Iterated Theory Base Change: A Computational Model

Mary-Anne Williams
Information Systems Group
Department of Management
University of Newcastle, NSW 2308
Australia

Abstract

The AGM paradigm is a formal approach to ideal and rational information change From a practical perspective it suffers from two shortcomings, the first involves difficulties with respect to the *finite representation* of information, and the second involves the lack of support for the iteration *of change operators* In this paper we show that these practical problems can be solved in theoretically satisfying ways wholely within the AGM paradigm

We introduce a *partial entrenchment ranking* which serves as a canonical representation for a theory base and a well-ranked episterruc entrenchment, and we provide a computational model for adjusting partial entrenchment rankings when they receive new information using a procedure based on the principle of minimal change

The connectioni between the standard AGM theory change operators and the theory base change operators developed herein suggest that the proposed computational model for iterated theory base change exhibits desirable behaviour

1 Introduction

Information systems can be used to represent a reasoning agent's view of the world Unless the agent has perfect and complete information it will require a mechanism to support the modification of its view as more information about the world is acquired Moreover, a computational model to support the acquisition of new information requires a finite representation of the information held by the agent This representation must capture the information content, the agent's committment to this information, and an encoding of how the information should change upon the intrusion of new information

The AGM paradigm was originally developed by Alchourron, Gardenfors and Makinson [1985] and has become one of the standard frameworks for information change It provides formal mechanisms for modeling the rational evolution of an ideal reasoner's view of the world In particular, it provides operators for modeling the revision and contraction of information Within the AGM paradigm the family of revision operators, and the family of contraction operators are circumscribed by sets of by rationality postulates The logical properties of a body of information are not strong enough to uniquely determine a revision or contraction operator, therefore the principal constructions for these operators rely on some form of underlying preference relation, such as a family of selection functions [Alchourron *et at* 1985], a system of spheres [Grove, 19881 a nice preorder on models [Katsuno and Mendelzon 1992, Peppas and Williams, 1995], or an epistemic entrenchment ordering [Gardenfors and Makinson, 1988]

From a theoretical point of view the AGM paradigm provides a very elegant and simple mechanism for rational and ideal change However, from a practical perspective its operators are insufficient because they essentially take a preference relation together with a sentence and produce a theory In other words, the underlying preference relation is lost This property is attractive in a theoretical context because it allows the resultant theory to adopt any preference relation depending on the desired dynamic behaviour In practice, however, a *policy for change* will be necessary A straightforward method of specifying such a policy is to impose constraints on the underlying preference relation which will in turn determine the dynamic behaviour of the system As noted above the AGM postulates do not uniquely determine revision and contraction operators, however Gardenfors and Makinson [1986] showed that an epistemic entrenchment ordering does According to Rott [1991a] what should be at focus is not theory revision but epistemic entrenchment revision

The underlying preference relation fully characterizes a information system's content, its committment to the information, and its desired dynamic properties We refer to the process of changing an information system's underlying preference relation as a *transmutation*

Williams [1994a] showed that allowing the *principle of minimal change* to command the policy for change results to two different forms of transmutations, conditionalization and adjustment Conditionalization was introduced by Spohn [1988] and is based on a *relative* measure of minimal change, on the other hand adjustment described by Williams [1994a] is based on an *absolute* measure of minimal change Adjustments are compared and contrasted with conditionalization in [Williams 1994a]

In our current *theory base* context we represent the underlying preference relation for a theory base as a *partial entrenchment ranking* which serves as a canonical representation of a well-ranked epistemic entrenchment We show it can be used to support transmutations of a theory base, and thus *iterated theory base change* Furthermore, we provide a computational model which modifies a partial entrenchment ranking using an absolute measure of minimal change Any theorem prover can be used to realize this model Consequently, we present a practical solution to the iterated theory base change problem, we note however that a full complexity analysis of our proposed procedure is yet to be conducted

We briefly outline some technical preliminaries and the AGM paradigm in section 2, this is a standard treatise where two extra postulates introduced in [Williams 1994a] are presented, the reader familiar with this work may skip to the next section In section 3, we describe the representation of theory bases using partial entrenchment rankings, and we demonstrate how they are related to epistemic entrenchment orderings In section 4 we describe iterated theory base change using transmutations of partial entrenchment rankings In section 5 we describe a simple computational model that implements a transmutation of a partial entrenchment ranking The transmutation employed is an adjustment hence the proposed computational model uses an absolute minimal measure of change In section 6 we explore the connection between our theory base change operators using an adjustment and the standard AGM theory change operators, and we argue that the established relationships demonstrate that our proposed theory base change operators exhibit desirable behaviour In particular they maintain as much of the theory base as would be retained in the corresponding theory change, thus propagating as much *explicit* information as possible Related work is discussed in section 7, and a summary of our results is given section 8

2 The AGM Paradigm

Let $\mathcal{L}$ denote a countable language which is closed under a complete set of Boolean connectives We will denote sentences in $\mathcal{L}$ by lower case Greek letters We assume $\mathcal{L}$ is governed by a logic that is identified with its consequence relation $\vdash$ which is assumed to satisfy the following conditions [Gardenfors 1988] (a) If α is a truth-functional tautology, then $\vdash \alpha$, (b) If $\vdash \alpha \rightarrow \beta$ and $\vdash \alpha$, then $\vdash \beta$, (c) $\vdash$ is consistent, that is, $\nvdash \perp$, where $\perp$ denotes the inconsistent theory, (d) $\vdash$ satisfies the deduction theorem, and (e) $\vdash$ is compact

The set of all logical consequences of a set $\Gamma \subseteq \mathcal{L}$, that is $\{\alpha \quad \Gamma \vdash \alpha\}$, is denoted by $\mathrm{Cn}(\Gamma)$ A *theory* of $\mathcal{L}$ is any subset of $\mathcal{L}$, closed under Cn Let $\mathcal{K}_{\mathcal{L}}$ denote the *set of all theories of* $\mathcal{L}$ For a theory T, if $T = \mathrm{Cn}(\Gamma)$, then we refer to Γ a *theory base* for T We denote the set of all consistent nontautological sentences in $\mathcal{L}$, that is, the set of contingent sentences, by $\mathcal{L}^{\bowtie}$ Finally, a *well ranked* preorder on a set Γ is a preorder such that every nonempty subset of Γ has a minimal member We denote the domain of a function $\mathbf{B}$, and the range of $\mathbf{B}$ by $\mathrm{dom}(\mathbf{B})$ and $\mathrm{range}(\mathbf{B})$, respectively

