Belief models: An order-theoretic investigation

Gert de Cooman (gert.decooman@ugent.be)
Universiteit Gent, Onderzoeksgroep SYSTeMS, Technologiepark – Zwijnaarde 914, 9052 Zwijnaarde, Belgium

September 7, 2010

Abstract. I show that there is a common order-theoretic structure underlying many of the models for representing beliefs in the literature. After identifying this structure, and studying it in some detail, I argue that it is useful. On the one hand, it can be used to study the relationships between several models for representing beliefs, and I show in particular that the model based on classical propositional logic can be embedded in that based on the theory of coherent lower previsions. On the other hand, it can be used to generalise the coherentist study of belief dynamics (belief expansion and revision) by using an abstract order-theoretic definition of the belief spaces where the dynamics of expansion and revision take place. Interestingly, many of the existing results for expansion and revision in the context of classical propositional logic can still be proven in this much more abstract setting, and therefore remain valid for many other belief models, such as those based on imprecise probabilities.

Keywords: Belief model, belief revision, classical propositional logic, imprecise probability, order theory, possibility measure, system of spheres

1. Introduction

It is often claimed that epistemic probability theory extends classical propositional logic from a logic of certainty to a logic of partial belief (see for instance [12]). I intend to argue here that in some definite sense this claim is invalid. But that is by no means my only, nor my main purpose with this paper. Indeed, I also want to point out that many of the models in the literature for representing a subject's beliefs (and utilities) share the same interesting order-theoretic framework.

In the first part of this paper, I identify this framework by concentrating on the inference methods behind such models. This leads me to the introduction of what I call *belief structures*. Roughly speaking, they are special collections of abstract entities called *belief models*, which share a number of (order-theoretic) properties, and which could also be seen as generalised 'epistemic states' [15]. I believe the study of these belief structures to be interesting and important, because it allows us to find out how different existing models for representing beliefs are related to each other.

In the second part, I suggest that these abstract belief models—or perhaps some more developed version of them—can form the basis for, or can at least be helpful in, a generalised study of the dynamics of epistemic states, which has received much attention in the AI literature since the publication of Gärdenfors's book [15] on belief change, already some time ago. What can be said about the dynamics of belief change when the space of epistemic states in which this dynamics takes place is more general than that considered by Gärdenfors and others? What happens when the epistemic states we are interested in are not sets of sentences in classical propositional logic, but possibility distributions [8] or ordinal conditional functions [31], sets of probability measures [17, 20], lower previsions [32], preference orderings on horse lotteries [29], and so on? Below, I make an attempt at beginning to provide a uniform answer to this question, by only assuming that the space of epistemic states satisfies the unifying properties of an abstract belief structure, and in this context deriving a number of interesting general results for belief expansion and belief revision. Of course, these results remain valid for the various instances of belief structures found in the literature.

The order-theoretic structures and notions that I want to draw attention to, are introduced and studied in Sections 2, 3 and 5. They are based directly on certain aspects of the inference mechanism present in classical propositional logic, but I show in Section 4 that there are numerous other important instances of these structures in the literature on uncertainty modelling. In Section 5, I introduce a number of notions that allow us to study the relationships between different belief structures, and I show in particular that the belief structure based on classical propositional logic can be embedded into that based on the theory of coherent lower previsions [32].

The rest of the paper deals with the dynamics of epistemic states: Section 6 is concerned with expansion of belief models. Revision is discussed in Section 7, and in Sections 8 and 9, which focus on specific ways to construct revision operators. Section 10 concludes the paper.

I will make no effort to define or explain the many mathematical notions borrowed from order theory, as most of them are (or deserve to be) well-known. I refer to a good introductory treatment (such as [6]) instead.

2. Belief structures

2.1. Classical propositional logic

Consider an object language **L** of well-formed formulae, or sentences, in classical propositional logic with the usual axiomatisation (see for instance [6, 28]). We call any subset K of **L**, *i.e.*, any set of sentences, a *belief model*.[1] Intuitively, a set of sentences K models the beliefs of a subject: it contains those sentences that the subject is certain are true. Of course, this is a very simple type of model, because it concentrates on certainty, or full belief. We shall want to study more general models, that are also able to represent a

[1] Gärdenfors [15] speaks of an *epistemic state*.

subject's uncertainty, or his partial beliefs. But for didactic reasons, I shall in this and the next section start with the well-known example of classical propositional logic, extract from it those aspects of its reasoning mechanism that are essential elements of a possible generalisation, and use them to formulate an abstract and more general notion of belief model. In Section 4, I shall list a multitude of concrete and interesting other instances of such models.

The collection of all belief models in classical propositional logic is the power set $\wp(\mathbf{L})$, consisting of all subsets of the object language $\mathbf{L}$. It can be partially ordered by set inclusion $\subseteq$, and the structure $\langle \wp(\mathbf{L}), \subseteq \rangle$ is a complete lattice, whose greatest element is $\mathbf{L}$ and whose smallest element is the empty set $\emptyset$. In this complete lattice, union plays the role of supremum, and intersection that of infimum.

There are two distinct ways in which a set of sentences K may be imperfect. It may (i) be contradictory or inconsistent, and it may (ii) not contain all the logical consequences of the sentences it contains. Both types of imperfection can be investigated more systematically if we look at the notion of *logical*, or *deductive*, *closure*.

Recall that a set of sentences is called *logically closed* if it is closed under conjunction and modus ponens, or in other words if it contains all the logical consequences of its sentences. It turns out—and this is of crucial importance for what follows—that the intersection (the infimum for $\subseteq$) of a collection of logically closed sets is still logically closed. This means that the logically closed sets constitute a Moore family [6, 28]. It is well known that we can associate a closure operator (or Moore closure[2].) with a Moore family. In the present case, this is nothing but the logical closure operator: the Moore closure of a set of sentences is indeed the smallest logically closed set of sentences that includes it, or equivalently, the intersection of all logically closed sets of sentences that include it. This closure operator represents the essence of the inference mechanism behind classical propositional logic at the level of belief models, *i.e.*, when we focus on sets of sentences rather than sentences themselves.

Among the logically closed sets of sentences, there is only one that represents contradiction: $\mathbf{L}$ contains every sentence as well as its negation! A subject whose belief model is $\mathbf{L}$, is certain of every sentence, and of its negation. This model is clearly to be avoided. For this reason, a set of sentences is often called (logically) *consistent* if its logical closure is different from $\mathbf{L}$. If a set of sentences is at the same time consistent and logically closed, I shall call it *coherent*. Coherence is the type of perfection that we are after. A coherent set of sentences is what logicians sometimes call a *theory*. We shall denote the collection of coherent sets of sentences by $\mathbf{C}_\mathbf{L}$ It is quite interesting to

[2] Observe that the properties of a Moore closure are in one-to-one correspondence to those of a Tarski consequence operation, see for instance [23]

observe that (i) the partial order $\subseteq$ for coherent models can be interpreted as 'is less informative than'; and that (ii) the intersection of a *non-empty* class of coherent sets of sentences is still coherent.

We are now ready to extend the notion of as belief model from a set of sentences to something more general and powerful, in a manner based on the observations just made.

2.2. Generalisation

Consider a non-empty set $\mathbf{S}$, whose elements are called *belief models*. They are partially ordered by a relation $\leq$ that is reflexive, transitive and antisymmetric, but need not be complete: it is not required that any two elements a and b of $\mathbf{S}$ should be comparable in the sense that $a \leq b$ or $b \leq a$. A first important assumption is that for any subset A of $\mathbf{S}$, its supremum $\sup A$ and infimum $\inf A$ with respect to this order exist, or in other words:[3]

S1. $\langle \mathbf{S}, \leq \rangle$ is a complete lattice.

Let us denote by $1_{\mathbf{S}}$ the top, or greatest element, $\sup \mathbf{S}$ of this complete lattice. Its bottom, or smallest element, $\inf \mathbf{S}$ is denoted by $0_{\mathbf{S}}$. Note that also $1_{\mathbf{S}} = \inf \emptyset$ and $0_{\mathbf{S}} = \sup \emptyset$. The supremum, or join, of two belief models a and b is also denoted by $a \smile b$ and their infimum, or meet, by $a \frown b$.

Among the belief models in $\mathbf{S}$, there is a subset $\mathbf{C} \subseteq \mathbf{S}$ of models that are called *coherent*. Coherent belief models are considered to be more perfect than the others, which will be called *incoherent*. It is obvious that $\mathbf{C}$ inherits the partial order $\leq$ from $\mathbf{S}$. A second central assumption is that:

S2. $\mathbf{C}$ is closed under arbitrary non-empty infima: for any non-empty subset C of $\mathbf{C}$, $\inf C \in \mathbf{C}$.

The belief model $1_{\mathbf{S}}$ will represent contradiction, so we assume that:

S3. The partially ordered set $\langle \mathbf{C}, \leq \rangle$ has no top. In particular, $1_{\mathbf{S}}$ is not a coherent belief model: $1_{\mathbf{S}} \notin \mathbf{C}$.

This means that the ordered structure $\langle \mathbf{C}, \leq \rangle$ is a complete meet-semilattice but not a complete lattice: every non-empty subset of $\mathbf{C}$ has an infimum but not necessarily a supremum in this structure. On the other hand, the set $\overline{\mathbf{C}} = \mathbf{C} \cup \{1_{\mathbf{S}}\}$ provided with the ordering $\leq$ *is* a complete lattice, whose infimum (but not necessarily its supremum) coincides with the infimum of $\langle \mathbf{S}, \leq \rangle$. The relation $\leq$ on $\mathbf{C}$ is taken to mean roughly 'is less informative than'.

[3] We could require here that $\langle \mathbf{S}, \leq \rangle$ should be a completely distributive complete lattice. I have not done so, as I see no direct need for it in the context of the present work.

Definition 1 (Belief structure)**.** If $\langle \mathbf{S}, \leq \rangle$ and $\mathbf{C}$ satisfy requirements S1–S3, then we call the triple $\langle \mathbf{S}, \mathbf{C}, \leq \rangle$ a *belief structure*.

We can now introduce a *closure operator* $\mathrm{Cl}_{\mathbf{S}} \colon \mathbf{S} \to \mathbf{S}$ as follows: for any belief model b in $\mathbf{S}$,

$$\mathrm{Cl}_{\mathbf{S}}(b) = \inf\{c \in \overline{\mathbf{C}} \colon b \leq c\}.$$

Note that if b is dominated by some coherent belief model, then $\mathrm{Cl}_{\mathbf{S}}(b)$ is the smallest coherent belief model that dominates b. The operator $\mathrm{Cl}_{\mathbf{S}}$ has the following immediate properties.

Proposition 1. *Let $\langle \mathbf{S}, \mathbf{C}, \leq \rangle$ be a belief structure. For any belief models a and b in $\mathbf{S}$,*

1. $a \leq \mathrm{Cl}_{\mathbf{S}}(a)$*;*
2. *if $a \leq b$ then* $\mathrm{Cl}_{\mathbf{S}}(a) \leq \mathrm{Cl}_{\mathbf{S}}(b)$*;*
3. $\mathrm{Cl}_{\mathbf{S}}(\mathrm{Cl}_{\mathbf{S}}(a)) = \mathrm{Cl}_{\mathbf{S}}(a)$*;*
4. $\mathrm{Cl}_{\mathbf{S}}(a \smile b) = \mathrm{Cl}_{\mathbf{S}}(\mathrm{Cl}_{\mathbf{S}}(a) \smile \mathrm{Cl}_{\mathbf{S}}(b))$*;*
5. $\mathrm{Cl}_{\mathbf{S}}(a) = a$ *if and only if* $a \in \overline{\mathbf{C}}$*;*

This justifies our calling $\mathrm{Cl}_{\mathbf{S}}$ a closure operator. The associated set of *closed* belief models, *i.e.* those belief models a for which $a = \mathrm{Cl}_{\mathbf{S}}(a)$, is precisely $\overline{\mathbf{C}}$. The underlying idea is that for any belief model $a \in \mathbf{S}$, a and its closure $\mathrm{Cl}_{\mathbf{S}}(a)$ are equally informative. The closure $\mathrm{Cl}_{\mathbf{S}}$ takes any belief model a with $\mathrm{Cl}_{\mathbf{S}}(a) < 1_{\mathbf{S}}$ into a coherent belief model $\mathrm{Cl}_{\mathbf{S}}(a)$ that is equally informative.

Observe that we do not require that $\mathrm{Cl}_{\mathbf{S}}(0_{\mathbf{S}}) = 0_{\mathbf{S}}$: $b_v = \mathrm{Cl}_{\mathbf{S}}(0_{\mathbf{S}})$ is the smallest coherent belief model, also called the *vacuous belief model*.

The closure operator $\mathrm{Cl}_{\mathbf{S}}$ allows us to give an expression for the supremum in the complete lattice $\langle \overline{\mathbf{C}}, \leq \rangle$: for any subset C of $\overline{\mathbf{C}}$, its supremum in this structure is given by $\mathrm{Cl}_{\mathbf{S}}(\sup C)$, where $\sup C$ is the supremum of C in the complete lattice $\langle \mathbf{S}, \leq \rangle$.

Recall that the top $1_{\mathbf{S}}$ is assumed to represent contradiction, or inconsistency. The closure operator $\mathrm{Cl}_{\mathbf{S}}$ allows us to take this a step further.

Definition 2 (Consistency)**.** A belief model $a \in \mathbf{S}$ is called *consistent* if $\mathrm{Cl}_{\mathbf{S}}(a) \in \mathbf{C}$, that is, if $\mathrm{Cl}_{\mathbf{S}}(a) < 1_{\mathbf{S}}$. Two belief models a and b in $\mathbf{S}$ are said to be *consistent* (with one another) if $a \smile b$ is consistent. More generally, a collection $S \subseteq \mathbf{S}$ of belief models is called *consistent* if $\sup S$ is a consistent belief model.

