability. University of Chicago Press, Chicago, Illinois.

CHOQUET G. (1953) Theory of capacities. Annales de l'Institut Fourier, Université de Grenoble, 5 131-296.

DEGROOT M.H. (1970) Optimal statistical decisions. McGraw-Hill, New York.

DELGADO M. and MORAL S. (1987) On the concept of possibility-probabilty consistency. Fuzzy Sets and Systems 21: 311-3018.

DUBOIS D., GARBOLINO P., KYBURG H.E., PRADE H. and SMETS Ph. (1991) Quantified Uncertainty. J. Applied Non-Classical Logics 1:105-197.

DUBOIS D. and PRADE H. (1986) A set theoretical view of belief functions. Int. J. Gen. Systems, 12:193-226.

DUBOIS D. and PRADE H. (1987) The principle of minimum specificity as a basis for evidential reasoning. in: Uncertainty in knowledge-based systems, BOUCHON B. and YAGER R. eds, Springer Verlag, Berlin, pg.75-84.

DUBOIS D., PRADE H. and SMETS Ph. (1995) Representing partial ignorance. under press, IEEE System Machine and Cybernetic.

EKELOF P.O. (1982) Rättegång IV. Fifth edition, Stockholm.

FINE T. (1973) Theories of probability. Academic Press, New York.

GÄRDENFORS P., HANSSON B. and SAHLIN N.E. (1983) Evidentiary value: philosophical, judicial and psychological aspects of a theory. C.W.K. Gleerups, Lund,

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

GILES R. (1982) Foundation for a possibility theory. in Fuzzy Information and Decision Processes, GUPTA M.M. and SANCHEZ E. eds. North Holland, Amsterdam, 183-195.

HSIA Y.-T. (1991) Characterizing Belief with Minimum Commitment. IJCAI-91:1184-1189.

JEFFREY R. (1988) Conditioning, kinematics, and exchangeability. in Causation, Chance, and Credence, Skyrms B. and Harper W.L. eds., Reidel, Dordrecht, vol.1, 221-255.

KENNES R.(1991) Evidential Reasoning in a Categorial Perspective : Conjunction and Disjunction of Belief Functions. in D'Ambrosio B., Smets P., and Bonissone P.P. eds, Uncertainty in AI 91,Morgan Kaufmann, San Mateo, Ca, USA, 1991, pg.174-181

KENNES R. (1992) Computational aspects of the Moebius transform of a graph. IEEE-SMC, 22: 201-223.

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.

KOHLAS J. and MONNEY P.A. (1995) A Mathematical Theory of Hints. An Approach to Dempster-Shafer Theory of Evidence. Lecture Notes in Economics and Mathematical Systems No. 425. Springer-Verlag.

KOOPMAN B.O. (1940) The bases of probability. Bull. Amer. Math. Soc. 46:763-774.

KRUSE R. and SCHWECKE E. (1990) Specialization: a new concept for uncertainty handling with belief functions. Int. J. Gen. Systems 18:49-60.

PAL N., BEZDEK J. and HEMASINHA R. (1992) Uncertainty measures for evidential reasoning I: a review. Int. J. Approx. Reas. 7:165-183.