2 1 The Postulates

Within the AGM paradigm a body of information is represented as a theory, and changes are regarded as transformations on theories There are three principal types of AGM transformations, expansion, contraction and revision We note that Peppas and Williams [1995] have shown that a form of *update* [Katsuno and Mendelzon, 1992] can be incorporated into the AGM paradigm These transformations allow us to model changes of information based on the *principle of minimal change* [Gardenfors, 1988] Expansion is the simplest change, it models the acquiescence of information without the removal of any inconsistent information More formally, the *expansion* of a theory T with respect to a sentence α, denoted as T^+_α, is defined to be the logical closure of T and α, that is $T^+_\alpha = \mathrm{Cn}(T \cup \{\alpha\})$

A *contraction* of T with respect to α, T^-_α, involves the removal of a set of sentences from T so that α is no longer implied Formally, a *well-behaved contraction operator* $^-$ *is* any function from $\mathcal{K}_{\mathcal{L}} \times \mathcal{L}$ to $\mathcal{K}_{\mathcal{L}}$, mapping (T, α) to T^-_α which satisfies the postulates ($^-$1) – ($^-$9), below, and a *very well-behaved contraction operator* $^-$ satisfies the postulates ($^-$1) – ($^-$10)

($^-$1) For any $\alpha \in \mathcal{L}$ and any $T \in \mathcal{K}_{\mathcal{L}}$, $T^-_\alpha \in \mathcal{K}_{\mathcal{L}}$
($^-$2) $T^-_\alpha \subseteq T$
($^-$3) If $\alpha \notin T$ then $T^-_\alpha = T$
($^-$4) If $\nvdash \alpha$ then $\alpha \notin T^-_\alpha$
($^-$5) $T \subseteq (T^-_\alpha)^+_\alpha$
($^-$6) If $\vdash \alpha \equiv \beta$ then $T^-_\alpha = T^-_\beta$
($^-$7) $T^-_\alpha \cap T^-_\beta \subseteq T^-_{\alpha\wedge\beta}$
($^-$8) If $\alpha \notin T^-_{\alpha\wedge\beta}$ then $T^-_{\alpha\wedge\beta} \subseteq T^-_\alpha$
($^-$9) For every nonempty set Γ of nontautological sentences, there exists a sentence $\alpha \in \Gamma$ such that $\alpha \notin T^-_{\alpha\wedge\beta}$ for every $\beta \in \Gamma$
($^-$10) For every nonempty set Γ of nontautological sentences, there exists a sentence $\alpha \in \Gamma$ such that $\beta \notin T^-_{\alpha\wedge\beta}$ for every $\beta \in \Gamma$

A *revision* attempts to transform a theory in order to incorporate a sentence so that the resultant theory is consistent Formally, a *well-behaved revision operator* * is any function from $\mathcal{K}_{\mathcal{L}} \times \mathcal{L}$ to $\mathcal{K}_{\mathcal{L}}$, mapping (T, α) to T^*_α which satisfies the postulates (*1) – (*9), below, and a *very well-behaved revision operator* * satisfies the postulates (*1) – (*10)

(*1) For any $\alpha \in \mathcal{L}$ and any $T \in \mathcal{K}_{\mathcal{L}}$, $T^*_\alpha \in \mathcal{K}_{\mathcal{L}}$
(*2) $\alpha \in T^*_\alpha$
(*3) $T^*_\alpha \subseteq T^+_\alpha$
(*4) If $\neg\alpha \notin T$ then $T^+_\alpha \subseteq T^*_\alpha$
(*5) $T^*_\alpha = \perp$ if and only if $\vdash \neg\alpha$
(*6) If $\vdash \alpha \equiv \beta$ then $T^*_\alpha = T^*_\beta$
(*7) $T^*_{\alpha\wedge\beta} \subseteq (T^*_\alpha)^+_\beta$
(*8) If $\neg\beta \notin T^*_\alpha$ then $(T^*_\alpha)^+_\beta \subseteq T^*_{\alpha\wedge\beta}$
(*9) For every nonempty set Γ of nontautological sentences, there exists a sentence $\alpha \in \Gamma$ such that $\alpha \notin T^*_{\neg\alpha\vee\neg\beta}$ for every $\beta \in \Gamma$
(*10) For every nonempty set Γ of nontautological sentences, there exists a sentence $\alpha \in \Gamma$ such that $\beta \notin T^*_{\neg\alpha\vee\neg\beta}$ for every $\beta \in \Gamma$

2 2 Epistemic Entrenchment

An epistemic entrenchment [Garden fors and Makinson, 1988] is an ordering or the sentences in *C* which attempts to capture the relative importance of information in the face of change In order to determine a unique revision or contraction operation a theory is endowed with an epistemic entrenchment ordering, which can be used to determine the information to be retracted, retained, and acquired during the contraction of old information and the acceptance of new information

Definition *Given a theory T of $\mathcal{L}$, an* epistemic entrenchment related to T *is any binary relation $\leq$ on $\mathcal{L}$ satisfying* (EE1) – (EE5), *below*

(EE1) *If $\alpha \leq \beta$ and $\beta \leq \gamma$, then $\alpha \leq \gamma$*
(EE2) *For all α, $\beta \in \mathcal{L}$, if $\alpha \vdash \beta$ then $\alpha \leq \beta$*
(EE3) *For all α, $\beta \in \mathcal{L}$, $\alpha \leq \alpha \wedge \beta$ or $\beta \leq \alpha \wedge \beta$*
(EE4) *When $T \neq \perp$, $\alpha \notin T$ if and only if $\alpha \leq \beta$ for all $\beta \in \mathcal{L}$*
(EE5) *If $\beta \leq \alpha$ for all $\beta \in \mathcal{L}$, then $\vdash \alpha$*

We define $\alpha < \beta$, as $\alpha \leq \beta$ and not $\beta \leq \alpha$ An epistemic entrenchment that possesses a finite number of natural partitions is *finite*

Gardenfors and Makinson [1988] showed that, for every contraction operator $^-$ there exists an epistemic entrenchment $\leq$ related to T such that the condition (E^-) below, is true for every $\alpha \in \mathcal{L}$, and conversely

$$(E^-) \quad T^-_\alpha = \begin{cases} \{\beta \in T \quad \alpha < \alpha \vee \beta\} & \text{if } \nvdash \alpha \\ T & \text{otherwise} \end{cases}$$

Their representation result was extended by Williams [1994c] to well-behaved and very well-behaved operators in the following theorem

Theorem 1 *Let T be a theory of $\mathcal{L}$ For every well behaved contraction (very well-behaved contraction) operator $^-$ for T there exists a well-ranked epistemic entrenchment (finite epistemic entrenchment) $\leq$ related to T such that (E^-) is true for every $\alpha \in \mathcal{L}$, and conversely*