Coherent belief models are in particular consistent. In fact, they are the ones that are both consistent and closed. The following proposition explicates a further relationship between coherence and consistency: the consistent belief models are the ones that are below some coherent belief model.

Proposition 2. *Let $\langle \mathbf{S}, \mathbf{C}, \leq \rangle$ be a belief structure. For any belief model a in* $\mathbf{S}$*, the following statements are equivalent:*

1. *a is consistent;*

2. $\mathrm{Cl}_{\mathbf{S}}(a) < 1_{\mathbf{S}}$*;*

3. $a \leq b$ *for some coherent belief model* $b \in \mathbf{C}$.

Proof. Assume that a is consistent. Then $\mathrm{Cl}_{\mathbf{S}}(a) \in \mathbf{C}$ whence $\mathrm{Cl}_{\mathbf{S}}(a) < 1_{\mathbf{S}}$. Next assume that $\mathrm{Cl}_{\mathbf{S}}(a) < 1_{\mathbf{S}}$. Then $a \leq \mathrm{Cl}_{\mathbf{S}}(a)$ and $\mathrm{Cl}_{\mathbf{S}}(a) \in \mathbf{C}$, which means that the third statement holds. Finally, assume that the third statement holds. Then by Proposition 1, $\mathrm{Cl}_{\mathbf{S}}(a) \leq \mathrm{Cl}_{\mathbf{S}}(b) = b < 1_{\mathbf{S}}$, whence $\mathrm{Cl}_{\mathbf{S}}(a) \in \mathbf{C}$, so a is consistent.

A belief model is inconsistent if closure takes it to the contradictory model $1_{\mathbf{S}}$. Note that $1_{\mathbf{S}}$ is the only contradictory or inconsistent model in $\overline{\mathbf{C}}$. In summary, the idea behind closure is that it takes $\mathbf{S}$ to the informationally equivalent structure $\overline{\mathbf{C}}$, where $1_{\mathbf{S}}$ is the only inconsistent model. Also note that if a and b are consistent, then $\mathrm{Cl}_{\mathbf{S}}(a \smile b)$ is a coherent belief model, and it is the supremum of $\mathrm{Cl}_{\mathbf{S}}(a)$ and $\mathrm{Cl}_{\mathbf{S}}(b)$ in the complete meet-semilattice $\langle \mathbf{C}, \leq \rangle$. It can be interpreted as the least informative coherent belief model that is at least as informative as a and b.

3. Strong belief structures and their duals

3.1. Classical propositional logic

If we start with a logically closed set of sentences K, then we can imagine trying to make it 'more informative' by adding sentences to it. At some point in this process, however, it will become impossible to add another sentence without creating a contradiction, *i.e.*, making the result logically inconsistent. To render this idea more precise, a *maximal* logically closed set of sentences, or a *complete* theory, K is defined as a logically closed set of sentences such that any set of sentences that strictly includes it is logically inconsistent.

By applying the Boolean Ultrafilter Theorem [6, 28, 32] to the Lindenbaum algebra associated with $\mathbf{L}$ [6, 28], we may infer (i) that there indeed are such maximal coherent sets of sentences; and (ii) that every coherent set of sentences is the intersection of the maximal coherent sets including it. This observation marks the starting point for a further refinement of the abstract notion of a belief structure.

3.2. GENERALISATION

Consider a belief structure $\langle \mathbf{S}, \mathbf{C}, \leq \rangle$. Recall that there is no greatest (or most informative) coherent belief model: the partially ordered set $\langle \mathbf{C}, \leq \rangle$ has no top. But it may have maximal elements, that is, elements m that are not dominated by any other element of $\mathbf{C}$. I denote by $\mathbf{M}$ the (possibly empty) set of these maximal elements:

$$\mathbf{M} = \{m \in \mathbf{C} \colon (\forall c \in \mathbf{C})(m \leq c \Rightarrow m = c)\}.$$

We can render the notion of a belief structure much more powerful by making an extra assumption, which concerns precisely these maximal elements. We may require that they can be used to construct any coherent belief model:

S4. The partially ordered set $\langle \mathbf{C}, \leq \rangle$ is *dually atomic*: $\mathbf{M} \neq \emptyset$ and for all $c \in \mathbf{C}$,

$$c = \inf\{m \in \mathbf{M} \colon c \leq m\}.$$

Definition 3 (Strong belief structure)**.** A belief structure $\langle \mathbf{S}, \mathbf{C}, \leq \rangle$ for which the additional requirement S4 is satisfied, is called a *strong belief structure*.

We also introduce the following notation: for any belief model $b \in \mathbf{S}$,

$$\mathscr{M}(b) = \{m \in \mathbf{M} \colon b \leq m\}$$

is the set of all maximal belief models dominating b. $\mathscr{M}(\cdot)$ can be interpreted as a map from $\mathbf{S}$ to the power set $\wp(\mathbf{M})$ of $\mathbf{M}$. It will play an important part in the investigation of the structure of $\langle \mathbf{S}, \mathbf{C}, \leq \rangle$.

If $b \in \mathbf{C}$ then S4 implies that $\mathscr{M}(b) \neq \emptyset$. For if $\mathscr{M}(b) = \emptyset$, then $\inf \mathscr{M}(b) = 1_{\mathbf{S}} > b$. In other words, in a strong belief structure every coherent belief model is dominated by at least one maximal coherent belief model. Moreover, there is the following extension of Proposition 2.

Proposition 3. *Let* $\langle \mathbf{S}, \mathbf{C}, \leq \rangle$ *be a strong belief structure, and let a be a belief model in* $\mathbf{S}$*. Then each of the three statements in Proposition 2 is equivalent to* $\mathscr{M}(a) \neq \emptyset$*.*

Proof. It suffices to prove that $\mathscr{M}(a) \neq \emptyset$ is equivalent to the third statement. Assume that there is a $b \in \mathbf{C}$ such that $a \leq b$. It follows from the definition of $\mathscr{M}(\cdot)$ that $\mathscr{M}(b) \subseteq \mathscr{M}(a)$. We have argued above that for $b \in \mathbf{C}$, S4 implies that $\mathscr{M}(b) \neq \emptyset$, whence $\mathscr{M}(a) \neq \emptyset$. Conversely, if $\mathscr{M}(a) \neq \emptyset$, then $a \leq \inf \mathscr{M}(a)$ and $\inf \mathscr{M}(a) \in \mathbf{C}$.

There is a very close relationship between the closure operator $\mathrm{Cl}_{\mathbf{S}}$ and the map $\mathscr{M}(\cdot)$.

Proposition 4. *Let $\langle \mathbf{S}, \mathbf{C}, \leq \rangle$ be a strong belief structure. Then for all $a \in \mathbf{S}$:*

1. $\mathscr{M}(a) = \mathscr{M}(\mathrm{Cl}_{\mathbf{S}}(a))$*;*
2. $\mathrm{Cl}_{\mathbf{S}}(a) = \inf \mathscr{M}(a)$*;*
3. $a \in \overline{\mathbf{C}} \Leftrightarrow a = \inf \mathscr{M}(a)$.

Proof. We begin with the first statement. It follows from Proposition 1 that for all $b \in \mathbf{C}$, $a \leq b \Leftrightarrow \mathrm{Cl}_{\mathbf{S}}(a) \leq b$, and since $\mathbf{M} \subseteq \mathbf{C}$ it follows that for all $m \in \mathbf{M}$, $m \in \mathscr{M}(a) \Leftrightarrow m \in \mathscr{M}(\mathrm{Cl}_{\mathbf{S}}(a))$. We continue with the second statement. First, assume that a is inconsistent. Then on the one hand $\mathrm{Cl}_{\mathbf{S}}(a) = 1_{\mathbf{S}}$ and on the other hand $\mathscr{M}(a) = \emptyset$, by Propositions 2 and 3. So in this case, $\inf \mathscr{M}(a) = \inf \emptyset = 1_{\mathbf{S}} = \mathrm{Cl}_{\mathbf{S}}(a)$. Next, assume that a is consistent. Then the first statement $\mathscr{M}(a) = \mathscr{M}(\mathrm{Cl}_{\mathbf{S}}(a))$ implies that $\inf \mathscr{M}(a) = \inf \mathscr{M}(\mathrm{Cl}_{\mathbf{S}}(a)) = \mathrm{Cl}_{\mathbf{S}}(a)$, taking into account S4 and the fact that $\mathrm{Cl}_{\mathbf{S}}(a) \in \mathbf{C}$, by Propositions 2 and 3. The third statement is an immediate consequence of the second.

It will be very important to pay special attention to the direct images of the sets $\mathbf{C}$ and $\overline{\mathbf{C}}$ under the map $\mathscr{M}(\cdot)$:

$$\overline{\mathfrak{M}} = \mathscr{M}(\overline{\mathbf{C}}) = \{\mathscr{M}(c) \colon c \in \overline{\mathbf{C}}\} = \mathscr{M}(\mathbf{S})$$
$$\mathfrak{M} = \mathscr{M}(\mathbf{C}) = \{\mathscr{M}(c) \colon c \in \mathbf{C}\}.$$

Both $\mathfrak{M}$ and $\overline{\mathfrak{M}}$ are subsets of $\wp(\mathbf{M})$, *i.e.*, sets of subsets of $\mathbf{M}$. Moreover, $\mathscr{M}(b_v) = \mathbf{M}$, so $\mathbf{M} \in \mathfrak{M}$. Also, $\mathscr{M}(1_{\mathbf{S}}) = \emptyset$ belongs to $\overline{\mathfrak{M}}$ but not to $\mathfrak{M}$, whence $\mathfrak{M} = \overline{\mathfrak{M}} \setminus \{\emptyset\}$.

A crucial property of $\overline{\mathfrak{M}}$ is that it is an intersection structure with top $\mathbf{M}$, or in other words that it is closed under arbitrary (also empty) intersections. Consequently, the partially ordered set $\langle \overline{\mathfrak{M}}, \subseteq \rangle$ is a complete lattice, where intersection has the role of infimum. This is made more explicit in the following theorem.

Theorem 5. *Let $\langle \mathbf{S}, \mathbf{C}, \leq \rangle$ be a strong belief structure. Then the following propositions hold.*

1. *$\overline{\mathfrak{M}}$ is a Moore collection of subsets of $\mathbf{M}$ [6, 28]: it is closed under arbitrary (and therefore also empty) intersections.*
2. *The complete lattices $\langle \overline{\mathbf{C}}, \leq \rangle$ and $\langle \overline{\mathfrak{M}}, \subseteq \rangle$ are dually order-isomorphic, with dual order isomorphism $\mathscr{M}(\cdot)$.*
3. *Consider the operator $\mathrm{Cl}_{\mathbf{M}} \colon \wp(\mathbf{M}) \to \wp(\mathbf{M})$ that is defined by $\mathrm{Cl}_{\mathbf{M}}(\mathscr{N}) = \mathscr{M}(\inf \mathscr{N})$ for all $\mathscr{N} \subseteq \mathbf{M}$. Then $\mathrm{Cl}_{\mathbf{M}}$ is a Moore closure [6, 28] and $\overline{\mathfrak{M}}$ is the associated set of closed sets: $\overline{\mathfrak{M}} = \{\mathscr{N} \subseteq \mathbf{M} \colon \mathrm{Cl}_{\mathbf{M}}(\mathscr{N}) = \mathscr{N}\}$.*

4. All singletons $\{m\}$, $m \in \mathbf{M}$, are closed.

Proof. We start with the first statement. Let $\{\mathscr{N}_j\colon j \in J\}$ be a family of elements of $\overline{\mathfrak{M}}$. If $J = \emptyset$, then $\bigcap_{j\in J}\mathscr{N}_j = \mathbf{M} = \mathscr{M}(b_v) \in \overline{\mathfrak{M}}$. If $J \neq \emptyset$, then let $b_j = \inf\mathscr{N}_j$ for all $j \in J$. Obviously $b_j \in \overline{\mathbf{C}}$, whence by Proposition 4, $\mathscr{N}_j = \mathscr{M}(b_j)$ for all $j \in J$ [observe that since $\mathscr{N}_j \in \overline{\mathfrak{M}}$, $\mathscr{N}_j = \mathscr{M}(c_j)$ for some c_j in $\overline{\mathbf{C}}$, and use the proposition to show that $c_j = b_j$]. Consequently,

$$\bigcap_{j\in J}\mathscr{N}_j = \bigcap_{j\in J}\mathscr{M}(b_j) = \mathscr{M}(\sup_{j\in J} b_j) \in \overline{\mathfrak{M}},$$

since $\sup_{j\in J} b_j \in \mathbf{S}$. We now turn to the second statement. Consider b_1 and b_2 in $\overline{\mathbf{C}}$. First, if $b_1 \leq b_2$ then obviously $\mathscr{M}(b_2) \subseteq \mathscr{M}(b_1)$. Conversely, if $\mathscr{M}(b_2) \subseteq \mathscr{M}(b_1)$, it follows from Proposition 4 that

$$b_1 = \inf\mathscr{M}(b_1) \leq \inf\mathscr{M}(b_2) = b_2.$$

So we conclude that

$$b_1 \leq b_2 \Leftrightarrow \mathscr{M}(b_2) \subseteq \mathscr{M}(b_1). \tag{1}$$

