AIDAN SIMS AND BRAILEY SIMS

ASSESSING THE MINIMALITY OF CHANGE IN BELIEF REVISION: CAPACITIES

1 INTRODUCTION

The process of belief revision [Alchourrón *et al.*, 1985] was developed to model the effect of accepting incoming information into a knowledge base. It can be based on a total preordering of the set of beliefs held by the system. When this ordering relates in a natural way to the semantics of the underlying knowledge base, it is dubbed an epistemic entrenchment ordering [Gärdenfors and Makinson, 1988]. This ranks beliefs in order of the agent's reluctance to give them up. Thus the belief *the sun will rise tomorrow* would be ranked higher than *the universe started with a big bang*. Such an ordering leads to the construction of a unique AGM transformation on the knowledge base to allow for the acceptance of new information; information which may be inconsistent with the current beliefs. This original work was purely theoretical, concerning itself only with one revision. The outcome was a revised set of beliefs, but the question of how those beliefs could be fitted with a "revised" ranking was not addressed. This question is, however, not only theoretically interesting but also vital for implementation since, in that situation, it is necessary to iterate the procedure and so a new epistemic entrenchment needs to be the outcome of each revision.

Later work extended this system to accomodate iterated belief revision, the process of moving under revision from one entrenchment to another being termed a transmutation. Following the lead of earlier work on Bayesian and Jeffrey conditionalisation, and Spohn's ordinal conditional functions [Spohn, 1988], this was effected by allowing new information to be inserted with a rank allocated prior to insertion [Williams, 1994a]. This provided revision operations for ensconcements [Williams, 1994b], which allowed for an actual computer-based implementation of iterated belief revision [Williams, 1995; Williams, 1997a; Williams, 1997b].

Since then, a number of algorithms for determining the set of beliefs to be discarded in the event of a revision have been proposed and implemented with the focus generally either on speed of execution (for example Linear revision [Nebel, 1994], and Adjustment [Williams, 1995]), or intuitiveness of the output (for example Maxi Adjustment [Williams, 1997a]). Typically, the algorithm moves through the ranking until it determines the largest 'cut' in which no inconsistency exists, and then locates a minimal set of beliefs — either in the next rank, or in the rest of the ensconcement — whose removal eliminates inconsistencies.

In every case, the heart of the algorithm lies in the method for deciding on a set of beliefs to discard. Intuitively, this set should be in some sense minimal, but various algorithms have differed considerably as to how this should be interpreted.

M.-A. Williams and H. Rott (eds.), Frontiers in Belief Revision, 163–177.

Qualitatively, a notion of minimality was captured by imposing heuristically plausible constraints on the transformations to be applied. However, little has been done to provide quantative measures of this minimality in this context.

We introduce the notion of capacity for a belief set. In order that a capacity reflects the underlying desire for minimal change, it is tied monotonically to the set-theoretic structure of the knowledge base, as well as to the preference ordering associated to it. The axioms have, however, been kept to a minimum so as to permit as much freedom as possible in the construction of capacities.

Examples showing how capacities may be constructed are presented. A theory of capacities is then developed; in particular some general properties of capacties are examined, and general techniques for constructing new capacities from given ones are explored. We then go on to demonstrate that there is a capacity associated to each epistemic entrenchment which exposes the sense in which the beliefs removed in the corresponding AGM transformation are minimal.

Our final consideration of capacities relates to the problem of simultaneous revision with respect to multiple beliefs. The behaviours of the existing standard belief revision strategies when used to revise with respect to multiple beliefs depend quite heavily on the order in which the new beliefs are added to the system, and require that a position in the ranking for each new belief be provided. While this can be a beneficial aspect of belief revision, it can also be restrictive in that it requires some external system (either user or calling program) to determine where new beliefs should lie in the ranking. Capacities can be used to provide an alternative to these iterated belief revision strategies; one that does not rely on a rank being preassigned to each new belief, and that is not dependent upon the order in which the beliefs are added.

The aim of capacities, then, is to formalise the concept of minimality by describing a set of axioms that circumscribe a family of functions that can be used naturally to identify minimal sets of beliefs.

2 CAPACITIES

2.1 *Definitions*

Before presenting the definition of capacities, we recall the definitions of a number of well-known concepts.

DEFINITION 1. A *theory* is a set T of beliefs (logical sentences) that is closed under the operation of entailment, that is, $A \subseteq T$ and $A \vdash b$ implies $b \in T$. The *theory closure* of a set X of beliefs, denoted $\mathrm{Th}(X)$, is the closure of X under entailment. That is, $\mathrm{Th}(X) := \{y | X \vdash y\}$.

DEFINITION 2. Following [Gärdenfors and Makinson, 1988], an *epistemic entrenchment* is a pair $(T, \leq)$ where T is a theory, and $\leq$ is a total preordering on T in which tautology is ranked strictly higher than any other belief and which is such that, for any $b \in T$, $\mathrm{Th}(\{a \in T | a \geq b\}) \cap \{c \in T | c < b\} = \emptyset$.