From the work of Gardenfors and Makinson [1988] and the Levi Identity it is straightforward to derive a similar representation result for revision using (E^*) below

$$(E^*) \quad T^*_\alpha = \begin{cases} \{\beta \in \mathcal{L} \quad \neg\alpha < \neg\alpha \vee \beta\} & \text{if } \nvdash \neg\alpha \\ \bot & \text{otherwise} \end{cases}$$

Peppas [1993] was the first to identify the class of well-behaved revision operators and he obtains an analogous representation result for well behaved revision based on a well-ordered system of spheres [Grove, 1988]

3 Representing Theory Bases

Identical canonical representations or epistemic entrenchment ordenngs were developed independently by Rott [1991a] and Williams [1992, 1994c] Rott's specification IS an E-Base and Williams' an ensconcement

An E-Base is a preference relation that provides a canonical representation of an epistemic entrenchment and a theory base Williams showed that an E-Base is capable of uniquely determining theory base change operators, as well as theory change operators However, the theory base operations provided do not take an E-Base to an E-Base In other words, the preference relation on the theory base is not propagated during the process of change, and hence iteration is not supported

Williams [1994a] extended the work of Spohn [1988], based on observations in [Gardenfore, 1988], to arbitrary iterated revision operators for theories, and referred to the process of changing an epistemic entrenchment ordering as a *transmutation*

In this section we demonstrate that a *partial entrenchment ranking* which is a ranking of a theory base can specify a well ranked (or finite) epistemic entrenchment ordering on a theory

A partial entrenchment ranking formally defined below, maps a set of sentences in *C* to ordinals Intuitively, the higher the ordinal assigned to a sentence the more firmly held it is Throughout the remainder of the paper it will be understood that *0* is an ordinal chosen to be sufficiently large for the purpose of the discussion

Definition *A* partial entrenchment ranking *is a function* **B** *from a subset of sentences in $\mathcal{L}$ into $\mathcal{O}$ such that the following conditions are satisfied for all* $\alpha \in \text{dom}(\mathbf{B})$

(PER1) $\{\beta \in \text{dom}(\mathbf{B}) \quad \mathbf{B}(\alpha) < \mathbf{B}(\beta)\} \nvdash \alpha$,
(PER2) *If* $\vdash \neg\alpha$, *then* $\mathbf{B}(\alpha) = 0$
(PER3) $\mathbf{B}(\alpha) = \mathcal{O}$ *if and only if* $\vdash \alpha$

(PER1) says that the sentences which are assigned ordinals higher than an arbitrary sentence α, do not entail α, and (PER2) says inconsistent sentences are assigned zero (PER3) says that tautologies in the domain of **B** are assigned the largest ordinal, this is not unduly restrictive as the language is countable, hence the cardinality of the range of **B** will always be bounded If Γ is a nonempty set of sentences, then we define $\mathbf{B}(\Gamma) = \min(\{\mathbf{B}(\alpha) \quad \alpha \in \Gamma\})$

As noted in numerous other works (see [Gardenfors and Makinson, 1994]) the ordinal assignment can be viewed in two distinct ways (i) *qualitatively*, where the relative ordering of sentences is used, or (ii) *quantitatively*, where ordinals have some extra meaning, and a calculus based on their numerical value adopted

We refer to a partial entrenchment ranking with a finite range as a *finite partial entrenchment ranking* A finite ranking does *not* imply that the language is finite We denote the family of all partial entrenchment rankings by $\mathcal{B}$

The intended interpretation of a partial entrenchment ranking is that sentences mapped to ordinals greater than zero represent the *explicit* beliefs of the reasoning agent, and their logical closure represents the agent's set of *implicit* beliefs

Definition *Define the* explicit information set *represented by* $\mathbf{B} \in \mathcal{B}$ *to be* $\{\alpha \in \text{dom}(\mathbf{B}) \quad \mathbf{B}(\alpha) > 0\}$, *and denote it by* $\exp(\mathbf{B})$ *Similarly, define the* implicit information set *represented by* $\mathbf{B} \in \mathcal{B}$ *to be* $\text{Cn}(\exp(\mathbf{B}))$, *and denote it by* $\text{imp}(\mathbf{B})$ *A sentence α is an* explicit belief *if and only if* $\alpha \in \exp(\mathbf{B})$

We now describe the relationship between a partial entrenchment ranking and an epistemic entrenchment ordering First we define the notion of a cut which has been used by numerous authors under various guises, most notably by Rott [1991a] and Gardenfors and Makinson [1994]

Definition *For a partial entrenchment ranking,* **B**, *and a sentence* $\alpha \in \text{imp}(\mathbf{B})$ *we define the* α cut *of* **B** *as* $\{\beta \in \exp(\mathbf{B}) \quad \{\gamma \in \exp(\mathbf{B}) \quad \mathbf{B}(\beta) < \mathbf{B}(\gamma)\} \nvdash \alpha\}$, *and denote it by* $\text{cut}(\mathbf{B}, \alpha)$

Cuts are always subsets of $\exp(\mathbf{B})$, and for contingent $\alpha \in \exp(\mathbf{B})$ the α cut of **B** contains all sentences in $\exp(\mathbf{B})$ that are assigned ordinals at least as large as $\mathbf{B}(\alpha)$, thus if $\alpha \in \exp(\mathbf{B})$, then α is in its own cut Note, if $\vdash \alpha$ then $\text{cut}(\mathbf{B}, \alpha) = \emptyset$

The following theorem provides an explicit construction of an epistemic entrenchment ordering from a partial entrenchment ranking based on cuts

Theorem 2 *For* $\mathbf{B} \in \mathcal{B}$ *and* $\alpha, \beta \in \mathcal{L}$, *define* $\leq_\mathbf{B}$ *to be given by* $\alpha \leq_\mathbf{B} \beta$ *if and only if either (i)* $\alpha \notin \text{imp}(\mathbf{B})$, *or (ii)* $\text{cut}(\mathbf{B}, \beta) \subseteq \text{cut}(\mathbf{B}, \alpha)$ *Then* $\leq_\mathbf{B}$ *is a well ranked epistemic entrenchment ordering related to* $\text{imp}(\mathbf{B})$

We refer to $\leq_\mathbf{B}$ as *the epistemic entrenchment ordering generated from* **B** From Theorem 2 we see that the tautologies are maximal, and sentences not in $\text{imp}(\mathbf{B})$ are minimal with respect to $\leq_\mathbf{B}$ If **B** is finite then the epistemic entrenchment ordering it generates is finite An epistemic entrenchment ordering is *finitely representable* if and only if every cut is finitely axiomatizable Consequently, $\leq$ is finitely representable if and only if there exists a **B** such that $\leq = \leq_\mathbf{B}$ and $\exp(\mathbf{B})$ is finite

Intuitively, Theorem 3, below, says given a well-ranked epistemic entrenchment ordering there exists a partial entrenchment ranking that can be used to generate it Theorems 2 and 3 tell us that a partial entrenchment ranking can be considered to be a (possibly minimal) specification of a well-ranked epistemic entrenchment ordering