PEARL J. (1988) Probabilistic reasoning in intelligent systems: networks of plausible inference. Morgan Kaufmann Pub. San Mateo, Ca, USA.
RAMSEY F.P. (1931) Truth and probability. in Studies in subjective probability, eds. KYBURG H.E. and SMOKLER H.E., p. 61-92. Wiley, New York.
RUSPINI E.H. (1986) The logical foundations of evidential reasoning. Technical note 408, SRI International, Menlo Park, Ca.
SAVAGE L.J. (1954) Foundations of Statistics. Wiley, New York.
SHAFER G. (1976) A mathematical theory of evidence. Princeton Univ. Press. Princeton, NJ.
SHAFER G. (1992) Rejoinder to Comments on "Perspectives in the theory and practice of belief functions". Intern. J. Approx. Reasoning, 6:445-480.
SHAFER G. and TVERSKY A. (1985) Laages and designs for probability. Cognitive Sc. 9:309-339.
SHAPLEY L.S. (1953) A value for n-person games. In Contributions to the Theory of Games, vol. 2, eds. H. Kuhn and A.W. Tucker. Princeton University Press, pp. 307-317.
SMETS Ph. (1978) Un modèle mathématico-statistique simulant le processus du diagnostic médical. Doctoral dissertation, Université Libre de Bruxelles, Bruxelles, (Available through University Microfilm International, 30-32 Mortimer Street, London W1N 7RA, thesis 80-70,003)
SMETS Ph. (1981) Medical diagnosis: fuzzy sets and degree of belief. Fuzzy Sets and Systems 5:259-266.
SMETS Ph. (1983) Information Content of an Evidence. Int. J. Man Machine Studies, 19: 33-43.
SMETS Ph. (1988) Belief functions. in SMETS Ph, MAMDANI A., DUBOIS D. and PRADE H. ed. Non standard logics for automated reasoning. Academic Press, London pg 253-286.
SMETS Ph. (1990a) The combination of evidence in the transferable belief model. IEEE-Pattern analysis and Machine Intelligence, 12:447-458.
SMETS Ph. (1990b) Constructing the pignistic probability function in a context of uncertainty. Uncertainty in Artificial Intelligence 5, Henrion M., Shachter R.D., Kanal L.N. and Lemmer J.F. eds, North Holland, Amsterdam, , 29-40.
SMETS Ph. (1991) Probability of provability and belief functions. Logique et Analyse, 133-134:177-195.
SMETS P. (1992a) The nature of the unnormalized beliefs encountered in the transferable belief model. in Dubois D., Wellman M.P., d'Ambrosio B. and Smets P. Uncertainty in AI 92. Morgan Kaufmann, San Mateo, Ca, USA, 1992, pg.292-297.
SMETS P. (1992b) The transferable belief model and random sets. Int. J. Intell. Systems 7:37-46.
SMETS P. (1992c) The concept of distinct evidence. IPMU 92 Proceedings.pg. 789-794.
SMETS P. (1992d) The transferable belief model for expert judgments and reliability problmes. Reliability Engineering and System Safety. 38:59-66.
SMETS Ph. (1993a) No Dutch Book can be built against the transferable belief model even though update is not obtained by Bayes rule of conditioning. SIS, Workshop on Probabilisitic Expert Systems, (ed. R. Scozzafava), Roma, pg. 181-204..
SMETS P. (1993b) Belief functions: the disjunctive rule of combination and the generalized Bayesian theorem. Int. J. Approximate Reasoning 9:1-35.
SMETS P. (1993c) An axiomatic justifiaction for the use of belief function to quantify beliefs. IJCAI'93 (Inter. Joint Conf. on AI), San Mateo, Ca, pg. 598-603.
SMETS Ph. (1994) What is Dempster-Shafer's model? in Advances in the Dempster-Shafer Theory of Evidence. Yager R.R., Kacprzyk J. and Fedrizzi M., eds, Wiley, New York, pg. 5-34.
SMETS Ph. (1995) The representation of quantified belief by belief functions: an axiomatic justificiation. TR/IRIDIA/95.8.2. Submitted for publication.
SMETS Ph. and KENNES R. (1994) The transferable belief model. Artificial Intelligence 66:191-234.

SMETS Ph. and KRUSE R. (1993) The transferable belief model for belief representation. in Smets Ph. and Motro A. eds. (1993) Uncertainty Management in information systems: from needs to solutions. UMIS Workshop, Puerto Andraix.Volume 2. pg 91-111.

SMITH C.A.B. (1961) Consistency in statistical inference and decision. J. Roy. Statist. Soc. B23:1-37.

TELLER P. (1973) Conditionalization and Observation. Synthesis 26:218-258.

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.

WILSON N. (1993) Decision making with belief functions and pignistic probabilities. in Symbolic and Quantitative Approaches to Reasoning and Uncertainty. Clarke M., Kruse R. and Moral S. (eds.), Springer Verlag, Berlin, pg. 364-371.

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.

XU H. (1992) An efficient tool for reasoning with belief functions. IPMU 92 Proceedings.pg. 65-68.

XU H., HSIA Y.T. and SMETS Ph. (1993) A belief function based decision support system. Uncertainty in AI 93. Heckerman D. and Mamdani A. eds. Morgan Kaufmann, San Mateo, Ca, USA, 535-542.

XU H. and SMETS Ph. (1994) Evidential reasoning with conditional belief functions. in Heckerman D., Poole D. and Lopez de Mantaras R. (eds.) Uncertainty in AI 94, Morgan Kaufmann, San Mateo, Ca.

YAGER R. (1986) The entailment principle for Dempster-Shafer granules. Int. J. Intell. Systems 1:247-262