REMARK 3. To avoid techinical distractions, we will restrict our attention to finitely representible epistemic entrenchments [Williams, 1994b]. That is, an epistemic entrenchment $(T, \leq)$ in which there are only a finite number of distinct subtheories ('cuts') of the form $C_a := \{b \in T | a \leq b\}$ where $a \in T$, each of which is finitely axiomatisible. Such entrenchments are of prime importance from the point of view of implementation of AGM belief revision. Thus, throughout what follows, epistemic entrenchment should be understood to mean finitely representible epistemic entrenchment.

REMARK 4. In an epistemic entrenchment $(T, \leq)$, it is not uncommon to associate a number in $[0, 1]$ with the beliefs in each of the ranks, where 1 is reserved for tautologies, whilst 0 is reserved for contradictions, and also assigned to non-beliefs; that is, beliefs not in T.

DEFINITION 5. An *entrenchment basis* for a finitely representible epistemic entrenchment $(T, \leq)$ is a finite subset $B \subseteq T$, together with the restriction of $\leq$ to B (denoted $(B, \leq)$), such that for each $a \in T, C_a = \mathrm{Th}(B \cap C_a)$. Such bases always exist, and have elsewhere been referred to as ensconcements [Williams, 1994b].

For our purposes, it will be convenient to introduce a more general concept.

DEFINITION 6. A *belief ranking* is a finite set X of beliefs with a total preordering $\leq$. We will refer to such a ranking by the pair $(X, \leq)$.

REMARK 7. It should be noted that there is no requirement that a belief ranking relate to the underlying logical structure of the set X, and so the concept of a belief ranking is more general than that of an entrenchment basis. Thus every entrenchment basis is a belief ranking, but the converse is not, in general, true.

DEFINITION 8. A function $m : \mathcal{P}(X) \to [0, \infty)$ is called a *capacity* on the belief ranking $(X, \leq)$ if for all $A, B \subseteq X$,

a) $m(\emptyset) = 0$;
b) $A \subseteq B \Rightarrow m(A) \leq m(B)$; and
c) $m(A) \leq m((A \setminus \{x\}) \cup \{y\}) \Leftrightarrow x \leq y$ whenever $x \in A$ and $y \notin A$.

REMARK 9. The requirement that $m(\emptyset) = 0$ is not entirely necessary, but it does lead to nicer behaviour. In particular, as we shall see later, it ensures that the complement of the complement of a capacity is, in fact, equal to the original capacity.

REMARK 10. Although we define capacities as functions from $\mathcal{P}(X)$ to $[0, \infty)$, it should be clear that every capacity m is equivalent in every respect to the capacity m' which maps from $\mathcal{P}(X)$ to $[0, 1]$ defined by $m'(A) := m(A)/m(X)$. For this reason, we shall assume that every capacity maps to $[0, 1]$ rather than to $[0, \infty)$, and that $m(X) = 1$ for every capacity m on a belief ranking $(X, \leq)$.

DEFINITION 11. Suppose $(X, \leq)$ is a belief ranking, and define $m : \mathcal{P}(X) \to [0, 1]$ by $m(A) = |A|/|X|$ for every $A \subseteq X$. This will be referred to as *counting measure*, and is useful in the construction of various capacities.

We now give two illustrative examples of capacities.

EXAMPLE 12. Suppose $(X, \leq)$ is a belief ranking, and define a function $r : X \to \mathbb{N}$ by defining $r(x)$ to be the number of distinct ranks below the one containing x. For example, $r(\perp) = 0$. That is, the function r assigns to each belief its rank counted up from the bottom of the belief ranking. We can then define a capacity m by

$$m(A) := \sum_{x \in A} r(x) \Big/ \sum_{x \in X} r(x)$$

Note that this example reduces to the counting measure when there is only one rank in the belief ranking.

EXAMPLE 13. Again suppose $(X, \leq)$ is a belief ranking, and define $r : X \to \mathbb{N}$ as in 12. Then for each $y \in (0, 1]$ we can define a capacity m_y by

$$m_y(A) := \sum_{x \in A} y^{-r(x)} \Big/ \sum_{x \in X} y^{-r(x)}$$

We refer to this set of capacities as the *geometric capacities*, since they are based on a geometric sequence.

2.2 *Capacity Constructions*

Given the definition of capacities, we now begin to investigate some methods for constructing new capacities from old ones. The first of these is the complement of a capacity.

DEFINITION 14. Let $(X, \leq)$ be a belief ranking, and let m be a capacity on $(X, \leq)$. We define the *complement* of m, denoted m^* by

$$m^*(A) := m(X) - m(X \setminus A)$$

for every $A \subseteq X$.