Theorem 3 *Let T be a theory in $\mathcal{L}$, and $\Gamma \subseteq T$ Let $\leq$ be a well-ranked epistemic entrenchment related to T Let $T_0, T_1, T_2,$ be the natural partitions of $\leq$ indexed* *so that $\alpha \in T_i$ and $\beta \in T_j$ whenever $i < j$ implies $\alpha < \beta$. For all $\alpha \in T$, if $\{\beta \in \Gamma : \alpha \leq \beta\} \vdash \alpha$, then the mapping of each sentence in Γ from the natural partition T_i to i for all i is a partial entrenchment ranking, $\mathbf{B}$, that generates $\leq$, that is, $\leq_{\mathbf{B}} = \leq$*

The T referred to in the theorem above is the explicit information set of the partial entrenchment ranking that generates the epistemic entrenchment and while obviously not unique, it must contain enough sentences of 'the right stuff' in each natural partition of the epistemic entrenchment ordering to enable its regeneration In particular, it must satisfy the condition in the theorem, at least one such *T* exists, namely *T* itself

Stepwise constructions of epistemic entrenchment orderings from partial entrenchment rankings, and vice versa, can be derived from the those provided in [Rott, 1991b, Williams, 1994c]

Definition *We define, the* degree of acceptance *of $\alpha \in \mathcal{L}$ with respect to $\mathbf{B} \in \mathcal{B}$ to be,* $\text{degree}(\mathbf{B}, \alpha) = \mathbf{B}(\text{cut}(\mathbf{B}, \alpha))$ *whenever $\not\vdash \alpha$ and $\alpha \in \text{imp}(\mathbf{B})$. If $\vdash \alpha$ then $\text{degree}(\mathbf{B}, \alpha) = \mathcal{O}$, and if $\alpha \notin \text{imp}(\mathbf{B})$ then* $\text{degree}(\mathbf{B}, \alpha) = 0$

Note, $\text{degree}(\mathbf{B}, \alpha) \leq \text{degree}(\mathbf{B}, \beta)$ if and only if $\alpha \leq_{\mathbf{B}} \beta$

EXAMPLE Let $\mathbf{B}$ be given by $\mathbf{B}(\alpha \rightarrow \beta) = 3$, $\mathbf{B}(\gamma) = 2$, $\mathbf{B}(\alpha) = 1$ Then we have $\text{degree}(\mathbf{B}, \beta) = 1$, $\text{degree}(\mathbf{B}, \alpha \rightarrow \gamma) = 2$, $\text{degree}(\mathbf{B}, \neg\alpha) = 0$, $\text{degree}(\mathbf{B}, \neg\beta) = 0$, and $\text{degree}(\mathbf{B}, \neg\gamma) = 0$

4 Iterated Change for Theory Bases

Recall from our previous discussion that from a practical perspective the AGM paradigm does not provide a policy to support the iteration of its change operations In this section we show how transmutations of partial entrenchment rankings can be used to support iterated theory base change, and we provide existence theorems that demonstrate how transmutations are related to the AGM theory revision and theory contraction operators

For revision and contraction operators the information input is a sentence We now define a *transmutation of partial entrenchment rankings* where the information input is a contingent sentence a and an ordinal i The interpretation [Gardenfore 1986] of this is that the sentence a is the information to be accepted, and i is the degree of firmness with which it is to be incorporated into the transmuted partial entrenchment ranking

Definition *We define a* transmutation schema *for partial entrenchment rankings, *, to be an operator from $\mathcal{B} \times \mathcal{L}^{\bowtie} \times \mathcal{O}$ to $\mathcal{B}$, such that $(\mathbf{B}, \alpha, i) \mapsto \mathbf{B}^\star(\alpha, i)$ which satisfies (i) $\mathbf{B}^\star(\alpha, i)(\alpha) = i$, and (ii) $\text{imp}(\mathbf{B}^\star(\alpha, i))$*

$$= \begin{cases} \text{Cn}(\{\beta \in \exp(\mathbf{B}) : \text{degree}(\mathbf{B}, \neg\alpha) < \mathbf{B}(\beta)\} \cup \{\alpha\}) & \textit{if } i > 0 \\ \text{Cn}(\{\beta \in \exp(\mathbf{B}) : \{\gamma \in \exp(\mathbf{B}) : \text{degree}(\mathbf{B}, \alpha) < \mathbf{B}(\gamma)\} \cup \{\neg\alpha\} \vdash \beta\}) & \textit{otherwise} \end{cases}$$

We say $\mathbf{B}^\star(\alpha, i)$ is an (α, i)–transmutation of $\mathbf{B}$, and we note that a transmutation is not defined for inconsistent, or tautological sentences An inconsistent sentence is not considered acceptable information [Spohn 1988], and the exclusion of tautological sentences is for technical convenience allowing us to focus on the principal case, however the definition can easily be extended to include both tautological and inconsistent sentences Intuitively, if $i > 0$ then $\mathbf{B}^\star(\alpha, i)$ embodies a theory base revision, and similarly $\mathbf{B}^\star(\alpha, 0)$ represents a theory base contraction It is interesting to note that $\text{imp}(\mathbf{B}) \mapsto \text{imp}(\mathbf{B}^\star(\alpha, 0))$ defines a *withdrawal function* [Makinson, 1987]

If $\text{eks}(\mathbf{B})$ is a theory then the transmutation schema for partial entrenchment rankings is precisely the *same* as the transmutation schema for ordinal epistemic entrenchment functions which is given in [Williams 1994a] for the theory change setting Theorems 4 and 5, below, demonstrate how transmutations of partial entrenchment rankings are related to theory revision and theory contraction operators within the AGM paradigm They are existence theorems, we provide explicit constructions based on a particular form of transmutation in section 6

Theorem 4 *Let T be a theory in $\mathcal{L}$ For every (very) well-behaved theory revision operator * for T, there exists a (finite) $\mathbf{B} \in \mathcal{B}$ such that $\text{imp}(\mathbf{B}) = T$ and $T^*_\alpha = \text{imp}(\mathbf{B}^\star(\alpha, i))$ for every $\alpha \in \mathcal{L}^{\bowtie}$ and $0 < i < \mathcal{O}$ Conversely, for every (finite) $\mathbf{B} \in \mathcal{B}$, there exists a (very) well-behaved theory revision operator * for $\text{imp}(\mathbf{B})$ such that $(\text{imp}(\mathbf{B}))^*_\alpha = \text{imp}(\mathbf{B}^\star(\alpha, i))$ is true for every $\alpha \in \mathcal{L}^{\bowtie}$ and $0 < i < \mathcal{O}$*