This means that $\mathscr{M}(\cdot)$ is a dual order embedding of $\langle\overline{\mathbf{C}},\leq\rangle$ into $\langle\overline{\mathfrak{M}},\subseteq\rangle$. It is furthermore surjective, since $\overline{\mathfrak{M}} = \mathscr{M}(\overline{\mathbf{C}})$. We conclude that $\mathscr{M}(\cdot)$ is indeed a dual order isomorphism. To prove the third statement, we first show that $\mathrm{Cl}_{\mathbf{M}}$ satisfies the defining properties of a Moore closure. Consider a subset $\mathscr{N}$ of $\mathbf{M}$. For any $m \in \mathscr{N}$ we have that $\inf\mathscr{N} \leq m$ so $m \in \mathscr{M}(\inf\mathscr{N})$, and therefore $\mathscr{N} \subseteq \mathrm{Cl}_{\mathbf{M}}(\mathscr{N})$. Moreover, $a = \inf\mathscr{N} \in \overline{\mathbf{C}}$, so $a = \inf\mathscr{M}(a) = \inf\mathscr{M}(\inf\mathscr{N})$ by Proposition 4. Consequently, $\mathrm{Cl}_{\mathbf{M}}(\mathscr{N}) = \mathscr{M}(a) = \mathscr{M}(\inf\mathscr{M}(\inf\mathscr{N})) = \mathrm{Cl}_{\mathbf{M}}(\mathrm{Cl}_{\mathbf{M}}(\mathscr{N}))$. Finally, for any subsets $\mathscr{N}$ and $\mathscr{S}$ of $\mathbf{M}$ such that $\mathscr{N} \subseteq \mathscr{S}$, we have $\inf\mathscr{S} \leq \inf\mathscr{N}$, whence $\mathscr{M}(\inf\mathscr{N}) \subseteq \mathscr{M}(\inf\mathscr{S})$, or $\mathrm{Cl}_{\mathbf{M}}(\mathscr{N}) \subseteq \mathrm{Cl}_{\mathbf{M}}(\mathscr{S})$. This means that $\mathrm{Cl}_{\mathbf{M}}$ is indeed a Moore closure. We now look at its associated set of closed sets. Consider a subset $\mathscr{N}$ of $\mathbf{M}$. Then it follows from $\mathscr{N} = \mathrm{Cl}_{\mathbf{M}}(\mathscr{N})$ that $\mathscr{N} = \mathscr{M}(\inf\mathscr{N})$, so $\mathscr{N} \in \overline{\mathfrak{M}}$ since $\inf\mathscr{N} \in \overline{\mathbf{C}}$. Conversely, if $\mathscr{N} \in \overline{\mathfrak{M}}$, then $\mathscr{N} = \mathscr{M}(a)$ for some $a \in \overline{\mathbf{C}}$, and by Proposition 4, $a = \inf\mathscr{M}(a) = \inf\mathscr{N}$. Consequently, $\mathscr{N} = \mathscr{M}(a) = \mathscr{M}(\inf\mathscr{N}) = \mathrm{Cl}_{\mathbf{M}}(\mathscr{N})$. This means that $\overline{\mathfrak{M}} = \{\mathscr{N} \subseteq \mathbf{M}\colon \mathrm{Cl}_{\mathbf{M}}(\mathscr{N}) = \mathscr{N}\}$. The fourth statement follows at once from $\mathscr{M}(m) = \{m\}$ for all $m \in \mathbf{M}$.

The elements of $\overline{\mathfrak{M}}$ are therefore the *closed* sets of maximal elements of $\langle\mathbf{C},\leq\rangle$. Since $1_{\mathbf{S}}$ and $\emptyset$ correspond in the dual order isomorphism, the partially ordered sets $\langle\mathbf{C},\leq\rangle$ and $\langle\mathfrak{M},\subseteq\rangle$ are dually order-isomorphic as well, with essentially the same dual order isomorphism $\mathscr{M}(\cdot)$. Other correspondences are b_v and $\mathbf{M}$.

The complete lattice $\langle\overline{\mathfrak{M}},\subseteq\rangle$ is called the *dual belief structure*, or simply the *dual*, of $\langle\mathbf{S},\mathbf{C},\leq\rangle$. Elements of $\overline{\mathfrak{M}}$ will also be called *spheres*. As $1_{\mathbf{S}}$ is the only inconsistent belief model in $\overline{\mathbf{C}}$, so $\emptyset$ is the only inconsistent sphere in $\overline{\mathfrak{M}}$, and it represents contradiction. $\mathbf{M}$ is called the *vacuous sphere*, and it corresponds to the least informative coherent belief model b_v. Singletons $\{m\}$ correspond to the maximally informative coherent belief models $m \in \mathbf{M}$.

We have seen that taking infima is very easy in the structure $\overline{\mathbf{C}}$: they coincide with infima in $\mathbf{S}$. But for taking suprema in $\overline{\mathbf{C}}$, we need to invoke the closure operator $\mathrm{Cl}_{\mathbf{S}}$. As an example, for two coherent belief models a and b, their supremum in $\overline{\mathbf{C}}$ is given by $\mathrm{Cl}_{\mathbf{S}}(a \smile b)$. But the corresponding operation is much easier in the dual structure: $\mathscr{M}(\mathrm{Cl}_{\mathbf{S}}(a \smile b)) = \mathscr{M}(a \smile b) = \mathscr{M}(a) \cap \mathscr{M}(b)$, or in other words, we just have to take intersections! To summarise, the most informative coherent belief model that is at most as informative as a and b is better described in the 'direct structure' $[a \frown b]$, than in the dual structure $[\mathrm{Cl}_{\mathbf{M}}(\mathscr{M}(a) \cup \mathscr{M}(b))]$; and the least informative coherent belief model that is at least as informative as a and b is has a more convenient representation $[\mathscr{M}(a) \cap \mathscr{M}(b)]$ in the dual than in the direct structure $[\mathrm{Cl}_{\mathbf{S}}(a \smile b)]$.

4. Examples of belief structures

Most of the mathematical models for representing beliefs (or uncertainty) in the literature that I am aware of constitute belief structures, apart from the ones that enforce precision or completeness, such as the Bayesian probability model. Many important ones even give rise to strong belief structures. In this section, I briefly discuss a number of examples, without aiming at completeness. They provide the main justification for the introduction and study of the abstract notions in the previous sections.

Classical propositional logic

We have seen above that the structure $\langle \mathbf{S}, \mathbf{C}, \leq \rangle = \langle \wp(\mathbf{L}), \mathbf{C}_{\mathbf{L}}, \subseteq \rangle$ that appears in the context of classical propositional logic is a *strong belief structure*. To summarise Section 2.1, the closure operator $\mathrm{Cl}_{\mathbf{S}}$ is here the logical, or deductive, closure operator; consistency is logical consistency; and coherence is logical consistency together with logical closure.

It is interesting to note that, in this special case, the elements of $\mathbf{M}$, *i.e.*, the maximal coherent belief models or sets of sentences, which constitute the building blocks for the dual structure $\overline{\mathfrak{M}}$, are sometimes called *(possible) worlds*. They also correspond to the possible valuations, and constitute the elements of the so-called dual *Boolean space* of the Lindenbaum algebra for $\mathbf{L}$ (which is a Boolean algebra). If K is a set of sentences, then $\mathscr{M}(K)$ can be interpreted as the set of those worlds where all the sentences in K are true.

It is well-known that there is a topology on $\mathbf{M}$ that turns $\mathbf{M}$ into a zero-dimensional, compact Hausdorff space (also called a *Stone space*); and that the Lindenbaum algebra for $\mathbf{L}$ is order-isomorphic to the Boolean algebra of the clopen subsets of the Stone space $\mathbf{M}$ (see for instance [6, 28] for more details). It is not too difficult to show that the elements of the set $\overline{\mathfrak{M}}$ are precisely the closed subsets of $\mathbf{M}$ in this topology, and that $\mathrm{Cl}_{\mathbf{M}}$ is therefore the *topological*[4] closure operator associated with it! In other words, the union of two closed sets of possible worlds is closed. Needless to say, this special property of the closure operator $\mathrm{Cl}_{\mathbf{M}}$ for classical propositional logic makes the reasoning mechanism behind it (and its direct counterpart $\mathrm{Cl}_{\mathbf{S}}$) rather special. We shall come back to this observation a number of times in what follows.

IMPRECISE PROBABILITY MODELS

In his important work on imprecise probabilities [32], Walley discusses a number of essentially equivalent imprecise probability models: lower previsions, upper previsions, sets of almost-desirable gambles, sets of strictly desirable gambles, almost-preference and strict preference relations. Lower and upper probabilities are special cases of these, and are less expressive. I shall give a very brief description of two of these models: sets of almost-desirable gambles and lower previsions, but related considerations can be made for the other models. For a much more detailed discussion, I refer to [32].

Sets of almost-desirable gambles

Consider a non-empty set Ω. We could interpret Ω as a set of possible states of the world, or of possible outcomes of some experiment. A bounded real-valued map X on Ω is called a *gamble*, and it represents a (possibly negative!) uncertain reward: if the actual state of the world turns out to be ω, then the award will be $X(\omega)$, expressed in units of some linear utility. So the reward is uncertain because the actual state of the world is. The set of all gambles on Ω will be denoted by $\mathscr{L}(\Omega)$. It is a linear space under the point-wise addition of gambles and the point-wise scalar multiplication of gambles with real numbers.

A subject can model his beliefs about the state of the world by specifying a set $\mathscr{D}$ of so-called *almost-desirable* gambles, *i.e.*, gambles X such that he accepts the uncertain reward $X+\varepsilon$ for any $\varepsilon>0$. Any subset of $\mathscr{L}(\Omega)$ is therefore a potential belief model, and consequently it makes good sense to let the power set $\wp(\mathscr{L}(\Omega))$ of $\mathscr{L}(\Omega)$ be the collection $\mathbf{S}$ of all belief models in this context. This set can be partially ordered by set inclusion, and the structure $\langle\wp(\mathscr{L}(\Omega)),\subseteq\rangle$ is or course a complete lattice, where intersection plays the role of infimum, and union that of supremum. Since specifying more

[4] A topological closure operator is a Moore closure with the additional property that the closure of any union is the union of the closures.

almost-desirable gambles leads to a more informative model, set inclusion $\subseteq$ indeed can be given the interpretation 'is less informative than'.

Since rewards are expressed in units of a linear utility, we see that any non-negative linear combination of a finite number of acceptable gambles is acceptable too. A set $\mathscr{D}$ of almost-desirable gambles is said to *avoid sure loss* if for any finite number of (not necessarily different) gambles $X_1, \ldots, X_n$ in $\mathscr{D}$ and non-negative real numbers $\lambda_1, \ldots, \lambda_n$, we have that

$$\sup_{\omega \in \Omega} \left[\sum_{k=1}^{n} \lambda_k X_k(\omega) \right] \geq 0.$$

Indeed, if this condition does not hold, then there are $X_1, \ldots, X_n$ in $\mathscr{D}$, non-negative real $\lambda_1, \ldots, \lambda_n$, and $\varepsilon > 0$ such that for all $\omega \in \Omega$, $\sum_{k=1}^{n} \lambda_k [X_k(\omega) + \varepsilon] < -\varepsilon$. In other words, the acceptable gamble $\sum_{k=1}^{n} \lambda_k [X_k + \varepsilon]$ (a non-negative linear combination of acceptable gambles) leads to a sure loss of a least ε, whatever the state of the world.

Taking non-negative linear combinations is the reasoning mechanism that allows us to infer new acceptable gambles from old ones. Moreover, non-negative gambles are clearly almost-desirable. It is not difficult to see that a set of gambles $\mathscr{D}$ is closed under this type of inference if and only if it is a convex cone of the linear space $\mathscr{L}(\Omega)$ that contains all non-negative gambles and is furthermore closed in the supremum-norm topology on $\mathscr{L}(\Omega)$. Any such set of gambles $\mathscr{D}$ is called *coherent* if it furthermore avoids sure loss (*i.e.*, does not contain any uniformly negative gambles). These coherent sets of gambles constitute a belief structure: the intersection of a non-empty class of coherent sets is again coherent. The associated *consistent* sets of gambles are precisely the ones that avoid sure loss. Walley defines the natural extension of a consistent set of gambles $\mathscr{D}$ to be smallest coherent set that includes it; natural extension therefore coincides with the closure operation of this belief structure, and it is very closely related to convex closure (under the supremum-norm topology).

The maximal coherent sets of almost-desirable gambles are the closed semispaces of $\mathscr{L}(\Omega)$ that contain all non-negative gambles. Using a separation theorem equivalent to the Ultrafilter Theorem [28, 32], it can be shown that any coherent set of almost-desirable gambles is the intersection of the maximal coherent sets it is included in. This tells us that coherent sets of almost-desirable gambles constitute a strong belief structure.

Observe that there is a one-to-one relationship between the closed positive semispaces of $\mathscr{L}(\Omega)$ and the so-called *linear previsions* P on $\mathscr{L}(\Omega)$, *i.e.*, the linear maps from $\mathscr{L}(\Omega)$ to the set $\mathbb{R}$ of real numbers that are furthermore positive (if $X \geq 0$ then $P(X) \geq 0$) and have unit norm ($P(1) = 1$): any such semispace is the set of all gambles X for which the corresponding linear prevision $P(X)$ is non-negative. Linear previsions are the precise probability models: they can be interpreted as coherent previsions, or fair prices, in the sense of de Finetti [12]. Thus maximal coherent sets of almost-desirable gambles can be identified with linear previsions. In the dual belief structure, closed sets of maximal elements (spheres) can then be identified with (weak*-)closed convex sets of linear previsions: closure in the dual structure is not topological, essentially because the union of two convex sets need not be convex.

Lower previsions
A lower prevision $\underline{P}$ on $\mathscr{K}$ (a belief model) is a map from a set of gambles $\mathscr{K}$ to the extended real interval $[-\infty,+\infty]$. For any gamble X, its lower prevision $\underline{P}(X)$ is interpreted as a subject's supremum acceptable buying price for X, *i.e.*, the greatest α such that the subject accepts to buy the uncertain reward X for any price $p<\alpha$. Let us denote by $\mathbb{F}(\mathscr{K})$ the set of all lower previsions with domain $\mathscr{K}$.