PROPOSITION 15. *Let $(X, \leq)$ be a belief ranking, and let m be a capacity on $(X, \leq)$. Then m^* is also a capacity on $(X, \leq)$.*

Proof.

a) $m^*(\emptyset) = m(X) - m(X \setminus \emptyset) = 0$

b)

$$\begin{aligned} A \subseteq B \quad &\Rightarrow \quad X \setminus B \subseteq X \setminus A \\ &\Rightarrow \quad m(X \setminus B) \leq m(X \setminus A) \\ &\Rightarrow \quad m^*(A) \leq m^*(B) \end{aligned}$$

c) Let $x \in A$ and $y \notin A$ and let $B := (A \setminus \{x\}) \cup \{y\}$. We need $x \leq y \Leftrightarrow m^*(A) \leq m^*(B)$. For this, first suppose $x \leq y$. Then:

$$\begin{aligned}
m^*(A) &= m(X) - m(X \setminus A) \\
&\leq m(X) - m(((X \setminus A) \setminus \{y\}) \cup \{x\}) \qquad (y \notin A, x \in A). \\
&= m(X) - m(X \setminus ((A \setminus \{x\}) \cup \{y\})) \\
&= m^*(B)
\end{aligned}$$

Now, suppose $m^*(A) \leq m^*(B)$. Then we have:

$$\begin{aligned}
m^*(A) \leq m^*(B) &\Rightarrow m(X \setminus A) \geq m(X \setminus B) \\
&\Rightarrow m(X \setminus A) \geq m(X \setminus ((A \setminus \{x\}) \cup \{y\})) \\
&\Rightarrow m(X \setminus A) \geq m(((X \setminus A) \setminus \{y\}) \cup \{x\}) \\
&\Rightarrow x \leq y \qquad \text{(since } m \text{ is a capacity)}
\end{aligned}$$

■

PROPOSITION 16. *Let $(X, \leq)$ be a belief ranking, and let m be a capacity on $(X, \leq)$. Then $(m^*)^* = m$.*

Proof.

$$\begin{aligned}
(m^*)^*(A) &= m^*(X) - m^*(X \setminus A) \\
&= \big(m(X) - m(\emptyset)\big) - \big(m(X) - m(X \setminus (X \setminus A))\big) \\
&= m(A) - m(\emptyset) \\
&= m(A) \text{ for all } A \subseteq X
\end{aligned}$$

■

REMARK 17. Philosophically, rather than removing a set A of minimal capacity to acheive minimal change, one might argue for the removal of a set A' whose complement has maximal capacity. 15 and 16 taken in conjunction with 14 show that there is no essential difference; it is really a choice between the use of two well-defined capacities. The latter strategy applied to the capacity m is identical to the first strategy applied to m^* and vice-versa.

In fact, the question of whether the complement of a capacity is necessary at all is not an obvious one. One would like to discover that every capacity m has the property that the ranking of the sets of beliefs given by m^* is the same as that given by m. We call capacities for which this is true *complementary*. It is not, however, the case that all capacities are complementary, though we do get the following result in the case where there is a minimum element.

LEMMA 18. *Let* $(X, \leq)$ *be a belief ranking, and* m *a capacity on* $(X, \leq)$. *If* $\emptyset \subset A \subseteq X$ *is such that* $m(A) < m(B)$ for all $B \in \mathcal{P}(X) \setminus \{A\}$, *then* $m(X \setminus A) \geq m(C)$ for all $C \subset X$.

Proof. Clearly $|A| = 1$ or else every singleton set that is a subset of A is at least as small as A by axiom (a). Hence, $A = \{a\}$ for some $a \in X$. Now suppose that $C \subset X$. If $C \neq X \setminus A$, then clearly there is a set $C' \neq X \setminus A$ with $|C'| = |X| - 1$ and $C \subseteq C'$, so that, in particular, $m(C) \leq m(C')$. But now, $C' = X \setminus \{c\}$ for some $c \in X$ so $a \in C'$, and hence $C' = (X \setminus \{c\}) \cup \{a\}$.

Then,

$$\begin{aligned} m(X \setminus A) &= m(X \setminus \{a\}) \\ &\geq m(((X \setminus \{a\}) \setminus \{c\}) \cup \{a\}), \quad \text{by axiom (b)} \\ &= m(X \setminus \{c\}) \\ &= m(C') \\ &\geq m(C) \end{aligned}$$

as required. But if $C = X \setminus A$ the result is trivially true. ■

COROLLARY 19. *Let* $(X, \leq)$ *be a belief ranking, and* m *a capacity on* $(X, \leq)$. *If* $A \subseteq X$ *is such that* $m(A) > m(B)$ for all $B \in \mathcal{P}(X) \setminus \{A\}$ *then* $m(X \setminus A) \leq m(C)$ for all $C \subseteq X$.

Proof. Since A is the maximum under m we know by the definition of the complement capacity that $X \setminus A$ is the minimum under m^*. Applying 18 to the complement capacity, this tells us that A is maximal under m^*, and again by the definition of m^*, this implies that $X \setminus A$ is minimal under m. ■

However, as the following example demonstrates, m and m^* can behave quite differently when there are no single maximum or minimum beliefs.