Theorem 5 *Let T be a theory in $\mathcal{L}$ For every (very) well-behaved theory contraction operator $^-$ for T, there exists a (finite) $\mathbf{B} \in \mathcal{B}$ such that $\text{imp}(\mathbf{B}) = T$, and $T^-_\alpha = \text{Cn}(\exp(\mathbf{B}^\star(\alpha, 0)) \cup \{\neg\alpha\}) \cap T)$ for every $\alpha \in \mathcal{L}^{\bowtie}$ Conversely, for every (finite) $\mathbf{B} \in \mathcal{B}$, there exists a (very) well-behaved theory contraction operator $^-$ for $\text{imp}(\mathbf{B})$ such that $(\text{imp}(\mathbf{B}))^-_\alpha = \text{Cn}(\exp(\mathbf{B}^\star(\alpha, 0)) \cup \{\neg\alpha\}) \cap \text{imp}(\mathbf{B})$ for every $\alpha \in \mathcal{L}^{\bowtie}$*

The definition of a transmutation schema forces the contraction of a theory base to be a subset of the original theory base, that is, for all $\alpha \in \mathcal{L}^{\bowtie}$, $\exp(\mathbf{B}^\star(\alpha, 0)) \subseteq \exp(\mathbf{B})$ Consequently, the recovery postulate ($^-5$) is not necessarily satisfied [Makinson, 1987] If it is desirable to satisfy recovery then certain sentences identified by Nebel [1991] and Foo can be added to the partial entrenchment ranking in a contraction to ensure its satisfaction For a more elaborate discussion of this issue see Williams [1994c] We note that the definitions and theorems presented in this paper can be modified to capture recovery in a straightforward manner We adopt the definition as given because it is more consistent with our interpretation of a partial entrenchment ranking being an *explicit* representation of an agent's information

5 A Computational Model

In this section we describe a particular type of transmutation called an adjustment which uses a *policy for change* based on an *absolute* minima*l* *measure* In particular, it involves the absolute minimal change of a partial entrenchment ranking that is required to incorporate the desired new information A semantic account of adjustments can be found in [Williams, 1994 a]

We present a computational model for adjustments which can be used as the basis for a computer-based implementation Any theorem prover can be used to realize it The model itself is stated in its most perspicuous form, and can be optimised in several obvious ways, see [Williams, 1993] for details For the purpose of our model we focus on finite epistemic rankings, and the following theorem describes the degree of acceptance of sentences with respect to finite rankings

Theorem 6 *Let a be a nontaulologtcal sentence If* B€B is *finite, then* degree(B, a)

$$= \begin{cases} \textit{largest } j \textit{ such that } \{\beta \in \exp(\mathbf{B}) \ \ \mathbf{B}(\beta) \geq j\} \vdash \alpha \\ \qquad\qquad\qquad\qquad\qquad\qquad \textit{if } \alpha \in \operatorname{imp}(\mathbf{B}) \\ 0 \qquad\qquad\qquad\qquad\qquad\quad \textit{otherwise} \end{cases}$$

Baaed on Theorem 6 the following function uaes a ▪imple procedural algorithm to determine the degree of a sentence with respect to a finite partial entrenchment ranking The function takes two input parameters, namely a finite partial entrenchment ranking, B, and a nontautological sentence, a, it calculates and returns the degree of acceptance of a in D The algorithm requires the support of your favourite theorem prover to implement the logical implication relation, I-

```
FUNCTION degree(B,a)
degree(B,a) ← max(range(B)) + 1
do degree(B,a) ← degree(B,a) - 1
until degree(B,a) ≤ 0 or
        {b ∈ dom(B)  B(b) ≥ degree(B,a)} ⊢ a
```

This algorithm can compute the generated epistemic entrenchment ordering of sentences based on the information encoded in a partial entrenchment ranking

A transmutation which is suitable for theory base change is an adjustment defined in the theorem below When D is finite then its definition constitutes a computational model which, loosely speaking, involves successive calls to the function degree for each sentence in the domain of D

Theorem 7 *Let* $\mathbf{B} \in \mathcal{B}$, $\alpha \in \mathcal{L}^{\bowtie}$, *and* $i < \mathcal{O}$ *Then* $\mathbf{B}^\star(\alpha, i)$ *defined below is a transmutation*

$$\mathbf{B}^\star(\alpha, i) = \begin{cases} (\mathbf{B}^-(\alpha, i)) & \textit{if } i < \operatorname{degree}(\mathbf{B}, \alpha) \\ (\mathbf{B}^-(\neg\alpha, 0))^+(\alpha, i) & \textit{otherwise} \end{cases}$$

where $\mathbf{B}^-(\alpha, i)(\beta)$

$$= \begin{cases} i & \textit{if } \operatorname{degree}(\mathbf{B}, \alpha) = \operatorname{degree}(\mathbf{B}, \alpha \vee \beta) \textit{ and } \mathbf{B}(\beta) > i \\ \mathbf{B}(\beta) & \textit{otherwise} \end{cases}$$

for all $\beta \in \operatorname{dom}(\mathbf{B})$, *and* $\mathbf{B}^+(\alpha, i)(\beta)$

$$= \begin{cases} \mathbf{B}(\beta) & \textit{if } \mathbf{B}(\beta) > i \\ i & \textit{if } \alpha \equiv \beta \textit{ or } \mathbf{B}(\beta) \leq i < \operatorname{degree}(\mathbf{B}, \neg\alpha \vee \beta) \\ \operatorname{degree}(\mathbf{B}, \neg\alpha \vee \beta) & \textit{otherwise} \end{cases}$$

for all $\beta \in \operatorname{dom}(\mathbf{B} \cup \{\alpha\})$

Hans Rott has pointed out that the $\mathbf{B}^+(\alpha, i)$ operation can be traced back to Rescher [1976] Intuitively, an (α, i)-adjustment of $\mathbf{B}$ involves minimal changes to $\mathbf{B}$ such that the sentence α is accepted with degree i In particular, each sentence $\beta \in \operatorname{dom}(\mathbf{B})$ is reassigned an ordinal closest to $\mathbf{B}(\beta)$ in the adjusted partial entrenchment ranking $\mathbf{B}^\star(\alpha, i)$ under the guiding principle that if we reduce the degree of an accepted sentence α to i, say, then we also reduce the degree of each sentence that would be retracted in α's contraction to i as well