Lower previsions with the same domain $\mathscr{K}$ can be partially ordered point-wise, and $\langle\mathbb{F}(\mathscr{K}),\leq\rangle$ is a complete lattice, where $\leq$ is the point-wise order of lower previsions. The ordering indeed has the interpretation 'is less informative than' or 'is less precise than': if a subject has a smaller lower prevision, this means that he is committed to pay less for the gambles in $\mathscr{K}$, and therefore displays a more conservative behaviour.

The coherent belief models are the lower previsions that are coherent in Walley's sense [32, Section 2.5]: for any finite number of (not necessarily different) gambles X_o, X_1, ..., X_n in $\mathscr{K}$ and non-negative real numbers λ_o, $\lambda_1,\dots,\lambda_n$, we require that

$$\sup_{\omega\in\Omega}\left[\sum_{k=1}^{n}\lambda_k[X_k(\omega)-\underline{P}(X_k)]-\lambda_o[X_o(\omega)-\underline{P}(X_o)]\right]\geq 0.$$

The point-wise infimum of a non-empty collection of coherent lower previsions is indeed coherent. If we denote the set of coherent lower previsions with domain $\mathscr{K}$ by $\underline{\mathbb{P}}(\mathscr{K})$, then we see that $\langle\mathbb{F}(\mathscr{K}),\underline{\mathbb{P}}(\mathscr{K}),\leq\rangle$ is a belief structure. The consistent models turn out to be the lower previsions that avoid sure loss [32, Section 2.4]: for any finite number of (not necessarily different) gambles X_1, ..., X_n in $\mathscr{K}$ and non-negative real numbers λ_1, ..., λ_n, it is required that

$$\sup_{\omega\in\Omega}\left[\sum_{k=1}^{n}\lambda_k[X_k(\omega)-\underline{P}(X_k)]\right]\geq 0.$$

According to Walley's definition [32, Section 3.1], the natural extension $\underline{E}$ of a lower prevision $\underline{P}$ that avoids sure loss is the smallest coherent lower prevision that point-wise dominates $\underline{P}$. It is clear that natural extension coincides with the closure operator in the present belief structure.

The maximal coherent belief models are the linear previsions on $\mathscr{K}$ (restrictions to $\mathscr{K}$ of the linear previsions on $\mathscr{L}(\Omega)$ defined above) [32, Section 2.8], which are the precise probability models. A coherent lower prevision is the point-wise infimum of its set of dominating linear previsions, so lower previsions constitute a strong belief structure. The spheres (closed sets of maximal elements in the dual belief structure) are the (weak*-)closed convex sets of linear previsions, and closure in the dual structure is therefore convex closure, and is not topological.

THE RELATION BETWEEN THESE MODELS

There is a one-to-one relationship between coherent sets of almost-desirable gambles and coherent lower previsions defined on $\mathscr{L}(\Omega)$. It follows naturally from the behavioural interpretations of these belief models.

Consider a subject's coherent set of almost-desirable gambles $\mathscr{D}$. Then the subject accepts to buy the gamble X for any price $p < \alpha$ if and only if the gamble $X - \alpha$ belongs to $\mathscr{D}$. It follows that his lower prevision $\underline{P}(X)$ for X is given by:

$$\underline{P}(X) = \sup\{\alpha \colon X - \alpha \in \mathscr{D}\}.$$

Conversely, consider a subject's coherent lower prevision $\underline{P}$ on $\mathscr{L}(\Omega)$. Then a gamble X is almost-desirable to him if he is willing to accept the gamble $X + \varepsilon$, or in other words to buy X for the price $-\varepsilon$, for all $\varepsilon > 0$. This is equivalent to $\underline{P}(X) \geq 0$, Consequently, his set of almost-desirable gambles is uniquely determined by

$$\mathscr{D} = \{X \in \mathscr{L}(\Omega) \colon \underline{P}(X) \geq 0\}.$$

This relationship is an example of the coherence isomorphisms that will be introduced in the next section.

SEVERAL OTHER MODELS, BRIEFLY

The *confidence relations* that I introduced and studied in [7] constitute a strong belief structure. So do Giles' so-called *possibility functions* [16]. With hindsight, these are precisely the coherent upper probabilities on a field of sets. *Ordinal possibility measures* [8, 14] lead to a belief structure that is not strong: in this structure the belief models are maps from the power set $\wp(\Omega)$ of some non-empty set Ω to a complete lattice (or chain) $\langle K, \preceq \rangle$. They can be ordered point-wise (we consider the dual, or reversed, ordering), and the coherent belief models are the normal K-valued possibility measures. These are closed under infima (*i.e.*, under point-wise suprema), and the corresponding closure operator can be related to possibilistic extension [4]. The same holds for Spohn's *ordinal conditional functions* [31], which are very closely related to, but less expressive than, ordinal possibility measures.

Among the hierarchical uncertainty models, the *price functions* introduced by Walley and myself [10] constitute a belief structure that is not strong. On the other hand, the more general *lower desirability functions* [9] do lead to a strong belief structure, for which the maximal coherent belief models are essentially the Bayesian second-order probabilities.

Aumann's *preference-or-indifference relations* defined on a mixture space [1, 2] lead to a belief structure, and so do the *preference relations* on horse lotteries studied by Seidenfeld *et al.* [29, 30]. In both cases, the belief structures seem not to be strong, but the authors do pay attention to representation of their belief models as intersections (*i.e.*, infima) of maximal belief models, and are able to derive interesting but partial representation results. In any case, and although the authors would probably object to this (see the discussion in [29, Section VI]), it is possible to get to a strong belief structure by looking at almost-preference rather than real preference,[5] in the spirit of [32, Sections 3.7 and 3.8]: one item is almost-preferred to a second item if it is the limit of a sequence of items that are really preferred to the second item. This amounts to replacing the Archimedean axioms in [1, 29] by a closedness axiom, and keeping all the other axioms.[6]

5. Belief substructures

In order to investigate the relations between the many belief structures in the literature, it is useful to be able to express that one belief structure is more general than another, or extends it in some way. This can be done with the notions of belief substructures, belief embeddings and belief isomorphisms. I want to stress that only a few basic notions are introduced here: the ideas hinted at below could (and probably should) be worked out and studied in much more detail.

In order to define belief substructures, we must keep in mind what is important about belief structures: the ordering between belief models and the set of coherent belief models . So a substructure of a belief structure will be a subset of the set belief models that 'inherits' its order and its coherent belief models. For strong belief substructures, also the maximal coherent belief models will have to be 'inherited'. More explicitly, we get the following.

Definition 4. Let $\langle \mathbf{S}_1, \mathbf{C}_1, \leq_1 \rangle$ and $\langle \mathbf{S}_2, \mathbf{C}_2, \leq_2 \rangle$ be belief structures. Then $\langle \mathbf{S}_1, \mathbf{C}_1, \leq_1 \rangle$ is called a *belief substructure* of $\langle \mathbf{S}_2, \mathbf{C}_2, \leq_2 \rangle$ if

1. $\langle \mathbf{S}_1, \leq_1 \rangle$ is a complete sublattice of $\langle \mathbf{S}_2, \leq_2 \rangle$, and consequently $\mathbf{S}_1 \subseteq \mathbf{S}_2$;

2. $\mathbf{C}_1 = \mathbf{C}_2 \cap \mathbf{S}_1$.

[5] This means that we look at a different notion of preference.

[6] See also [25, 26] for a more detailed study of such almost-preference relations over horse lotteries.

If $\langle \mathbf{S}_1, \mathbf{C}_1, \leq_1 \rangle$ and $\langle \mathbf{S}_2, \mathbf{C}_2, \leq_2 \rangle$ are strong, then $\langle \mathbf{S}_1, \mathbf{C}_1, \leq_1 \rangle$ is called a *strong belief substructure* of $\langle \mathbf{S}_2, \mathbf{C}_2, \leq_2 \rangle$ if in addition:

3. $\mathbf{M}_1 = \mathbf{M}_2 \cap \mathbf{S}_1$.

Belief embeddings, then, allow us to generalise the notion of a belief substructure, by allowing 'a copy' of the first belief structure to be a belief substructure of the second. This is done as follows.

Definition 5. Let $\langle \mathbf{S}_1, \mathbf{C}_1, \leq_1 \rangle$ and $\langle \mathbf{S}_2, \mathbf{C}_2, \leq_2 \rangle$ be belief structures. A map $\phi \colon \mathbf{S}_1 \to \mathbf{S}_2$ is called a *belief embedding* of $\langle \mathbf{S}_1, \mathbf{C}_1, \leq_1 \rangle$ in $\langle \mathbf{S}_2, \mathbf{C}_2, \leq_2 \rangle$ if

1. ϕ is an order embedding of $\langle \mathbf{S}_1, \leq_1 \rangle$ in $\langle \mathbf{S}_2, \leq_2 \rangle$, *i.e.*,
$$(\forall (s,t) \in \mathbf{S}_1^2)(s \leq_1 t \Leftrightarrow \phi(s) \leq_2 \phi(t)).$$
2. $\phi(\mathbf{C}_1) = \mathbf{C}_2 \cap \phi(\mathbf{S}_1)$.

If the belief structures $\langle \mathbf{S}_1, \mathbf{C}_1, \leq_1 \rangle$ and $\langle \mathbf{S}_2, \mathbf{C}_2, \leq_2 \rangle$ are strong, then the map $\phi \colon \mathbf{S}_1 \to \mathbf{S}_2$ is a *strong belief embedding* of $\langle \mathbf{S}_1, \mathbf{C}_1, \leq_1 \rangle$ in $\langle \mathbf{S}_2, \mathbf{C}_2, \leq_2 \rangle$ if in addition:

3. $\phi(\mathbf{M}_1) = \mathbf{M}_2 \cap \phi(\mathbf{S}_1)$.

If moreover the map ϕ is not only injective, but also surjective, and therefore an order isomorphism between the complete lattices $\langle \mathbf{S}_1, \leq_1 \rangle$ and $\langle \mathbf{S}_2, \leq_2 \rangle$, then ϕ is called a *(strong) belief isomorphism*, and the (strong) belief structures $\langle \mathbf{S}_1, \mathbf{C}_1, \leq_1 \rangle$ and $\langle \mathbf{S}_2, \mathbf{C}_2, \leq_2 \rangle$ are called *belief isomorphic*. In that case, $\phi(\mathbf{S}_1) = \mathbf{S}_2$, $\phi(\mathbf{C}_1) = \mathbf{C}_2$ and $\phi(\mathbf{M}_1) = \mathbf{M}_2$.

If ϕ is a (strong) belief embedding of $\langle \mathbf{S}_1, \mathbf{C}_1, \leq_1 \rangle$ in $\langle \mathbf{S}_2, \mathbf{C}_2, \leq_2 \rangle$, then the structures $\langle \mathbf{S}_1, \mathbf{C}_1, \leq_1 \rangle$ and $\langle \phi(\mathbf{S}_1), \phi(\mathbf{C}_1), \leq_2 \rangle$ are (strongly) belief isomorphic, and therefore copies of one another. And the copy $\langle \phi(\mathbf{S}_1), \phi(\mathbf{C}_1), \leq_2 \rangle$ is a (strong) belief substructure of $\langle \mathbf{S}_2, \mathbf{C}_2, \leq_2 \rangle$.

The most important part of a belief structure is its set $\mathbf{C}$ of coherent belief models. Thus it is possible that two belief structures have essentially the same set of coherent belief models, although they differ as far as their incoherent models are concerned. Since most types of reasoning only involve the coherent models—because inference amounts to taking closures—we need some way to recognise that these two structures are identical in what matters most.

Definition 6. Two belief structures $\langle \mathbf{S}_1, \mathbf{C}_1, \leq_1 \rangle$ and $\langle \mathbf{S}_2, \mathbf{C}_2, \leq_2 \rangle$ are called *coherence isomorphic* if the complete lattices $\langle \overline{\mathbf{C}}_1, \leq_1 \rangle$ and $\langle \overline{\mathbf{C}}_2, \leq_2 \rangle$ are order isomorphic.

This means that in the two belief structures, the coherent models, the closure operators, the sets of maximal elements and therefore also their dual structures (if they exist) are essentially the same.

As mentioned in Section 4, most of the imprecise probability models introduced by Walley [32] lead to strong belief structures: e.g., lower previsions, upper previsions, sets of almost-desirable gambles and almost-preference relations. These structures are all coherence isomorphic, and they have essentially the same dual structure: (weak*-)closed convex sets of linear previsions. Among these structures, the ones built on lower previsions and upper previsions are also (strongly) belief isomorphic.

The following important example explains how the strong belief structure built on classical propositional logic can be seen as a substructure of the one built on the above-mentioned imprecise probability models.

CLASSICAL LOGIC AND IMPRECISE PROBABILITIES

We intend to show that the strong belief structure based on classical propositional logic discussed in Section 4 can be embedded into the strong belief structure based on lower previsions described in that same section. We proceed in two steps.