EXAMPLE 20. Define m on $\mathcal{P}(\{a, b, c, d\})$ by:

$$\begin{aligned} &m(\emptyset) = 0.0 \\ &m(\{a\}) = m(\{b\}) = m(\{a, b\}) = 0.1 \\ &m(\{c\}) = m(\{d\}) = m(\{a, c\}) \\ &\quad = m(\{a, d\}) = m(\{b, c\}) = m(\{b, d\}) = 0.2 \\ &m(\{c, d\}) = 0.3 \\ &m(\{a, b, c\}) = m(\{a, b, d\}) = 0.4 \\ &m(\{a, c, d\}) = m(\{b, c, d\}) = 0.5 \\ &m(\{a, b, c, d\}) = 1.0 \end{aligned}$$

Note that the example is not without substance: $\{a, b, c, d\}$, with the ordering that arises from the capacities of the singleton sets in 20, is an entrenchment basis.

REMARK 21. We see in 20 that whilst $\{a, b\}$ is minimal, $X \setminus \{a, b\} = \{c, d\}$ is not maximal, and furthermore that m and m^* rank the subsets of X differently in the sense that (for example) $m(\{a, b\}) = 0.1 < m(\{c\}) = 0.2$ but $m^*(\{a, b\}) = 0.7 > m^*(\{c\}) = 0.6$.

It would be rather nice to know precisely when a capacity is complementary. A fairly obvious sufficient condition presents itself:

LEMMA 22. *Let m be a capacity on a belief ranking $(X, \leq)$. If there is a $k \in \mathbb{R}$ so that*

$$m(A) + m(X \setminus A) = k, \ \textit{for all } A \subseteq X,$$

then m is equivalent to m^* in the sense that $m(A) \leq m(B)$ if and only if $m^*(A) \leq m^*(B)$ for all $A, B \subseteq X$.

Proof. Suppose $m(A) + m(X \setminus A) = k \ \forall A \subseteq X$. Then let $k' = k - m(X)$, and we have

$$\begin{aligned} m(A) \leq m(B) &\Leftrightarrow m(A) - k' \leq m(B) - k' \\ &\Leftrightarrow m(X \setminus A) \geq m(X \setminus B) \\ &\Leftrightarrow m(X) - m(X \setminus A) \leq m(X) - m(X \setminus B) \\ &\Leftrightarrow m^*(A) \leq m^*(B) \end{aligned}$$

■

REMARK 23. It is clear that any additive capacity satisfies the hypotheses of 22, so the existence of non-complementary capacities (20) demonstrates that there exist capacities that are not equivalent to any additive capacity. Consequently the notion of a capacity is more general than that of a measure.

We now consider combinations of two capacities on the same belief ranking that allow us to combine the characteristics of pairs of existing capacities on a set of beliefs, as well as allowing us to construct a complementary capacity from any capacity at all. Before doing this, we first demonstrate that the set of capacities on a given belief ranking is convex.

LEMMA 24. *Let $(X, \leq)$ be a belief ranking, and suppose that m and m' are both capacities on $(X, \leq)$. Then for each $\lambda \in [0, 1]$ the sum $\lambda m + (1 - \lambda)m'$ defined on X by $(\lambda m + (1 - \lambda)m')(A) = \lambda m(A) + (1 - \lambda)m'(A)$ is also a capacity.*

Proof. It is clear that $\lambda m + (1 - \lambda)m'$ is 1 at X, 0 at $\emptyset$, and takes values between these two extremes. Since the remaining axioms, (b) and (c), for a capacity hold if and only if they hold for a positive multiple of the capacity, it is sufficient to show that (b) and (c) hold for the pointwise sum of any two capacities. For this, let m and m' be any two capacities on $(X, \leq)$.

b)

$$\begin{aligned} A \subseteq B \quad &\Rightarrow \quad m(A) \leq m(B) \wedge m'(A) \leq m'(B) \qquad (m, m' \text{ capacities}) \\ &\Rightarrow \quad m(A) + m'(A) \leq m(B) + m'(B) \end{aligned}$$

c) Let $x \in A$ and $y \notin A$ and let $B := (A \setminus \{x\}) \cup \{y\}$. We need $x \leq y \Leftrightarrow m(A) + m'(A) \leq m(B) + m'(B)$. For this, first suppose $x \leq y$. Then $m(A) \leq m(B)$ since m is a capacity, and $m'(A) \leq m'(B)$ since m' is a capacity. Thus $m(A) + m'(A) \leq m(B) + m'(B)$.

Now, suppose $(m+m')(A) \leq (m+m')(B)$. Then $m(A)+m'(A) \leq m(B) + m'(B)$, so either $m(A) \leq m(B)$ or $m'(A) \leq m'(B)$. In either case, it follows that $x \leq y$ since m and m' are capacities. ■

We now move on to the capacity sum we were aiming for.