The following theorem illustrates the interrelationships between theory base revision and theory base contraction based on adjustments In particular, Theorem 8(J) IS analogous to the Harper Identity and it captures the dependence of theory base contraction on the information content of the theory base, that is, the explicit information set, exp(B) Similarly Theorem 8(11) is analogous to the Levi identity, and 8(111) says that a Levi Identity also exists at the deeper preference relation level In particular, precisely the *same* partial entrenchment ranking is obtained when a partial entrenchment ranking is adjusted to accept new information *a* with firmness 0 < 1 < *O,* as when we adjust the partial entrenchment ranking to remove a, and then adjust the resultant partial entrenchment ranking to accept new information α with firmness i Similarly, 8(v) establishes a Harper Identity at the preference relation level where the functions, $\mathbf{B}$ and $\mathbf{B}^\star$, have been identified with the set of ordered pairs $(\alpha, \mathbf{B}(\alpha))$ and $(\alpha, \mathbf{B}^\star(\alpha))$, respectively, such that $\alpha \in \operatorname{dom}(\mathbf{B})$

Theorem 8 *Let* $\mathbf{B} \in \mathcal{B}$, *let the transmutation schema* * *be an adjustment, and let* $0 < i < \mathcal{O}$ *Then*
(i) $\exp(\mathbf{B}^\star(\alpha, 0)) = \exp(\mathbf{B}^\star(\neg\alpha, i)) \cap \exp(\mathbf{B})$
(ii) $\exp(\mathbf{B}^\star(\alpha, i)) = \exp(\mathbf{B}^\star(\neg\alpha, 0)) \cup \{\alpha\}$,
(iii) $\mathbf{B}^\star(\alpha, i) = (\mathbf{B}^\star(\neg\alpha, 0))^\star(\alpha, i)$, *and*
(iv) $\exp(\mathbf{B}^\star(\alpha, 0)) = \exp(\mathbf{B}^\star(\alpha, 0) \cap \mathbf{B})$
(v) $\mathbf{B}^\star(\alpha, 0) \cap \mathbf{B} = \mathbf{B}^\star(\neg\alpha, \max(1, \operatorname{degree}(\mathbf{B}, \neg\alpha))) \cap \mathbf{B}$

We illustrate the adjustment of a theory base in the example below

EXAMPLE Let $\mathbf{B}$ be given by $\mathbf{B}(\alpha \to \beta) = 3$, $\mathbf{B}(\gamma) = 2$, $\mathbf{B}(\alpha) = 1$

(a) Consider the reorganization of $\mathbf{B}$ where more compelling evidence for α is acquired, and we decide to increase the degree of acceptance of α from 1 to 3, that is, $\mathbf{B}^\star(\alpha, 3)$ Then we have $\mathbf{B}^\star(\alpha, 3)(\alpha) = 3$, $\mathbf{B}^\star(\alpha, 3)(\alpha \to \beta) = 3$, $\mathbf{B}^\star(\alpha, 3)(\gamma) = 2$, $\operatorname{degree}(\mathbf{B}^\star(\alpha, 3), \beta) = 3$, $\operatorname{degree}(\mathbf{B}^\star(\alpha, 3), \alpha \to \gamma) = 2$, $\operatorname{degree}(\mathbf{B}^\star(\alpha, 3), \neg\alpha) = 0$, $\operatorname{degree}(\mathbf{B}^\star(\alpha, 3), \neg\beta) = 0$, and $\operatorname{degree}(\mathbf{B}^\star(\alpha, 3), \neg\gamma) = 0$

(b) Now let's consider the contraction of α, that is, $\mathbf{B}^\star(\alpha, 0)$ Then we have $\mathbf{B}^\star(\alpha, 0)(\alpha) = 0$, $\mathbf{B}^\star(\alpha, 0)(\alpha \to \beta) = 3$, $\mathbf{B}^\star(\alpha, 0)(\gamma) = 2$, $\operatorname{degree}(\mathbf{B}^\star(\alpha, 0), \beta) = 0$, $\operatorname{degree}(\mathbf{B}^\star(\alpha, 0), \alpha \to \gamma) = 2$, $\operatorname{degree}(\mathbf{B}^\star(\alpha, 0), \neg\alpha) = 0$, $\operatorname{degree}(\mathbf{B}^\star(\alpha, 0), \neg\beta) = 0$, and $\operatorname{degree}(\mathbf{B}^\star(\alpha, 0), \neg\gamma) = 0$

(c) Finally, we consider the acceptance of $\neg\alpha$ with degree 4, that is, $\mathbf{B}^\star(\neg\alpha, 4)$ Then we have $\mathbf{B}^\star(\neg\alpha, 1)(\neg\alpha) = 4$, $\mathbf{B}^\star(\neg\alpha, 1)(\alpha \to \beta) = 4$, $\operatorname{degree}(\mathbf{B}^\star(\neg\alpha, 1), \gamma) = 2$, $\operatorname{degree}(\mathbf{B}^\star(\neg\alpha, 1), \beta) = 0$, $\operatorname{degree}(\mathbf{B}^\star(\neg\alpha, 1), \alpha \to \gamma) = 4$, $\operatorname{degree}(\mathbf{B}^\star(\neg\alpha, 1), \alpha) = 0$ $\operatorname{degree}(\mathbf{B}^\star(\neg\alpha, 1), \neg\beta) = 0$, $\operatorname{degree}(\mathbf{B}^\star(\neg\alpha, 1), \neg\gamma) = 0$

Adjustments use the relative ordering encoded in a partial entrenchment ranking, and they *preserve finiteness* with respect to both revisions and contractions, adjusting a finite partial entrenchment ranking results in a finite ranking, and if $\exp(\mathbf{B})$ is finite then $\exp(\mathbf{B}^\star(\alpha, i))$ is finite

6 Connections with Theory Change

The theorems in this section establish the explicit relationship between *theory base change* operators based on adjustments and *theory change* operators based on the generated epistemic entrenchment ordering

Theorem 9 *Let* $\mathbf{B} \in \mathcal{B}$, *let the transmutation schema* * *be an adjustment, and let* $0 < i < \mathcal{O}$ *Let* $^-$ *be the contraction operator for* $\operatorname{imp}(\mathbf{B})$ *uniquely determined by* (E^-) *and* $\leq_{\mathbf{B}}$ *Then (i)* $\exp(\mathbf{B}^\star(\alpha, 0)) = (\operatorname{imp}(\mathbf{B}))^-_\alpha \cap \exp(\mathbf{B})$, *(ii)* $(\operatorname{imp}(\mathbf{B}))^-_\alpha = \operatorname{Cn}(\exp(\mathbf{B}^\star(\alpha, 0)) \cup \{\neg\alpha\}) \cap \operatorname{imp}(\mathbf{B})$

Theorem 10 *Let* $\mathbf{B} \in \mathcal{B}$, *let the transmutation schema* * *be an adjustment, and let* $0 < i < \mathcal{O}$ *Let* * *be the revision operator for* $\operatorname{imp}(\mathbf{B})$ *uniquely determined by* $(E^\star)$ *and* $\leq_{\mathbf{B}}$ *Then (i)* $\exp(\mathbf{B}^\star(\alpha, i)) = (\operatorname{imp}(\mathbf{B}))^\star_\alpha \cap \{\exp(\mathbf{B}) \cup \{\alpha\}\}$, *and (ii)* $(\operatorname{imp}(\mathbf{B}))^\star_\alpha = \operatorname{imp}(\mathbf{B}^\star(\alpha, i))$