First of all, lower previsions are expressed in terms of some set Ω of possible states of the world (or possible worlds). We therefore proceed in the usual way in relating the propositional logic system **L** to such a set of possible worlds, and in relating sentences in **L** to certain subsets of Ω, called events. Indeed, by the Stone Representation Theorem applied to the Lindenbaum algebra (a Boolean algebra) of the system **L** [6], there is some set Ω (its dual Boolean space, or the set of two-valued Boolean homomorphisms, or its set of Boolean ultrafilters, also called its set of possible worlds) and a field $\mathscr{A}$ of subsets of Ω, such that there is a one-to-one correspondence between sentences in **L**—after identifying syntactically equivalent sentences—and elements of $\mathscr{A}$. An element A of $\mathscr{A}$ is called an *event*, and it corresponds to an element of the Lindenbaum algebra of **L**, or in other words to set of logically equivalent sentences in **L**. Moreover, it is easily checked that

(i) sets of sentences (*i.e.*, belief models) correspond to sets of events, or in other words to *subsets* of $\mathscr{A}$;

(ii) consistent and logically closed sets of sentences (*i.e.*, coherent belief models) correspond to *filters* of events, *i.e.*, subsets of $\mathscr{A}$ that are increasing with respect to set inclusion, and closed under intersection; and

(iii) maximal consistent and logically closed sets of sentences (*i.e.*, maximal coherent belief models) correspond to *ultrafilters* of events, *i.e.*, filters of events that are maximal elements with respect to set inclusion.

If we denote the set of subsets of $\mathscr{A}$ by $\wp(\mathscr{A})$, its set of filters by $\mathscr{F}(\mathscr{A})$ and its set of ultrafilters by $\mathscr{U}(\mathscr{A})$, then $\langle \wp(\mathscr{A}), \mathscr{F}(\mathscr{A}), \subseteq \rangle$ is a strong belief structure, with set of spheres $\mathscr{U}(\mathscr{A})$, which we have argued is belief isomorphic—after identifying logically equivalent sentences—to the strong belief structure based on the classical propositional logic system **L**. The set $\mathscr{U}(\mathscr{A})$ of maximal elements corresponds in this belief isomorphism to the set of complete theories of **L**, *i.e.*, the set of maximal, deductively closed sets of sentences in **L**.

As a second step, we consider the set of gambles $\mathscr{K}_{\mathscr{A}} = \{I_A\colon A \in \mathscr{A}\}$ on Ω, where the gamble I_A is the indicator function of the subset A of Ω, assuming the value 1 on A and 0 elsewhere. We define the map ϕ from $\wp(\mathscr{A})$ to the set $\mathbb{F}(\mathscr{K}_{\mathscr{A}})$ of lower previsions with domain $\mathscr{K}_{\mathscr{A}}$: for any $\mathscr{B} \subseteq \mathscr{A}$, $\phi(\mathscr{B})$ is the lower prevision on $\mathscr{K}_{\mathscr{A}}$, defined by

$$\phi(\mathscr{B})(I_A) = I_{\mathscr{B}}(A) = \begin{cases} 1 & \text{if } A \in \mathscr{B} \\ 0 & \text{if } A \notin \mathscr{B}. \end{cases}$$

Observe that $\phi(\mathscr{B})$ is essentially defined on the field $\mathscr{A}$ of events, and we can therefore call it a *lower probability*.[7] To interpret ϕ, consider a subset $\mathscr{B}$ of $\mathscr{A}$—which is equivalent to a specific set of sentences, or in other words, to a specific belief model. Suppose you believe these sentences to be true, or in the language of events, that you believe that all events in $\mathscr{B}$ occur. Now consider an event $A \in \mathscr{A}$. If $A \in \mathscr{B}$, then you believe that A occurs, and so you are willing to bet on the occurrence of A at all odds, or equivalently, to buy the uncertain reward I_A for any price up to one. This means that your lower prevision for I_A, or equivalently, your lower probability for A, is equal to one. If $A \notin \mathscr{B}$, then the only thing we can say is that it is not the case that you believe that A occurs (you might believe that A doesn't occur, or you might be ignorant about the occurrence of A). In that case you will not be prepared to bet on the occurrence of A at any non-trivial odds, so your lower prevision for I_A, or equivalently, your lower probability for A, is equal to zero. In other words, $\phi(\mathscr{B})$ is the lower probability that corresponds to your believing that all events in $\mathscr{B}$ occur, or that all the corresponding sentences are true, *and nothing more*.

It is easy to verify that ϕ is an order embedding of the complete lattice $\langle\wp(\mathscr{A}), \subseteq\rangle$ into the complete lattice $\langle\mathbb{F}(\mathscr{K}_{\mathscr{A}}), \leq\rangle$. $\phi(\wp(\mathscr{A}))$ is the set of the $0-1$-valued lower previsions on $\mathscr{K}_{\mathscr{A}}$ (or $0-1$-valued lower probabilities on $\mathscr{A}$). Moreover, we can reformulate Walley's results in [32, Section 2.9.8] to conclude that

(i) the lower prevision (probability) $\phi(\mathscr{B})$ avoids sure loss if and only if the set of propositions behind $\mathscr{B}$ is logically consistent;

(ii) the lower prevision (probability) $\phi(\mathscr{B})$ is coherent if and only if the set of propositions behind $\mathscr{B}$ is logically consistent and deductively closed;

[7] A (lower) probability is the restriction of a (lower) prevision to (indicators) of events.

(iii) the lower prevision (probability) $\phi(\mathscr{B})$ is a coherent prevision (finitely additive probability[8]) if and only if the propositions behind $\mathscr{B}$ make up a maximal, logically consistent and deductively closed set; and

(iv) taking the natural extension of the lower probability $\phi(\mathscr{B})$ that avoids sure loss leads to the lower probability $\phi(\mathscr{B}')$ where $\mathscr{B}'$ corresponds to the logical closure of the sentences corresponding to $\mathscr{B}$.

In other words, ϕ maps the filters of $\mathscr{A}$ onto the *coherent* $0-1$-valued lower probabilities on $\mathscr{A}$: $\phi(\mathscr{F}(\mathscr{A})) = \phi(\wp(\mathscr{A})) \cap \underline{\mathbb{P}}(\mathscr{K}_{\mathscr{A}})$; and in addition, ϕ maps the ultrafilters of $\mathscr{A}$ onto the $0-1$-valued finitely additive probabilities on $\mathscr{A}$. This tells us that the map ϕ is a strong belief embedding of $\langle \wp(\mathscr{A}), \mathscr{F}(\mathscr{A}), \subseteq \rangle$ into $\langle \mathbb{F}(\mathscr{K}_{\mathscr{A}}), \underline{\mathbb{P}}(\mathscr{K}_{\mathscr{A}}), \leq \rangle$. In this embedding, natural extension generalises the logical closure operator.

If we denote by $\mathbb{F}_{01}(\mathscr{A})$ the set of $0-1$-valued lower probabilities on $\mathscr{A}$, by $\underline{\mathbb{P}}_{01}(\mathscr{A})$ the set of coherent $0-1$-valued lower probabilities on $\mathscr{A}$, and by $\mathbb{P}_{01}(\mathscr{A})$ the set of $0-1$-valued finitely additive probabilities on $\mathscr{A}$, then the strong belief structure $\langle \mathbb{F}_{01}(\mathscr{A}), \underline{\mathbb{P}}_{01}(\mathscr{A}), \leq \rangle$ is actually belief isomorpic to the strong belief structure $\langle \wp(\mathscr{A}), \mathscr{F}(\mathscr{A}), \subseteq \rangle$, and therefore to the strong belief structure based on the classical propositional logic system **L**. The set of $0-1$-valued precise probabilities corresponds in this belief isomorphism to $\mathscr{U}(\mathscr{A})$, and therefore to the set of complete theories of **L**, *i.e.*, the set of maximal, deductively closed sets of sentences in **L**. These correspondences are summarised in Figure 1.

$$\begin{array}{ccccc} \langle \wp(\mathbf{L}), \mathbf{C}_{\mathbf{L}}, \subseteq \rangle & = & \langle \wp(\mathscr{A}), \mathscr{F}(\mathscr{A}), \subseteq \rangle & = & \langle \mathbb{F}_{01}(\mathscr{A}), \underline{\mathbb{P}}_{01}(\mathscr{A}), \leq \rangle \\ & & & & \downarrow \\ & & & & \langle \mathbb{F}(\mathscr{K}_{\mathscr{A}}), \underline{\mathbb{P}}(\mathscr{K}_{\mathscr{A}}), \leq \rangle \end{array}$$

Figure 1. Correspondences between the belief structures based on classical propositional logic and those based on $0-1$-valued lower probabilities. Strong belief embeddings are depicted using arrows, belief isomorphisms using double lines.

We conclude from these two steps that the strong belief structure built on classical propositional logic can be embedded in the one built on lower previsions. In this sense, the theory of coherent lower probabilities is a generalisation of classical propositional logic. In this embedding, precise previsions (or probabilities) play the role of maximal elements, and correspond to the maximal consistent logically closed sets of sentences. In this light, it seems strange that a number of Bayesians continue to claim that probability

[8] Finitely additive probabilities are the restrictions of linear, or coherent, previsions to events.

measures are the *only* reasonable extension of classical logic able to deal with partial beliefs: how many logicians would claim that the only rational logically closed sets of sentences are the maximal ones—or that the only rational theories are complete? This result furthermore tells us that, in a very definite sense, precise probability theory is not powerful enough to generalise all of classical propositional logic, but that imprecise probability theory (*i.e.*, the theory of coherent lower previsions) is!

6. Belief expansion

I now want to show that it can be useful to look at existing types of belief models as special cases of the abstract order-theoretic structures introduced above. This is because quite often the exact underlying details of how belief models are constructed, is not really of crucial importance; what matters is the reasoning, or inference, method and that is captured completely in the closure operator $\mathrm{Cl}_{\mathbf{S}}$ and its dual counterpart $\mathrm{Cl}_{\mathbf{M}}$ (if it exists). We shall see below that in a number of interesting cases, only the order-theoretic properties of these closure operators are relevant, and not the additional properties which they may derive from the underlying details of the belief models.

In the sections that follow, I generalise part of the work done by Gärdenfors [15] on belief expansion and revision of epistemic states in the context of classical propositional logic, where his so-called *epistemic states* are logically closed sets of sentences. In principle, nothing prevents us from considering as an epistemic state a more general type of belief model, such as the imprecise probability models or the preference orderings discussed in Section 4. Indeed, these models are also intended to represent the beliefs (and utilities) of some subject. But how do we then define belief expansion and revision, and how can Gärdenfors's coherentist axioms for belief change be generalised? Below I sketch how this could be done, and thereby generalise the work done by Moral and Wilson [24] on belief revision when the epistemic states are closed convex sets of probabilities, and to some extent the work of Benferhat *et al.* [3] and Dubois and Prade [13] for possibilistic belief revision.

I want to stress that it is not my aim in this paper to present a fully worked out theory of belief change for belief models. The discussion below is only intended to illustrate the usefulness of my belief structures. I shall therefore restrict myself to pointing out the more striking results for belief revision and expansion, and give only little further motivation. I will say only very little about the notion of contraction, because that seems more involved and outside the scope of this paper. I want to stress that definitions and results below (but certainly also the order-theoretic simplicity of their proofs!) should be compared with the discussion in Gärdenfors's book [15] in order to be fully understood. I also hasten to add that my concentrating on Gärdenfors's work does not imply that I think he had the final word in the matter of belief change, nor that I believe his approach and his axiom systems to be the only, or for that matter the most reasonable ones. Indeed, Gärdenfors's view of belief states as sets of propositional sentences has been criticised and expanded several times (see for instance [5, 19, 33]). But Gärdenfors's approach and axioms did seem simple enough to provide an elegant example of the power of the order-theoretic machinery introduced above.

Let me start with belief expansion. Assume that we have a coherent belief model $b \in \mathbf{C}$, and that new information is obtained, which can be represented by a (not necessarily coherent) belief model $\gamma \in \mathbf{S}$. This new information takes b to a new coherent belief model b'. We represent the action of new information $\gamma \in \mathbf{S}$ on the coherent belief model b by an operator $E(b;\cdot)\colon \mathbf{S} \to \mathbf{S}$, called *(belief) expansion operator*. In the spirit of the work of Gärdenfors, we may require that such an operator should satisfy the following postulates: for b and c in $\mathbf{C}$, and for all $\gamma \in \mathbf{S}$,

$E1$. $E(b;\gamma) \in \overline{\mathbf{C}}$;

$E2$. $\gamma \leq E(b;\gamma)$;

$E3$. $b \leq E(b;\gamma)$;

$E4$. if $\gamma \leq b$ then $E(b;\gamma) = b$;

$E5$. if $b \leq c$ then $E(b;\gamma) \leq E(c;\gamma)$;

$E6$. $E(b;\cdot)$ is the point-wise smallest (least informative) of all the operators satisfying $E1$–$E5$.

$E1$–$E6$ correspond one by one to Gärdenfors's expansion postulates $(\mathrm{K}^{+}1)$–$(\mathrm{K}^{+}6)$, in that order. The correspondence is obvious if we recall that expansion by a proposition has been generalised to expansion by a belief model.

Theorem 6. *Let $\langle \mathbf{S}, \mathbf{C}, \leq \rangle$ be a belief structure, and consider a coherent belief model $b \in \mathbf{C}$. Then the postulates $E1$–$E6$ single out a unique belief expansion operator $E(b;\cdot)$, given by:*

$$E(b;\gamma) = \mathrm{Cl}_{\mathbf{S}}(b \smile \gamma), \quad \gamma \in \mathbf{S}.$$

Proof. Note that $\mathrm{Cl}_{\mathbf{S}}(b \smile \cdot)$ obviously satisfies $E1$–$E5$. Moreover, for any $\gamma \in \mathbf{S}$ it follows from $E2$ and $E3$ that $b \smile \gamma \leq E(b;\gamma)$ and from $E1$ and Proposition 1 that $\mathrm{Cl}_{\mathbf{S}}(b \smile \gamma) \leq E(b;\gamma)$. From $E6$ we then deduce that $E(b;\gamma) = \mathrm{Cl}_{\mathbf{S}}(b \smile \gamma)$.