DEFINITION 25. Let m and m' be capacities on a belief ranking $(X, \leq)$. We define the *sum capacity* of m and m' by

$$(m \oplus m')(A) := \frac{m(A) + m'(A)}{2}, \quad \text{for all } A \subseteq X.$$

As an immediate consequence of 24, we have

COROLLARY 26. *Let $(X, \leq)$ be a belief ranking, and suppose that m and m' are both capacities on $(X, \leq)$. Then the sum capacity $m \oplus m'$ of m and m' is a capacity.*

REMARK 27. The division by two in the above is not necessary to define the capacity sum sensibly; the simple pointwise sum of the two functions would do equally well. The only reason for dividing by two is to conform to our convention that capacities map $\mathcal{P}(X)$ to $[0, 1]$ and that $m(X) = 1$.

Now, given any capacity m on a belief ranking, we can construct a complementary capacity from it as follows:

LEMMA 28. *Let $(X, \leq)$ be a belief ranking, and m a capacity on $(X, \leq)$. Then the sum capacity $m \oplus m^*$ is a complementary capacity.*

Proof.

$$\begin{aligned} (m \oplus m^*)(A) + (m \oplus m^*)(X \setminus A) &= \frac{1}{2}((m(A) + m^*(A)) \\ &\quad +(m(X \setminus A) + m^*(X \setminus A)) \\ &= \frac{1}{2}((m(A) + m(X) - m(X \setminus A)) \\ &\quad +(m(X \setminus A) + m(X) \\ &\quad -m(X \setminus (X \setminus A)))) \end{aligned}$$

$$\begin{aligned} &= \frac{1}{2}(2m(X)) \\ &= m(X) \\ &= 1 \end{aligned}$$

Hence, by 22 with $k = 1$ we have $(m \oplus m^*)$ a complementary capacity. ■

REMARK 29. Note that the capacity $m \oplus m^*$ may have desirable features in so far as minimising with respect to it would, in some sense, represent a balance between the two strategies discussed in 17.

Next we introduce the *extension capacity* which we use to extend an existing capacity to a slightly larger belief ranking. This capacity will be used in our development of a belief revision strategy based on capacities.

DEFINITION 30. Let m be a capacity on a belief ranking $(X, \leq)$, and suppose $a \notin X$. Suppose that $(X \cup \{a\}, \leq_a)$ is also a belief ranking, with $\leq_a$ extending $\leq$ on X. Then we define the *extension capacity* m_a on $(X \cup \{a\}, \leq_a)$ by

$$m_a(A) := \begin{cases} \max\limits_{b \leq_a a}(\max\{m(B) | B \subseteq X, (B \setminus \{b\}) \subseteq (A \setminus \{a\})\}) \text{ if } a \in A \\ m(A) \text{ if } a \notin A \end{cases}.$$

LEMMA 31. *Let m be a capacity on a belief ranking $(X, \leq)$, and suppose $a \notin X$. Suppose that $(X \cup \{a\}, \leq_a)$ is also a belief ranking, with $\leq_a$ extending $\leq$ on X. Then m_a defined above is a capacity on $(X \cup \{a\}, \leq_a)$ extending m.*

Proof. It is clear that m_a agrees with m on X, so that it is, in fact, an extension of m to $X \cup \{a\}$.

a) $m_a(\emptyset) = m(\emptyset) = 0$

b) Suppose $A \subseteq B$. Then $C \subseteq A \Rightarrow C \subseteq B$ (and specifically, $A \setminus \{x\} \subseteq B \setminus \{x\}$ for all $x \in X$), so that

$$\begin{aligned} m_a(A) &= max\{m(C) \mid C \subseteq (A \setminus \{a\})\} \\ &\leq max\{m(C) \mid C \subseteq (B \setminus \{a\})\} \\ &= m_a(B) \end{aligned}$$

c) Suppose $x \in A$, $y \notin A$, and let $B := (A \setminus \{x\}) \cup \{y\}$. In the case where $x \neq a$ and $y \neq a$ it is clear that the result follows from the definition of m_a and the fact that m is a capacity. Suppose now that $x = a$. Then

$$\begin{aligned} x \leq_a y &\Rightarrow m(B) \geq \max\{m((A \setminus \{a\}) \cup \{c\}) \mid c \leq_a a\} \\ &\Rightarrow m_a(B) \geq m_a(A) \end{aligned}$$

and

$$\begin{aligned} m_a(B) \le m_a(A) \quad &\Rightarrow \quad m_a(B) \le \max\{m(C) \mid \exists\, c \le_a a \text{ s.t. } \\ &\qquad\qquad A = (C \setminus \{c\}) \cup \{a\}\} \\ &\Rightarrow \quad m((A \setminus \{a\}) \cup \{y\}) \le m((A \setminus \{a\}) \cup \{c\}), \\ &\qquad\qquad (\text{for some } c \le_a a) \\ &\Rightarrow \quad y \le \max\{c \le_a a \mid A = (C \setminus \{c\}) \cup \{a\}\} \end{aligned}$$