Theorem 9(1) captures the dependence of a theory base contraction on the content of the theory base that is, exp(B) In particular, a sentence is retained in the theory base contraction if and only if it is member of the theory base and it would be retained in the corresponding theory contraction Theoren 10(1) establishes a similar result for revision This substantiates our claim that adjustments retain as 'much as possible' of the original theory base

Consequently, a partial entrenchment ranking can be used to model belief change for a limited reasoner In particular, partial entrenchment rankings can be used to represent the reasoner's explicit information, its commitment to that information, and an encapsulation of the desired dynamic behaviour In fact, depending on the nature of the theorem prover adopted for an implementation some resource bounds could also be introduced [Williams, 1993] Theorems 9(1) and 10(1) clearly demonstrate that theory base adjustments contain as many of the explicit beliefs as would be retained in the corresponding theory change operations In particular, we have established that each explicit belief retained in a theory change is also retained via theory base adjustments

Theorems 9(11) and 10(n) show that theory change operators can be formulated in terms of theory base change using adjustments, and they provide explicit constructions for the Theorems 4 and 5, respectively

We define two partial entrenchment rankings to be *equivalent* whenever they generate the same epistemic entrenchment ordering It is obvious that, if $\mathbf{B}_1$ and $\mathbf{B}_2$ are equivalent, then for all $\alpha \in \mathcal{L}$ $Cn(cut(\mathbf{B}_1, \alpha)) = Cn(cut(\mathbf{B}_2, \alpha))$, and $imp(\mathbf{B}_1) = imp(\mathbf{B}_2)$ For example, if $\mathbf{B}_1$ is given by $\mathbf{B}_1(\alpha \rightarrow \beta) = 2$, $\mathbf{B}_1(\alpha) = 2$, $\mathbf{B}_1(\gamma) = 1$, and $\mathbf{B}_2$ is given by $\mathbf{B}_2(\beta) = 2$, $\mathbf{B}_2(\alpha) = 2$, $\mathbf{B}_2(\alpha \rightarrow \gamma) = 1$, then $\mathbf{B}_1$ and $\mathbf{B}_2$ are equivalent

For revision the adjustments of equivalent partial entrenchment rankings result in equivalent implicit information sets Furthermore, it turns out that the second parts of Theorems 9 and 10 can be generalised to equivalent partial entrenchment rankings, however we provided the current readings to for the sake of simplicity

7 Related Work

Gardenfore and Rott [1992] provide a comprehensive analysis of various prominent approaches [Hansson 1989, Fuhrmann 1991, Nebel, 1991] to theory base change, and the interested reader should consult their work

Nayak [1993] and Boutilier [1993] explore iterated theOTy change, and their approaches are closely related to transmutations because they focus on changes at the preference relation level, and as a consequence if the underlying preference relation is well-ranked then both of their procedures can be expressed as transmutations Boutiher's *natural revision* being a special case of adjustment, and Nayak's method being a special case of conditionalization For example, if we use Theorem 3 to capture the relationship between an epistemic entrenchment ordering and a partial entrenchment ranking, then Boutilier's approach can be characterized as an (a, l)-adjustrnent

Boutilier's new information is always accepted minimally firmly [Spohn 1988, p 114], and Nayak's new information is always accepted maximally firmly [Spohn 1988, pII4] In other words, they both suffer from problems described by Spohn, since their representation is essentially a *simple conditional function,* and their information input a sentence alone Spohn claims that both of these extreme schemes are undesirable because we don't always want to accept new information with the same degree of firmness

In defense of Nayak and Boutilier we should point out that the degree of firmness for new information may not always be available, in which case we could adopt default firmnesses for newly acquired information when the quality of the evidence is unknown

If the underlying preference relation is well-ranked then transmutation schema* can capture any conceivable change within the AGM paradigm For the purpose of an implementation we will almost certainly be dealing with a well-ranked, probably finite, preference relation

AN alternative approach to the problem of iterated revision is adopted by Darwiche and Pearl [1994], and Freund and Lehmann [1994] In particular, they have studied the iterated revision problem at the axiomatic level, and they suggest new meta-postulates for iterated revision rather than constraints on the modification of preference relations

Williams *et al* [1995b] capture Spohn's notion of *reason for* [1983] using adjustments of entrenchment rankings, and the computational model provided in section 5 can be used to implement Spohman reasons

Finally, we highlight some connections with nonmonotonic reasoning Gardenfors and Makinson [1994] use an comparative expectation ordering to construct a nonmonotonic inference relation Williams [1995a] defines a *partial expectation ranking* to be is a function B from a subset of sentences in *C* into *O* such that (PERI) and (PER2) are satisfied We also note that Reseller's *plausibility indexes* [1976] are essentially partial expectation rankings Given a partial expectation ranking B, if we define a nonmonotonic inference relation h, as af—/3 if and only if exp(B*(α,i)) (- β where 0 < 1, then (-v is *consistency preserving* and *rational,* and the associated model structure is nice [Gardenfors and Makinson, 1994]

Williams [1995a] shows that adjustments can be applied to partial expectation rankings thus providing a mechanism for *changing nonmonotonic inference relations* in an absolute and minimal way

Boutiher [1993], Pearl [1994], Freund and Lehmann [1994] also address the idea of changing default information, arguing that default information is usually quite stable, for example although our information about the flight coefficient of a particular bird may change dramatically, the default that *typically birds fly* will invariably remain unchanged, therefore changes should maintain as much default information as possible An adjustment not only concords with this perspective, but it preserves the properties of consistency perservation and rationality of the inference relation

8 Discussion

We have shown that the AGM paradigm can be extended to solve two outstanding practical problems that arise in the development of a computational model for belief revision This extension focuses on a *finite representation* of an epistemic entrenchment ordering, and the determination of a *policy for change* based on the principle of minimal change

We established that partial entrenchment rankings can be used to construct theory base change operators, theory change operators, as well as nonmonotonic inference relations Moreover, we provided representation results for the construction of well-behaved and very well-behaved theory change operators based on partial entrenchment rankings

We used partial entrenchment rankings to represent a well-ranked epistemic entrenchment ordering, and we provided a *computational model* for modifying partial entrenchment rankings based on an *absolute* measure *of minima/ change* which dealt with the removal of old information and the acquiescence of new information

We established that theory base revision and theory base contraction operators based on adjustments are related via Levi and Harper Identities at both the information content and the preference relation levels Finally, we demonstrated that theory base operators based on an adjustment maintain as much *explicit* information as is retained by the corresponding theory change operator