It is interesting to note that if b and γ are consistent, then $\mathrm{Cl}_{\mathbf{S}}(b \smile \gamma)$ is the supremum of b and $\mathrm{Cl}_{\mathbf{S}}(\gamma)$ in the complete join-semilattice $\langle \mathbf{C}, \leq \rangle$: it is the smallest (least informative) coherent belief model that is at least as informative as b and $\mathrm{Cl}_{\mathbf{S}}(\gamma)$. In the dual structure (if it exists), expansion takes a very simple form: expanding the sphere $\mathscr{M}(b) \in \mathfrak{M}$ with the sphere $\mathscr{N} \in \overline{\mathfrak{M}}$ amounts to taking their intersection $\mathscr{M}(b) \cap \mathscr{N}$.

7. Belief revision

Let us now turn to belief revision, where a coherent belief model b is revised into a belief model b' under new information in the form of a belief model $\gamma \in \mathbf{S}$. We again represent the action of the new information $\gamma \in \mathbf{S}$ on the coherent belief model b by an operator $R(b;\cdot)\colon \mathbf{S} \to \mathbf{S}$, called *(belief) revision operator*. Inspired by Gärdenfors's work, we propose the following postulates for belief revision: for b in $\mathbf{C}$, and for all γ in $\mathbf{S}$,

*R*1. $R(b;\gamma) \in \overline{\mathbf{C}}$;

*R*2. $\gamma \leq R(b;\gamma)$;

*R*3. $R(b;\gamma) \leq E(b;\gamma)$;

*R*4. if b and γ are consistent then $E(b;\gamma) \leq R(b;\gamma)$;

*R*5. $R(b;\gamma)$ is inconsistent if and only if γ is inconsistent;

*R*6. $R(b;\gamma) = R(b;\mathrm{Cl}_{\mathbf{S}}(\gamma))$;

*R*7. $R(b;\gamma \smile \delta) \leq E(R(b;\gamma);\delta)$;

*R*8. if $R(b;\gamma)$ and δ are consistent then $E(R(b;\gamma);\delta) \leq R(b;\gamma \smile \delta)$.

*R*1–*R*8 again correspond one by one, and in that order, to Gärdenfors's revision postulates (K^*1)–(K^*8). Here too, the correspondence is straightforward if (i) we recall that revision by a proposition has been generalised to revision by a belief model, (ii) we invoke the notion of (in)consistency to capture the essence of the postulates (K^*4) and (K^*5) involving the negation of propositions, and (iii) we realise that the conjunction of two propositions corresponds in our language to the join of two belief models: a belief model generalises a set of propositions, and revision by a conjunction of two propositions means revision by both the propositions, *i.e.*, by their 'union'.

The more striking results can be derived if the belief structure $\langle \mathbf{S}, \mathbf{C}, \leq \rangle$ is strong. Let us reformulate these axioms into their dual versions. It should be noted that *R*1 and *R*6 are necessary for this to be possible, as we can only represent elements of $\overline{\mathbf{C}}$ by closed sets of maximal coherent belief models. So, whenever we work in the dual space, with a *dual revision operator* $\mathfrak{R}(\mathscr{M}(b);\cdot)\colon \overline{\mathfrak{M}} \to \overline{\mathfrak{M}}$, it is implicit that *R*1 and *R*6 hold. It is easily verified that the other postulates can be reformulated in the following way: for all $\mathscr{N}$ and $\mathscr{S}$ in $\overline{\mathfrak{M}}$,

$\mathfrak{R}2$. $\mathfrak{R}(\mathscr{M}(b);\mathscr{N}) \subseteq \mathscr{N}$;

$\mathfrak{R}3$. $\mathscr{M}(b) \cap \mathscr{N} \subseteq \mathfrak{R}(\mathscr{M}(b);\mathscr{N})$;

$\mathfrak{R}4$. if $\mathscr{M}(b) \cap \mathscr{N} \neq \emptyset$ then $\mathfrak{R}(\mathscr{M}(b);\mathscr{N}) \subseteq \mathscr{M}(b) \cap \mathscr{N}$;

$\mathfrak{R}5$. $\mathfrak{R}(\mathscr{M}(b);\mathscr{N}) = \emptyset$ if and only if $\mathscr{N} = \emptyset$;

$\mathfrak{R}7$. $\mathfrak{R}(\mathscr{M}(b);\mathscr{N}) \cap \mathscr{S} \subseteq \mathfrak{R}(\mathscr{M}(b);\mathscr{N} \cap \mathscr{S})$;

$\mathfrak{R}8$. if $\mathfrak{R}(\mathscr{M}(b);\mathscr{N}) \cap \mathscr{S} \neq \emptyset$ then $\mathfrak{R}(\mathscr{M}(b);\mathscr{N} \cap \mathscr{S}) \subseteq \mathfrak{R}(\mathscr{M}(b);\mathscr{N}) \cap \mathscr{S}$.

I now propose a very particular type of dual revision operator, which will turn out to be sufficiently general. The central idea behind it is that for every $b \in \mathbf{C}$ (or every $\mathscr{M}(b)$) there is a *selection function* $\mathfrak{S}_b \colon \overline{\mathfrak{M}} \to \overline{\mathfrak{M}}$ that selects for any $\mathscr{N} \in \overline{\mathfrak{M}}$ a subset $\mathfrak{S}_b(\mathscr{N})$ of $\mathscr{N}$ under the following conditions:

$\mathfrak{S}1$. if $\mathscr{M}(b) \cap \mathscr{N} \neq \emptyset$ then $\mathfrak{S}_b(\mathscr{N}) = \mathscr{M}(b) \cap \mathscr{N}$;

$\mathfrak{S}2$. if $\mathscr{M}(b) \cap \mathscr{N} = \emptyset$ and $\mathscr{N} \neq \emptyset$ then $\mathfrak{S}_b(\mathscr{N})$ is *some* non-empty closed subset of $\mathscr{N}$; and

$\mathfrak{S}3$. $\mathfrak{S}_b(\emptyset) = \emptyset$.

A dual revision operator $\mathfrak{R}(\mathscr{M}(b);\cdot)$ can now be defined as follows: for any $\mathscr{N}$ in $\overline{\mathfrak{M}}$,

$$\mathfrak{R}(\mathscr{M}(b);\mathscr{N}) = \mathfrak{S}_b(\mathscr{N}). \tag{2}$$

For the corresponding revision operator $R(b;\cdot)$ we then have:

$$R(b;\gamma) = \inf \mathfrak{S}_b(\mathscr{M}(\gamma)). \tag{3}$$

There is the following general representation theorem. In the dual structure, its proof is a matter of straightforward verification, and it is therefore omitted.

Theorem 7. *Let $\langle \mathbf{S}, \mathbf{C}, \leq \rangle$ be a strong belief structure and let $b \in \mathbf{C}$ be a coherent belief model. A dual revision operator $\mathfrak{R}(\mathscr{M}(b);\cdot)$ satisfies $\mathfrak{R}2$–$\mathfrak{R}5$ if and only if there is a selection function $\mathfrak{S}_b$ satisfying $\mathfrak{S}1$–$\mathfrak{S}3$ such that $\mathfrak{R}(\mathscr{M}(b);\cdot) = \mathfrak{S}_b(\cdot)$. Equivalently, a revision operator $R(b;\cdot)$ satisfies $R1$–$R6$ if and only if there is a selection function $\mathfrak{S}_b$ satisfying $\mathfrak{S}1$–$\mathfrak{S}3$ such that $R(b;\cdot) = \inf \mathfrak{S}_b(\mathscr{M}(\cdot))$. Moreover, $\mathfrak{R}(\mathscr{M}(b);\cdot)$ also satisfies $\mathfrak{R}7$–$\mathfrak{R}8$, and $R(b;\cdot)$ also satisfies $R7$–$R8$, if and only if the selection function $\mathfrak{S}_b$ satisfies, for all $\mathscr{N}$ and $\mathscr{S}$ in $\overline{\mathfrak{M}}$:*

$\mathfrak{S}4$. *if* $\mathscr{S} \cap \mathfrak{S}_b(\mathscr{N}) \neq \emptyset$ *then* $\mathfrak{S}_b(\mathscr{N} \cap \mathscr{S}) = \mathscr{S} \cap \mathfrak{S}_b(\mathscr{N})$.

Since a selection function is clearly not uniquely defined, the revision axioms allow for more than one type of revision. We explore a few interesting revision methods in the following sections.

8. Revision using linear orderings

In this section, I show how a revision operator can be constructed using a linear ordering on the set of maximal elements $\mathbf{M}$. The discussion here is inspired by Gärdenfors *relational partial meet contractions* [15, Section 4.4] and by the work of Moral and Wilson on revision based on linear orderings of probabilities [24].

Let us assume that the elements m of $\mathbf{M}$ are ordered by a complete pre-order, *i.e.*, a relation that is reflexive, transitive and complete, but not necessarily antisymmetrical. This is equivalent to assuming that there is a complete chain $\langle K, \preceq \rangle$ and a map $\pi\colon \mathbf{M} \to K$ which induces an ordering on $\mathbf{M}$ through the values it takes on K. We denote the top of $\langle K, \preceq \rangle$ by 1_K and its bottom by 0_K.

We can use the ordering induced on $\mathbf{M}$ to define a particular selection function $\mathfrak{S}_\pi$, as follows: for any $\mathscr{N}$ in $\overline{\mathfrak{M}}$,

$$\begin{aligned} \mathfrak{S}_\pi(\mathscr{N}) &= \{m \in \mathscr{N} : (\forall n \in \mathscr{N})(\pi(n) \preceq \pi(m))\} \\ &= \{m \in \mathscr{N} : \Pi(\mathscr{N}) \preceq \pi(m)\} \\ &= \{m \in \mathscr{N} : \Pi(\mathscr{N}) = \pi(m)\} \end{aligned} \tag{4}$$

where $\Pi(\mathscr{N}) = \sup_{m \in \mathscr{N}} \pi(m)$, so Π is the K-valued possibility measure, defined on $\wp(\mathbf{M})$, with distribution π [8]. We can now ask what properties π must have for $\mathfrak{S}_\pi$ to satisfy $\mathfrak{S}1$–$\mathfrak{S}3$. It is no essential restriction to assume that Π is *normal* in the sense that $\Pi(\mathbf{M}) = \sup_{m \in \mathbf{M}} \pi(m) = 1_K$.

Theorem 8. *Let* $\langle \mathbf{S}, \mathbf{C}, \leq \rangle$ *be a strong belief structure, and let* $b \in \mathbf{C}$ *be a coherent belief model. Let* π *be an* $\mathbf{M} - K$*-map such that the* K*-valued possibility measure* Π *with distribution* π *is normal. Consider the selection function* $\mathfrak{S}_\pi$ *defined by* (4). *Then* $\mathfrak{S}_\pi$ *satisfies* $\mathfrak{S}1$–$\mathfrak{S}3$ *if and only if*

$\pi 1$. $\mathscr{M}(b) = \{m \in \mathbf{M} \colon \pi(m) = 1_K\}$*; and*

$\pi 2$. *for every* $\mathscr{N}$ *in* $\mathfrak{M}$, $\{m \in \mathscr{N} : \pi(m) = \Pi(\mathscr{N})\}$ *is a non-empty closed subset of* $\mathscr{N}$.

In that case $\mathfrak{S}_\pi$ *automatically also satisfies* $\mathfrak{S}4$, *and the associated dual belief revision operator satisfies* $\mathfrak{R}2$–$\mathfrak{R}5$ *and* $\mathfrak{R}7$–$\mathfrak{R}8$. *The associated belief revision operator then satisfies* $R1$–$R8$.

Note that the second condition implies in particular that the map π assumes its supremum on every closed subset of $\mathbf{M}$.

Proof. Assume that $\mathfrak{S}_\pi$ satisfies $\mathfrak{S}1$–$\mathfrak{S}3$. Since $\mathbf{M}\cap\mathscr{M}(b)=\mathscr{M}(b)\neq\emptyset$, it follows from $\mathfrak{S}1$ that

$$\mathscr{M}(b)=\mathbf{M}\cap\mathscr{M}(b)=\mathfrak{S}_\pi(\mathbf{M})=\{m\in\mathbf{M}\colon \pi(m)=1_K\},$$

which tells us that $\pi1$ holds. Next, consider any $\mathscr{N}\in\mathfrak{M}$, then $\mathscr{N}\neq\emptyset$ and using $\mathfrak{S}1$ and $\mathfrak{S}2$, $\mathfrak{S}_\pi(\mathscr{N})=\{m\in\mathscr{N}\colon \pi(m)=\Pi(\mathscr{N})\}$ is a non-empty closed subset of $\mathscr{N}$, so $\pi2$ holds. Conversely, assume that $\pi1$ and $\pi2$ hold. Consider an element $\mathscr{N}$ of $\overline{\mathfrak{M}}$. If $\mathscr{N}=\emptyset$ then obviously $\mathfrak{S}_\pi(\mathscr{N})=\emptyset$. If $\mathscr{N}\cap\mathscr{M}(b)\neq\emptyset$, then it follows from $\pi1$ that on the one hand $\Pi(\mathscr{N})=1_K$, and consequently on the other hand

$$\mathfrak{S}_\pi(\mathscr{N})=\{m\in\mathscr{N}\colon \pi(m)=1_K\}=\mathscr{N}\cap\mathscr{M}(b).$$

If $\mathscr{N}\cap\mathscr{M}(b)=\emptyset$ and $\mathscr{N}\neq\emptyset$, we know from $\pi2$ that $\mathfrak{S}_\pi(\mathscr{N})=\{m\in\mathscr{N}\colon \pi(m)=\Pi(\mathscr{N})\}$ is a non-empty closed subset of $\mathscr{N}$. We conclude that $\mathfrak{S}_\pi$ satisfies $\mathfrak{S}1$–$\mathfrak{S}3$: $\mathfrak{S}_\pi$ is a selection function, and it follows from Theorem 7 that the associated dual belief revision operator $\mathfrak{R}(\mathscr{M}(b);\cdot)=\mathfrak{S}_\pi(\cdot)$ satisfies $\mathfrak{R}2$–$\mathfrak{R}5$. To prove that it also satisfies $\mathfrak{R}7$–$\mathfrak{R}8$, we must show that $\mathfrak{S}_\pi$ satisfies $\mathfrak{S}4$. Consider $\mathscr{N}$ and $\mathscr{S}$ in $\overline{\mathfrak{M}}$ and assume that $\mathscr{S}\cap\mathfrak{S}_\pi(\mathscr{N})\neq\emptyset$. This implies that there is an $m\in\mathscr{N}\cap\mathscr{S}$ such that $\pi(m)=\Pi(\mathscr{N})$, whence $\Pi(\mathscr{N})=\Pi(\mathscr{N}\cap\mathscr{S})$. Consequently,

$$\begin{aligned}\mathfrak{S}_\pi(\mathscr{N}\cap\mathscr{S})&=\{m\in\mathscr{N}\cap\mathscr{S}\colon \pi(m)=\Pi(\mathscr{N}\cap\mathscr{S})\}\\&=\{m\in\mathscr{N}\cap\mathscr{S}\colon \pi(m)=\Pi(\mathscr{N})\}\\&=\mathscr{S}\cap\{m\in\mathscr{N}\colon \pi(m)=\Pi(\mathscr{N})\}\\&=\mathscr{S}\cap\mathfrak{S}_\pi(\mathscr{N}),\end{aligned}$$

so $\mathfrak{S}_\pi$ satisfies $\mathfrak{S}4$. The rest of the proof is now immediate.