Finally, suppose $y = a$ and $x \neq a$. Then $x \le_a y$ implies that A is one of the sets the maximum capacity of which is assigned to B by m_a, so $m_a(B) \ge m_a(A)$. On the other hand, $m_a(A) = m(A)$, so

$$\begin{aligned} m_a(A) \le m_a(B) \quad &\Rightarrow \quad m(A) \le \max\{m((B \setminus \{a\}) \cup \{c\}) \mid c \le_a a\} \\ &\Rightarrow \quad m(A) \le m((B \setminus \{a\}) \cup \{c\}) \text{ for some } c \le_a a \\ &\Rightarrow \quad m((B \setminus \{a\}) \cup \{x\}) \le m((B \setminus \{a\}) \cup \{c\}) \\ &\Rightarrow \quad x \le c \qquad (\text{since } m \text{ is a capacity}) \\ &\Rightarrow \quad x \le_a a \qquad \text{by the definition of } \le_a . \end{aligned}$$

■

Finally, we introduce the *substitution capacity* on a belief ranking X with regard to a subset R of X and a belief $a \notin X$.

DEFINITION 32. Suppose that $(X, \le)$ is a belief ranking, that m is a capacity on X, that $R \subseteq X$ and that $a \notin X$. Then we define the *substitution capacity* $m_{a,R}$ on $(X \setminus R) \cup \{a\}$ via the extension capacity where a is ranked in $(X \setminus R) \cup \{a\}$ in such a way that $a \le b$ if and only if $m(R) \le m(\{b\})$ and $a \ge b$ if and only if $m(R) \ge m(\{b\})$.

LEMMA 33. *For any set $R \subseteq X$ and for any belief $a \in X$ the substitution capacity $m_{a,R}$ is a capacity.*

Proof. The proof is an immediate consequence of the fact that the new ranking is, in fact, a belief ranking, and 31. ■

3 CAPACITY BASED REVISION

3.1

Suppose that we are given an entrenchment basis $(X, \le)$, a belief a, and a position we would like to assign a relative to the elements of X in $X \cup \{a\}$. This gives rise to a total ordering $\le_a$ on $X \cup \{a\}$. Then to perform a capacity belief revision with a on $(X, \le)$ we begin by testing whether $a \in X$. If so, we do nothing.

If not, we must insert a into the belief ranking, and then find the minimal set of beliefs not containing a whose removal will eliminate any contradictions in the new belief ranking. To do this, of course, we require a capacity on the belief ranking $(X \cup \{a\}, \leq_a)$. If such a capacity already exists, we use this. If we do not already have a capacity on $(X \cup \{a\}, \leq_a)$, however, we can easily generate one, namely the extension m_a of m to $(X \cup \{a\}, \leq_a)$ as defined in 30.

We now select the subset $M \subseteq X$ (so $a \notin M$) such that $\perp \notin Th(X \setminus M)$ that is minimal under m_a. This set is removed, giving a new belief ranking $((X \cup \{a\}) \setminus M, \leq_a)$. To convert the belief ranking into an entrenchment, we move any beliefs that are ranked higher than a and which would be removed by a revision with $\neg a$ down to the same rank as a, and move any beliefs ranked below a that can now be proved from higher ranked beliefs up to the highest rank at which they can be proved. This done, we have a new ranking, which is clearly an epistemic entrenchment in which a is ranked as required.

Note that the smallest set may be non-unique — there may be more than one such set. Given that there is no way to distinguish them, a skeptical policy would consider only the union of equally-ranked sets for removal. Another policy is to choose randomly between them. A third is to maintain all equally acceptable adjustments until further information is introduced which allows a distinction to be made. Only the skeptical policy is discussed in this paper.

Under the skeptical policy, when a search reveals that several sets tie for minimal capacity, the removal of any particular one of them is not allowed; the only set derived from them which could, in keeping with the policy, be removed is their union. Their union may, however, have a capacity larger than that of some other set whose removal is allowable under the policy. Thus, the tied minimal sets are removed from consideration, and the search is reiterated on the sets that remain — one of which is obviously the union of the excluded tied sets.

Operation of the algorithm is made more clear by the following example. It should, however, be noted that considerable searching among the elements of the power set of B is required, but has been suppressed.