Acknowledgements

The author gratefully acknowledges, Peter Garden fore for suggesting how Spohn'B work might be used in the development of a computational model, David Makinson for a myriad of insightful suggestions concerning theory base change Many thanks are also due to Craig Boutilier, Didier DUBOIS, Norman Foo, Georg Gottlob, Abhaya Nayak, Bernhard Nebel, Erik Olsson, Maurice Pagnucco, Judea Pearl, Pavlos Pep pas, Wlodek Rabinowicz and Hans Rott for their helpful comments

References

[Alchourron *et al]* Alchourron, C , Gardenfora, P , Makinson, D , *On the Logic of Theory Change Partial Meet Functions for Contraction ana Revision,* Journal of Symbolic Logic, 50 510-530, 1985

[Boutilier, 1993] Boutiher, C , *Revision Sequences and Nested Conditionah,* Thirteenth International Joint Conference on Artificial Intelligence, Morgan Kaufmann, 519 - 525, 1993

[Darwiche and Pearl, 1994] Darwiche, D and Pearl, J , *On the Logic of Iterated Belief Revision,* in Fagin (ed), Proceedings of the 1994 Conference on Theoretical Aspects of Reasoning about Knowledge, 5 - 23, 1994

[Freund and Lehmann, 1994] Freund, M and Lehmann, D , *Belief Revision and Rational Inference,* Technical Report TR 94-16, Institute of Computer Science, The Hebrew University of Jerusalem, Israel, 1994

[Furhmann, 1991] Fuhrmann, A , *Theory Contraction through Base Contraction,* Journal of Philosophical Logic, 20 175 - 203, 1991

[Gardenfors, 1988] Gardenfors, P, *Knowledge* in *Flux Modeling the Dynamics of Epistemic States,* Bradford Books, The MIT Press, Cambridge Massachusetts, 1988

[Gardenfors and Makinson 1988] Gardenfors, P , and Makinson, D , *Revisions of Knowledge Systems using Epistemic Entrenchment,* Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge, p83 - 96, 1988

[Gardenfors and Makinson, 1994] Gardenfors, P and Makinson, D , Nonmonotonic *Inference Based on Expectations,* Artificial Intelligence Journal 65, 197 - 245 1994

[Gardenfors and Rott, 1992] [Gardenfors, P and Rott, H , *Belief Revision,* Report 30, Zentrum Philosophic und Wissenschaftstheone, Universitat Konstanz, Germany, 1992

[Grove, 1988] Grove, A , *Two Modellings for Theorg Change,* Journal of Philosophical Logic 17 157- 170, 1988

[Hansson, 1989] Hansson, S O , *New* Operators *for Theory Change,* Theona 55 115 - 132, 1989

[Katsuno and Mendelzon, 1992] Katsuno, H and Mendelzon, A O , *On the Difference between Updat ing a Knowledge Database and Revising it,* in Belief Revision, Gardenfors, P (ed), Cambridge Press, Cambridge, 1992

[Makinson, 1987] Makinson, D , *On the Status of the Postulate of Recovery* in *the Logic of Theory Change,* Journal of Philosophical Logic, 16 383 - 394, 1987

[Nayak, 1993] Nayak, A , *Studies* in *Belief Change,* PhD Dissertation, Department of Philosophy, University of Rochester, USA, 1993

[Nebel, 1991] Nebel, B , *A Knowledge Level Analysis of Belief Revision,* in Principles of Knowledge Representation and Reasoning Proceedings of the Third International Conference, Morgan Kaufmann, San Mateo, CA, 301 - 311, 1991

[Pearl, 19941 Pearl, J , *From Adams' Conditionals to Default Expressions, Causal Conditionals, and Counter]'actuals,* in E Eells and B Skyrma (eds), Probability and Conditionals, Cambridge University Press, 47 - 74, 1994

[Peppas, 1993] Peppas, P , *Belief Change and Reasoning about Action* PhD Dissertation, University of Sydney, Australia, 1993

[Peppas and Williams 1995] Peppas, P , and Williams, M A , *Constructive Modellings for Theory Change,* Notre Dame Journal of Formal Logic, to appear

[Rescher, 1976] Resetter, N , *Plausible Reasoning,* Van Gorcum, Assen, Amsterdam, 1976

[Rott, 1991a] Rott, H , *Two Methods of Constructing Contractions and Revisions of Knowledge Systems,* Journal of Philosphical Logic, 20 149 - 173, 1991

[Rott, 1991b] Rott, H , *A Nonmonotonic Conditional Logic for Belief* Revision /, in A Fuhrmann and M Morreau (eds), The Logic of Theory Change, Springer-Verlag, LNAI 465, Berlin, 135 - 183, 1991

[Spohn, 1983], Spohn, W , *Deterministic and Probabilistic Reasons and Causes,* Erkenntnis 19 371-396, 1983

[Spohn, 1988], Spohn, W, *Ordinal Conditional Functions A Dynamic Theory of Epistemtc States,* In Harper, W L , and Skyrms, B (eds), Causation in decision, belief change.and statistics, II, Kluwer Academic Publishers, pl05 - 134, 1988

[Williams, 1992] Williams M A , *Two operators for theory base change,* in the Proceedings of the 5th Australian Joint Conference on Artificial Intelligence. p259 - 265, 1992

[Williams, 1993] Williams, M A , *Implementing Knowledge* Tronsmutations, Information Systems Group Report ISRR-006-1993, University of Newcastle, 1993

[Williams, 1994a], Williams, M A , *Transmutations of Knowledge Systems,* in J Doyle, E Sandewall, and P Torasso (eds), Principles of Knowledge Representation and Reasoning Proceedings of the Fourth International Conference, Morgan Kaufmann, San Mateo, CA, 619 - 629, 1994

[Williams, 1994b] Williams, M A , Ezp/anaiion *and Theory Base Transmutations,* in the proceedings of the Eleventh European Conference on Artificial Intelligence, A Cohn (ed), Wiley, 346 - 350, 1994

[Williams, 1994c] Williams, M A , *On the Logic of Theory Base Change,* in Logics in Artificial Intelligence, C MacNish, D Pearce and L M Pereira (eds), LNCS No 835, 86 - 105, Spnngei Verlag, 1994

[Williams, 1995a] Williams, M A , *Changing Nonmonotonic Inference Relations,* in the Proceedings of the Second World Conference on the Foundations of Artificial Intelligence, Paris, 1995 (to appear)

[Williams *et al,* 1995b] Williams, M A , Pagnucco, M , Foo, N Y and Sims, B (1995), *Determining Explanations using Knowledge Transmutations,* in Proceedings of the International Joint Conference on Artificial Intelligence, (this volume)