9. Revision using a system of spheres

I have called the elements $\mathscr{N}$ of a dual belief structure $\langle\mathfrak{M},\subseteq\rangle$ *spheres* because they are natural generalisations of the spheres studied by Grove [18] in the context of belief revision in classical propositional logic (see also [15, Section 4.5]). Sets of such spheres also have an important part in conditional logic, and selection functions based on such sets of spheres are also very common in this framework (see for instance [27]). Indeed, Lewis [21, 22] was probably the first to introduce 'sets of spheres' for providing appropriate semantics in conditional logic.

In this section, I show that the *generalised spheres* can also be used to construct a revision operator.

Let $b \in \mathbf{C}$ be a coherent belief model, so $\mathscr{M}(b) \neq \emptyset$. We call $\sigma(b)$ the collection of spheres that include $\mathscr{M}(b)$:

$$\sigma(b) = \{\mathscr{N} \in \mathfrak{M} \colon \mathscr{M}(b) \subseteq \mathscr{N}\},$$

so the elements $\mathscr{N}$ of $\sigma(b)$ correspond to coherent belief models $\inf \mathscr{N} \leq b$ that are less informative than b. Note that $\sigma(b)$ is an intersection structure (Moore family) with bottom $\mathscr{M}(b)$ and top $\mathbf{M}$ (it is closed under arbitrary intersections). The following definition generalises Grove's notion of a system of spheres, but note that contrary to Grove, I do not require that the elements of a system of spheres should be linearly ordered by set inclusion.[9]

Definition 7. Let $\langle \mathbf{S}, \mathbf{C}, \leq \rangle$ be a strong belief structure and let $b \in \mathbf{C}$ be a coherent belief model, so that $\mathscr{M}(b) \neq \emptyset$. We call $\sigma \subseteq \mathfrak{M}$ a *system of spheres* around $\mathscr{M}(b)$ if

$\sigma 1$. $\sigma \subseteq \sigma(b)$, *i.e.* $(\forall \mathscr{N} \in \sigma)(\mathscr{M}(b) \subseteq \mathscr{N})$;

$\sigma 2$. $\mathscr{M}(b) \in \sigma$ and $\mathbf{M} \in \sigma$;

$\sigma 3$. $\bigcap\{\mathscr{N} \cap \mathscr{S} \colon \mathscr{S} \in \sigma \text{ and } \mathscr{N} \cap \mathscr{S} \neq \emptyset\} \neq \emptyset$ for all $\mathscr{N} \in \mathfrak{M}$.

Given a system of spheres σ around $\mathscr{M}(b)$, we define a selection function $\mathfrak{S}_\sigma$ in the spirit of Grove [15, 18]: for any $\mathscr{S} \in \overline{\mathfrak{M}}$,

$$\begin{aligned}\mathfrak{S}_\sigma(\mathscr{S}) &= \bigcap\{\mathscr{S} \cap \mathscr{N} \colon \mathscr{N} \in \sigma \text{ and } \mathscr{S} \cap \mathscr{N} \neq \emptyset\} \\ &= \mathscr{S} \cap \bigcap\{\mathscr{N} \in \sigma \colon \mathscr{S} \cap \mathscr{N} \neq \emptyset\}. \quad (5)\end{aligned}$$

This selection leads to a very convenient type of revision operator, as the following theorem shows.

Theorem 9. *Let* $\langle \mathbf{S}, \mathbf{C}, \leq \rangle$ *be a strong belief structure and let* $b \in \mathbf{C}$ *be a coherent belief model, so that* $\mathscr{M}(b) \neq \emptyset$*. Let* σ *be a system of spheres around* $\mathscr{M}(b) \neq \emptyset$ *and let* $\mathfrak{S}_\sigma$ *be the associated selection function, defined by* (5)*. Then* $\mathfrak{S}_\sigma$ *satisfies* $\mathfrak{S}1$–$\mathfrak{S}4$*, and the corresponding dual belief revision operator satisfies* $\mathfrak{R}2$–$\mathfrak{R}5$ *and* $\mathfrak{R}7$–$\mathfrak{R}8$*. The corresponding belief revision operator then satisfies* $R1$–$R8$*.*

Proof. We only have to prove that $\mathfrak{S}_\sigma$ satisfies $\mathfrak{S}1$–$\mathfrak{S}4$. Consider $\mathscr{N} \in \overline{\mathfrak{M}}$. If $\mathscr{N} \cap \mathscr{M}(b) \neq \emptyset$, then all elements of σ intersect with $\mathscr{N}$ [use $\sigma 1$], so $\bigcap\{\mathscr{S} \in \sigma \colon \mathscr{S} \cap \mathscr{N} \neq \emptyset\} = \bigcap \sigma = \mathscr{M}(b)$ [use $\sigma 1$ and $\sigma 2$], and consequently

[9] Indeed, this requirement seems unnecessary, and even tends to hide interesting structure, as it emerges in Theorem 10.

$\mathfrak{S}_\sigma(\mathscr{N}) = \mathscr{N} \cap \mathscr{M}(b)$, so $\mathfrak{S}1$ holds. Obviously, if $\mathscr{N} = \emptyset$ then $\mathfrak{S}_\sigma(\mathscr{N}) = \emptyset$, so $\mathfrak{S}3$ holds. Assume that $\mathscr{N} \neq \emptyset$ and $\mathscr{N} \cap \mathscr{M}(b) = \emptyset$. Then $\mathfrak{S}_\sigma(\mathscr{N})$ is non-empty [use $\sigma 3$], a subset of $\mathscr{N}$, and closed as an intersection of closed sets. We conclude that $\mathfrak{S}3$ holds. Let $\mathscr{N}$ and $\mathscr{S}$ be elements of $\overline{\mathfrak{M}}$ such that $\mathscr{S} \cap \mathfrak{S}_\sigma(\mathscr{N}) \neq \emptyset$. On the one hand, since $\{\mathscr{N}' \in \sigma\colon \mathscr{N}' \cap \mathscr{S} \cap \mathscr{N} \neq \emptyset\} \subseteq \{\mathscr{N}' \in \sigma\colon \mathscr{N}' \cap \mathscr{N} \neq \emptyset\}$, it follows that

$$\begin{aligned}
\mathscr{S} \cap \mathfrak{S}_\sigma(\mathscr{N}) &= \mathscr{S} \cap \mathscr{N} \cap \bigcap \{\mathscr{N}' \in \sigma\colon \mathscr{N}' \cap \mathscr{N} \neq \emptyset\} \\
&\subseteq \mathscr{S} \cap \mathscr{N} \cap \bigcap \{\mathscr{N}' \in \sigma\colon \mathscr{N}' \cap \mathscr{S} \cap \mathscr{N} \neq \emptyset\} \\
&= \mathfrak{S}_\sigma(\mathscr{S} \cap \mathscr{N}).
\end{aligned}$$

Conversely, call $\mathscr{N}_o = \bigcap \{\mathscr{N}' \in \sigma\colon \mathscr{N}' \cap \mathscr{N} \neq \emptyset\}$. Then $\mathfrak{S}_\sigma(\mathscr{N}) = \mathscr{N} \cap \mathscr{N}_o \neq \emptyset$, and it follows from the assumption that $\mathscr{S} \cap \mathscr{N} \cap \mathscr{N}_o = \mathscr{S} \cap \mathfrak{S}_\sigma(\mathscr{N})$ is non-empty. Consequently, $\bigcap \{\mathscr{N}' \in \sigma\colon \mathscr{N}' \cap \mathscr{S} \cap \mathscr{N} \neq \emptyset\} \subseteq \mathscr{N}_o$, whence $\mathfrak{S}_\sigma(\mathscr{S} \cap \mathscr{N}) \subseteq \mathscr{S} \cap \mathscr{N} \cap \mathscr{N}_o = \mathscr{S} \cap \mathfrak{S}_\sigma(\mathscr{N})$.

It is not clear to me whether any belief revision operator satisfying $R1$–$R8$, or any selection function satisfying $\mathfrak{S}1$–$\mathfrak{S}4$, can be generated by a system of spheres, or in other words, whether Grove's characterisation result [18] for belief revision in classical propositional logic can be extended (but see Proposition 12). The following results should be seen as a first step toward answering this interesting open question.

Theorem 10. *Let $\mathfrak{S}_b$ be a selection function satisfying $\mathfrak{S}1$–$\mathfrak{S}3$ and define $\sigma_o \subseteq \sigma(b)$ as*[10]

$$\bigcap_{\substack{\mathscr{S} \in \mathfrak{M} \\ \mathscr{S} \cap \mathscr{M}(b) = \emptyset}} \{\mathscr{N} \in \sigma(b)\colon \mathscr{N} \cap \mathscr{S} \neq \emptyset \Rightarrow \mathfrak{S}_b(\mathscr{S}) \subseteq \mathscr{N} \cap \mathscr{S}\}.$$

Then σ_o is a system of spheres around $\mathscr{M}(b)$ and it is the greatest (finest) such system for which $\mathfrak{S}_b(\mathscr{S}) \subseteq \mathfrak{S}_{\sigma_o}(\mathscr{S})$ for all $\mathscr{S} \in \overline{\mathfrak{M}}$, with equality if $\mathscr{S} = \emptyset$ or $\mathscr{S} \cap \mathscr{M}(b) \neq \emptyset$. Consequently, there is a system of spheres that generates $\mathfrak{S}_b$ if and only if σ_o generates $\mathfrak{S}_b$, i.e., if $\mathfrak{S}_b = \mathfrak{S}_{\sigma_o}$.

Proof. It is obvious that σ_o satisfies $\sigma 1$ and $\sigma 2$. It is also clear from the definition of σ_o that for all $\mathscr{S} \in \overline{\mathfrak{M}}$:

$$\mathfrak{S}_b(\mathscr{S}) \subseteq \bigcap \{\mathscr{N} \cap \mathscr{S}\colon \mathscr{N} \in \sigma_o \text{ and } \mathscr{N} \cap \mathscr{S} \neq \emptyset\}. \tag{6}$$

Since for $\mathscr{S} \in \mathfrak{M}$, $\mathscr{S} \neq \emptyset$ and therefore $\mathfrak{S}_b(\mathscr{S}) \neq \emptyset$ [use $\mathfrak{S}1$–$\mathfrak{S}3$], it follows that $\bigcap \{\mathscr{S} \cap \mathscr{N}\colon \mathscr{N} \in \sigma_o \text{ and } \mathscr{S} \cap \mathscr{N} \neq \emptyset\} \neq \emptyset$, so σ_o satisfies $\sigma 3$ and is therefore a system of spheres around $\mathscr{M}(b)$. It also follows from (5) and (6)

[10] Note that σ_o is closed under arbitrary intersections.

that for the associated selection function $\mathfrak{S}_{\sigma_o}$: $\mathfrak{S}_b(\mathscr{S}) \subseteq \mathfrak{S}_{\sigma_o}(\mathscr{S})$ for all $\mathscr{S} \in \overline{\mathfrak{M}}$. Clearly, equality holds if $\mathscr{S} = \emptyset$ or if $\mathscr{S} \cap \mathscr{M}(b) \neq \emptyset$. Now let σ be a system of spheres around $\mathscr{M}(b)$ such that $\mathfrak{S}_b(\mathscr{S}) \subseteq \mathfrak{S}_\sigma(\mathscr{S})$ for all $\mathscr{S} \in \overline{\mathfrak{M}}$. Let $\mathscr{N}$ be an arbitrary element of σ and let $\mathscr{S} \in \mathfrak{M}$ such that $\mathscr{S} \cap \mathscr{M}(b) = \emptyset$. If $\mathscr{N} \cap \mathscr{S} \neq \emptyset$ then it follows from (5) and the assumption that $\mathfrak{S}_b(\mathscr{S}) \subseteq \mathfrak{S}_\sigma(\mathscr{S}) \subseteq \mathscr{N} \cap \mathscr{S}$, whence $\mathscr{N} \in \sigma_o$. We conclude that $\sigma \subseteq \sigma_o$. The rest of the proof is now trivial.