EXAMPLE 34. Consider the geometric capacity (13) where $y = \frac{1}{2}$ acting on the entrenchment basis B =

$$\begin{array}{c} \neg\alpha \\ \epsilon \rightarrow \alpha \\ \beta \rightarrow \alpha, \gamma \rightarrow \beta, \gamma \\ \gamma \rightarrow \alpha, \delta \rightarrow \alpha \\ \delta \\ \kappa \rightarrow \epsilon, \kappa \end{array}$$

Where the beliefs on a given line are ranked higher than those on the line below. We begin by determining the sets which generate contradictions:

$$\{\neg\alpha, \beta \rightarrow \alpha, \gamma \rightarrow \beta, \gamma\}, \{\neg\alpha, \gamma, \gamma \rightarrow \alpha\}, \{\neg\alpha, \delta \rightarrow \alpha, \delta\},$$

$$\{\neg\alpha, \epsilon \rightarrow \alpha\ , \kappa \rightarrow \epsilon, \kappa\}.$$

Clearly we need to find a set of beliefs of minimal capacity containing at least one belief from each of the above sets. It is readily checked that there are two sets meeting this requirement ($\{\gamma, \delta, \kappa \to \epsilon\}, \{\gamma, \delta, \kappa\}$). Each has a capacity of $\frac{22}{178}$, so they are indistinguishable. Hence we can only consider their union, $\{\gamma, \delta, \kappa \to \epsilon, \kappa\}$ for removal. Comparing it to the next best set (one of $\{\gamma, \delta \to \alpha, \kappa \to \epsilon\}, \{\gamma, \delta \to \alpha, \kappa\}$, each with capacity $\frac{26}{178}$), we note that its capacity $\frac{24}{178}$ is now uniquely minimal amongst those sets left under consideration, so we remove it. The revised entrenchment basis is:

$$\begin{gathered}\neg\alpha \\ \epsilon \to \alpha \\ \beta \to \alpha, \gamma \to \beta \\ \gamma \to \alpha, \delta \to \alpha\end{gathered}$$

4. AGM REVISION FROM A CAPACITY

4.1

To simulate the AGM entrenchment construction [Gärdenfors and Makinson, 1988] of a theory via capacities, we show how to construct an appropriate capacity on an entrenchment basis. Revising the basis by minimising with respect to this capacity implements revision of the basis, yeilding a new basis whose theory closure is the AGM revision of the initial basis' theory closure [Williams, 1994b; Williams, 1995]. To acheive this we merely need to construct a capacity which throws away any and all beliefs which are ranked at or below the first point in the belief ranking at which a contradiction will follow from the incoming information. This is shown to be equivalent to base revision in [Williams, 1994b]. The appropriate capacity is the one defined as follows:

DEFINITION 35. Let $(X, \leq)$ be a belief ranking. Define a funtion $f : X \to \{0, 1\}$ by

$$f(x) := \begin{cases} 0 \text{ if } \{y \in X \mid x \leq y\} \Rightarrow \perp \\ 1 \text{ otherwise} \end{cases}$$

Then we define the *AGM-capacity* by

$$m_{AGM}(A) := \sum_{y \in A} f(y) / \sum_{y \in X} f(y).$$

LEMMA 36. *m_{AGM} is a capacity on $(X, \leq)$.*

Proof. a)

$$\begin{aligned} m_{AGM}(\emptyset) &= \sum_{y \in \emptyset} f(y) / \sum_{y \in X} f(y) \\ &= 0 \end{aligned}$$

b) is trivial as we are taking a sum over a smaller set, and c) is trivial because the capacity is defined additively. ■

THEOREM 37. *Let a be a belief and let $(T, \leq)$ be a finitely representible epistemic entrenchment. If $(X, \leq)$ is an entrenchment basis for $(T, \leq)$, then the theory closure of the AGM-capacity belief revision with a of $(X, \leq)$ is the same as the standard AGM revision (with a on T) based on the entrenchment construction.*

Proof. In view of the above discussion, it is sufficient to show that the beliefs removed from X are precisely those that are ranked at or below the first point in the ranking at which a contradiction follows from the inclusion of the belief to be incorporated.

Clearly if all beliefs x such that $m_{AGM}(x) = 0$ are removed, then the belief ranking will become consistent. No less than all of them will be removed because they are indistinguishable. On the other hand, no other beliefs will be selected for removal because doing so would necessarily increase the capacity of the removed set to at least 1. Consequently the set of such beliefs is exactly the set removed, as required. ■

5 MEASURE BASED TRANSMUTATION

It has been noted that, in some situations, the fact that a rank needs to be provided for each new belief can be a drawback to the methodology of standard belief revision. Another drawback is the fact that only revision with a single belief at a time is well-defined.

Capacities provide a method for avoiding the first of these difficulties using substitution capacities (32) to iterate the transmutation process on a belief ranking. At the n^{th} step, we progress from a belief ranking $(X_n, \leq_n)$ with a capacity m_n to a new belief ranking $(X_{n+1}, \leq_{n+1})$ and a capacity m_{n+1} which are all defined as in definition 32 with $R \subseteq X_n$ the subset that is minimal (in the sense dictated by the skeptical policy) with respect to m_n subject to $\neg a_n \notin Th(X_n \setminus R_n)$. It is clear that at each step we are left with a belief ranking and a capacity, and that we can iterate for as many beliefs as required.

Furthermore, if an entrenchment is to be maintained throughout the iteration, it should be clear that at each step we can use the current capacity to determine the smallest set of beliefs that needs to be moved up to the same ranking as a to maintain the entrenchment property (that is, the analagous procedure to that employed in 3).

Capacities also allow us to deal with the second problem in one of two ways: the first is to take the minimum of $m_n(R_n)$, over all permutations of the finite set A of beliefs to be inserted, as the revision to accept. The second is simply find the set $R \subseteq X$ that is minimal with respect to m subject to $Th(X \setminus R) \cap \{\neg a : a \in A\} = \emptyset$. This second solution, however, does not give a belief ranking on the new set, but it does allow us instead to avoid the iteration by revising with all the required beliefs at once.

6 DISCUSSION

This paper defines capacities, a family of functions that provide a general method for capturing the principle of minimal change. Some preliminary results concerning the behaviour of capacities were presented, and a capacity that can be used to implement a standard AGM revision based on the entrenchment construction [Gärdenfors and Makinson, 1988] was introduced. A capacity-based method of iterating transmutations on a belief ranking, or indeed, an entrenchment basis, to incorporate new beliefs in such a way that a rank for each new belief need not be provided was also outlined, as was a method for revising to accept a set of beliefs rather than an individual belief.

Since we have been concerned with transmutations on entrenchment bases, we have restricted ourselves to finite belief sets. However, it is clear that much of the theory extends to the case of an infinite belief set with an infinite number of distinct ranks provided that convergence of the appropriate sums is accounted for.

Future work will look at capacity transmutations that correspond to other existing strategies for nonmonotonic reasoning [Brewka, 1989] and belief revision [Nebel, 1989; Williams, 1994a; Williams, 1997a], will determine the relationship between the operators defined by capacity transmutations and the AGM rationality postulates [Alchourrón *et al.*, 1985], and will investigate topologies on the set of capacities of a belief ranking $(X, \leq)$ and attempt to identify its extreme points; that is, fundamental generating sets of capacities.

ACKNOWLEDGEMENTS

The authors wish to thank Dr. Williams for a number of very insightful discussions regarding the material, and the subject of Belief Revision in general.

The authors are also grateful to the referees for drawing to our attention a number of points which needed clarification; in particular, and oversight in the construction of 20. As a result of suggestions made, we feel that a considerable improvement in the exposition has been achieved.

Finally, the authors would like to acknowledge the contributions of Steven Kucera in the early genesis of the idea of a capacity.

Aidan Sims
School of Management, Newcastle University, Australia.

Brailey Sims
Department of Mathematics, Newcastle University, Australia.

REFERENCES

[Alchourrón *et al.*, 1985] C. Alchourrón, P. Gärdenfors and D. Makinson, On the logic of theory change: partial meet functions for contraction and revision, *Journal of Symbolic Logic*, **50**, 510–530, 1985.

[Brewka, 1989] G. Brewka. Preferred subtheories: an extended logical framework for default reasoning , In *Proceedings of the International Joint Conference on Artificial Intelligence*, pp. 1043–1048, 1989.

[Gärdenfors and Makinson, 1988] P. Gärdenfors and D. and Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, pp. 83–96, 1988.

[Nebel, 1989] B. Nebel. A knowledge level analysis of belief revision. In *Principles of Knowledge Representation and Reasoning: Proceedings of the First International Conference*, R. Brachman, H. Levesque and R. Reiter, eds. pp. 301–311. Morgan Kaufmann, San Mateo, CA, 1989.

[Nebel, 1994] B. Nebel. Base revision operations and schemes: semantics, representation, and complexity. In *Proceedings of the European Conference on Artificial Intelligence*, pp. 341–345. John Wiley and Sons, 1994.

[Spohn, 1988] W. Spohn. Ordinal conditional functions: a dynamic theory of epistemic states. In *Causation in Decision, Belief Change, and Statistics II*, W. L. Harper and B. Skyrms, eds. pp. 105–134. Kluwver Academic Publishers, 1988.

[Williams, 1994a] M. A. Williams. Transmutations of knowledge systems. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourth International Conference*, J. Doyle, E. Sandewall and P. Torasso, es. pp. 619–629. Morgan Kaufmann, San Mateo , CA, 1994.

[Williams, 1994b] M. A. Williams. On the logic of theory base change. In *Logics in Artificial Intelligence*, C. MacNish, D. Pearce and L. M. Pereira, eds. pp. 86–105. LNCS No 835, Springer Verlag, 1994.

[Williams, 1995] M. A. Williams. Iterated theory base change: a computational model. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, Montréal, pp. 1541–1550, 1995.

[Williams, 1997a] M. A. Williams. Anytime belief revision. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pp. 74–79. Morgan Kaufmann, 1997.

[Williams, 1997b] M. A. Williams. Implementing belief revision. In *Nonmonotonic Reasoning*, G. Antoniou, ed. MIT Press 1997.