Corollary 11. *A selection function $\mathfrak{S}_b$ satisfying* $\mathfrak{S}1$–$\mathfrak{S}4$ *can be generated by some system of spheres if and only if for all $\mathscr{S} \in \mathfrak{M}$ such that $\mathscr{S} \cap \mathscr{M}(b) = \emptyset$ there is an $\mathscr{N} \in \sigma_o$ such that $\mathscr{S} \cap \mathscr{N} = \mathfrak{S}_b(\mathscr{S})$.*

Proposition 12 gives a simple necessary condition for a revision operator to be generated by a system of spheres. This condition is satisfied for any revision operator satisfying $R1$–$R8$ in the case of belief models based on classical propositional logic, as in that case the union of two spheres is a sphere (the closure operator $\mathrm{Cl}_{\mathbf{M}}$ is topological).

Proposition 12. *A necessary condition for a selection function $\mathfrak{S}_b$ that satisfies* $\mathfrak{S}1$–$\mathfrak{S}4$ *to be generated by some system of spheres is that for all $\mathscr{N} \in \overline{\mathfrak{M}}$:*

$$\mathfrak{S}_b(\mathscr{N}) = \mathscr{N} \cap \mathrm{Cl}_{\mathbf{M}}(\mathscr{M}(b) \cup \mathfrak{S}_b(\mathscr{N})). \tag{7}$$

Proof. Assume that $\mathfrak{S}_b$ satisfies $\mathfrak{S}1$–$\mathfrak{S}4$ and that it is generated by some system of spheres σ. Consider $\mathscr{N} \in \overline{\mathfrak{M}}$. It is clear that (7) holds if $\mathscr{N} = \emptyset$ [use $\mathfrak{S}3$] or if $\mathscr{N} \cap \mathscr{M}(b) \neq \emptyset$ [use $\mathfrak{S}1$]. Assume therefore that $\mathscr{N} \neq \emptyset$ and $\mathscr{N} \cap \mathscr{M}(b) = \emptyset$. Then we know, using (5), that $\mathfrak{S}_b(\mathscr{N}) = \mathscr{N} \cap \bigcap\{\mathscr{S} \in \sigma \colon \mathscr{S} \cap \mathscr{N} \neq \emptyset\}$. Since $\sigma(b)$ is closed under arbitrary intersections, this means that there is an $\mathscr{S} \in \sigma(b)$ such that $\mathfrak{S}_b(\mathscr{N}) = \mathscr{S} \cap \mathscr{N}$. As a consequence, $\mathfrak{S}_b(\mathscr{N}) \subseteq \mathscr{S}$ and $\mathscr{M}(b) \subseteq \mathscr{S}$, whence, since $\mathscr{S}$ is closed,

$$\mathscr{M}(b) \cup \mathfrak{S}_b(\mathscr{N}) \subseteq \mathrm{Cl}_{\mathbf{M}}(\mathscr{M}(b) \cup \mathfrak{S}_b(\mathscr{N})) \subseteq \mathscr{S},$$

and if we take the intersection with $\mathscr{N}$, taking into account that $\mathscr{N} \cap \mathscr{M}(b) = \emptyset$ and $\mathfrak{S}_b(\mathscr{N}) \subseteq \mathscr{N}$ [use $\mathfrak{S}1$–$\mathfrak{S}3$],

$$\mathfrak{S}_b(\mathscr{N}) \subseteq \mathscr{N} \cap \mathrm{Cl}_{\mathbf{M}}(\mathscr{M}(b) \cup \mathfrak{S}_b(\mathscr{N})) \subseteq \mathscr{N} \cap \mathscr{S} = \mathfrak{S}_b(\mathscr{N}),$$

which completes the proof.

Example 1. Consider the smallest (or coarsest) system of spheres around $\mathscr{M}(b)$: $\sigma = \{\mathscr{M}(b), \mathbf{M}\}$. The corresponding selection function is given by

$$\mathfrak{S}_\sigma(\mathscr{S}) = \begin{cases} \mathscr{S} \cap \mathscr{M}(b) & \text{if } \mathscr{S} \cap \mathscr{M}(b) \neq \emptyset \\ \mathscr{S} & \text{if } \mathscr{S} \cap \mathscr{M}(b) = \emptyset. \end{cases}$$

so we find for the corresponding revision operator:

$$R(b;\gamma) = \begin{cases} \mathrm{Cl}_{\mathbf{S}}(b \smile \gamma) & \text{if } b \text{ and } \gamma \text{ are consistent} \\ \gamma & \text{if } b \text{ and } \gamma \text{ are inconsistent,} \end{cases}$$

In the spirit of Gärdenfors's work [15], we could call this $R(b;\cdot)$ a 'full meet revision'.

Example 2. Consider a normal possibility distribution $\pi\colon \mathbf{M} \to K$ on the set of maximal coherent belief models $\mathbf{M}$. We assume that it satisfies $\pi 1$ and that its cut sets are closed: $\pi_\alpha = \{m \in \mathbf{M}\colon \alpha \preceq \pi(m)\} \in \mathfrak{M}$ for all $\alpha \in K$. This implies in particular that $\pi 2$ is also satisfied. Define the following collection of closed subsets of $\mathbf{M}$:

$$\sigma_\pi = \{\pi_\alpha\colon \alpha \in K\}.$$

It follows from $\pi 1$ that for all $\alpha \in K$, $\pi_\alpha \supseteq \pi_{1_K} = \mathscr{M}(b)$. Since moreover $\pi_{0_K} = \mathbf{M}$, we see that σ_π satisfies $\sigma 1$ and $\sigma 2$. Next, consider $\mathscr{N} \in \mathfrak{M}$. Since it follows from $\pi 2$ that π assumes its supremum on every closed set $\mathscr{N} \in \mathfrak{M}$, we have for all $\alpha \in K$ that $\mathscr{N} \cap \pi_\alpha \neq \emptyset$ if and only if $\alpha \preceq \Pi(\mathscr{N})$, whence

$$\bigcap\{\pi_\alpha\colon \mathscr{N} \cap \pi_\alpha \neq \emptyset\} = \bigcap\{\pi_\alpha\colon \alpha \preceq \Pi(\mathscr{N})\} = \{m \in \mathbf{M}\colon \Pi(\mathscr{N}) \preceq \pi(m)\}$$

and taking into account $\pi 2$ and (4),

$$\begin{aligned} \mathfrak{S}_{\sigma_\pi}(\mathscr{N}) &= \mathscr{N} \cap \bigcap\{\pi_\alpha\colon \mathscr{N} \cap \pi_\alpha \neq \emptyset\} \\ &= \{m \in \mathscr{N}\colon \Pi(\mathscr{N}) = \pi(m)\} = \mathfrak{S}_\pi(\mathscr{N}) \neq \emptyset. \end{aligned}$$

This proves that $\sigma 3$ holds, so σ_π is a system of spheres around $\mathscr{M}(b)$. We find for the corresponding selection operator that $\mathfrak{S}_{\sigma_\pi} = \mathfrak{S}_\pi$.

10. Conclusion

I am convinced that the study of belief structures, their mathematical properties and their mutual relationships, can help us relate the many belief models that have been proposed in the literature. I am aware that the present study is far from complete, and that refinements and even small modifications may be necessary. One topic where this may be the case, is belief contraction. We have seen that for belief expansion and revision, many of Gärdenfors's results are valid in a broader context. Although his proofs use the details of the underlying logical language, I have shown that this is not necessary, and that simpler and more powerful proofs can be found by using a few general unifying properties. It turns out, however, that in Gärdenfors's discussion of contraction crucial steps are taken which are very specific to classical logic (using the topological nature of the closure $\mathrm{Cl}_{\mathbf{M}}$, for one thing); and which are hard, if not impossible, to generalise directly. For one thing, preserving the relationship between contraction and revision (Levi's and Harper's identities) becomes problematical. More effort should be invested in finding out what can said about belief contraction for more general belief models, what can be preserved in the generalisation, and how.

Acknowledgements

This paper is an updated and extended version of a paper presented at the Second International Symposium on Imprecise Probabilities and Their Applications (ISIPTA '01, Cornell University, 26–29 June 2001). It presents research results of the Belgian Program on Interuniversity Poles of Attraction initiated by the Belgian state, Prime Minister's Office for Science, Technology and Culture; and of research project nr. G.0139.01 of the Fund for Scientific Research, Flanders (Belgium). The scientific responsibility rests with me, the author.

I am indebted to Serafín Moral, Peter Walley and Teddy Seidenfeld, who have read and commented on an earlier version of the paper. I would also like to thank a number of anonymous reviewers for helpful suggestions, and for additional pointers to the relevant literature.

Finally, and before anything else, I wish to dedicate this paper to Minou, the well and the rose-tree.

References

1. R. J. Aumann. Utility theory without the completeness axiom. *Econometrica*, 30:445–462, 1962.
2. R. J. Aumann. Utility theory without the completeness axiom: a correction. *Econometrica*, 32:210–212, 1964.
3. S. Benferhat, D. Dubois, and H. Prade. Possibilistic and standard probabilistic semantics of conditional knowledge bases. *Journal of Logic and Computation*, 9:873–895, 1999.
4. L. Boyen, G. de Cooman, and E. E. Kerre. On the extension of P-consistent mappings. In G. de Cooman, D. Ruan, and E. E. Kerre, editors, *Foundations and Applications of Possibility Theory – Proceedings of FAPT '95*, pages 88–98, Singapore, 1995. World Scientific.
5. A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.
6. B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, Cambridge, 1990.
7. G. de Cooman. Confidence relations and ordinal information. *Information Sciences*, 104:241–278, 1997.
8. G. de Cooman. Possibility theory I–III. *International Journal of General Systems*, 25:291–371, 1997.

9. G. de Cooman. Precision–imprecision equivalence in a broad class of imprecise hierarchical uncertainty models. *Journal of Statistical Planning and Inference*, 105:175–198, 2002.
10. G. de Cooman and P. Walley. A possibilistic hierarchical model for behaviour under uncertainty. *Theory and Decision*, 52:327–374, 2002.
11. B. de Finetti. *Teoria delle Probabilità*. Einaudi, Turin, 1970.
12. B. de Finetti. *Theory of Probability*, volume 1. John Wiley & Sons, Chichester, 1974. English Translation of [11].
13. D. Dubois and H. Prade. A survey of belief revision and updating rules in various uncertainty models. *International Journal of Intelligent Systems*, 9:61–100, 1994.
14. D. Dubois and H. Prade. Possibility theory: qualitative and quantitative aspects. In Ph. Smets, editor, *Handbook on Defeasible Reasoning and Uncertainty Management Systems*, volume Volume 1: Quantified Representation of Uncertainty and Imprecision, pages 169–226. Kluwer Academic Publishers, Dordrecht, 1998.
15. P. Gärdenfors. *Knowledge in Flux – Modeling the Dynamics of Epistemic States*. The MIT Press, Cambridge, MA, 1988.
16. R. Giles. Foundations for a theory of possibility. In M. M. Gupta and E. Sanchez, editors, *Fuzzy Information and Decision Processes*, pages 183–195. North-Holland, Amsterdam, 1982.
17. F. J. Giron and S. Rios. Quasi-Bayesian behaviour: A more realistic approach to decision making? In J. M. Bernardo, M. H. DeGroot, D. V. Lindley, and A. F. M. Smith, editors, *Bayesian Statistics*, pages 17–38. Valencia University Press, Valencia, 1980.
18. A. J. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, pages 157–170, 1988.
19. G. Kern-Isberner. The principle of conditional preservation in belief revision. In *Proceedings of the Second International Symposium on Foundations of Information and Knowledge Systems (FoIKS 2002)*, Lecture Notes in Computer Science 2284, pages 105–129. Springer, 2002.
20. I. Levi. *The Enterprise of Knowledge*. MIT Press, London, 1980.
21. D. K. Lewis. *Counterfactuals*. Harvard University Press, Cambridge, MA, 1973.
22. D. K. Lewis. Counterfactuals and comparative possibility. *Journal of Philosophical Logic*, 2, 1973.
23. D. Makinson. General patterns in non-monotonic reasoning. In D. M. Gabbay, C. H. Hogger, and J. A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3, pages 35–110. Oxford University Press, Oxford, 1994.
24. S. Moral and N. Wilson. Revision rules for convex sets of probabilities. In G. Coletti, D. Dubois, and R. Scozzafava, editors, *Mathematical Models for Handling Partial Knowledge in Artificial Intelligence*, pages 113–128. Plenum Press, New York, 1995.
25. R. Nau. The shape of incomplete preferences. In J.-M. Bernard, T. Seidenfeld, and M. Zaffalon, editors, *ISIPTA '03 – Proceedings of the Third International Symposium on Imprecise Probabilities and Their Applications*, pages 408–422. Carleton Scientific, 2003.
26. R. Nau. The shape of incomplete preferences. Technical report, Decision Analysis Working Paper WP030009, September 2003. Extended version of [25].
27. D. Nute. *Topics in Conditional Logic*. D. Reidel Publishing Company, Dordrecht, Holland, 1980.
28. E. Schechter. *Handbook of Analysis and Its Foundations*. Academic Press, San Diego, CA, 1997.
29. T. Seidenfeld, M. J. Schervish, and J. B. Kadane. A representation of partially ordered preferences. *The Annals of Statistics*, 23:2168–2217, 1995. Reprinted in [30], pp. 69–129.

30. T. Seidenfeld, M. J. Schervish, and J. B. Kadane. *Rethinking the Foundations of Statistics*. Cambridge University Press, Cambridge, 1999.
31. W. Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In W. L. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change and Statistics*, volume 2, pages 105–134. Reidel, Dordrecht, 1987.
32. P. Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London, 1991.
33. M. A. Williams. Transmutation of knowledge systems. In *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning*, pages 619–629. Morgan Kaufmann, 1994.