

 Open access • Journal Article • DOI:10.1007/BF00248151

Iterated revision and minimal change of conditional beliefs — [Source link](#)

[Craig Boutilier](#)

Institutions: [University of British Columbia](#)

Published on: 01 Jun 1996 - [Journal of Philosophical Logic](#) (Kluwer Academic Publishers)

Topics: [Belief revision](#) and [Ramsey RESET test](#)

Related papers:

- [On the logic of theory change: Partial meet contraction and revision functions](#)
- [On the logic of iterated belief revision](#)
- [Ordinal Conditional Functions: A Dynamic Theory of Epistemic States](#)
- [Two modellings for theory change](#)
- [Iterated belief change based on epistemic entrenchment](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/iterated-revision-and-minimal-change-of-conditional-beliefs-4diah64zoi>

Iterated Revision and Minimal Change of Conditional Beliefs

Craig Boutilier
University of British Columbia

Abstract

We describe a model of iterated belief revision that extends the AGM theory of revision to account for the effect of a revision on the conditional beliefs of an agent. In particular, this model ensures that an agent makes as few changes as possible to the conditional component of its belief set. Adopting the Ramsey test, minimal conditional revision provides acceptance conditions for arbitrary right-nested conditionals. We show that problem of determining acceptance of any such nested conditional can be reduced to acceptance tests for unnested conditionals. Thus, iterated revision can be accomplished in a “virtual” manner, using uniterated revision.

1 Introduction

The acceptance conditions for subjunctive conditionals have been widely studied, but no criterion has drawn more attention recently than the *Ramsey test*. Supposing an agent to possess some belief set K , it is instructed by Stalnaker to accept the conditional $A > B$ iff it satisfies the following test:

First add the antecedent (hypothetically) to your stock of beliefs; second make whatever adjustments are required to maintain consistency (without modifying the hypothetical belief in the antecedent); finally, consider whether or not the consequent is true. [32, p.44]

The crucial step in the Ramsey test is the revision of the belief set. The notion of revision adopted will determine which conditionals are accepted and rejected. Conversely, given a fixed (complete) set of accepted conditionals, the revision function adopted by an agent will also be determined.

The *AGM theory of revision* due to Alchourrón, Gärdenfors and Makinson [2] provides guidelines for the revision of objective belief sets (containing no conditional sentences). These are usually presented as postulates that constrain the behavior of revision functions. We use K_A^* to denote the

belief set resulting from the revision of K by A . Unfortunately, the theory does not address the issue of revising conditional belief sets. Gärdenfors [14] proposes that belief sets be extended to include conditional beliefs, requiring only that the following postulate, reflecting the Ramsey test, be respected:¹

(RT) $A > B \in K$ iff $B \in K_A^*$.

This description of the revision of conditional belief sets says nothing about the preservation of conditionals in a belief set during the course of revision. A hallmark of the AGM theory is its commitment to the principle of *informational economy*: beliefs are only given up when there are no less entrenched candidates. Yet this principle is ignored in this formulation of conditional belief revision, as well as many related proposals. Given an ordering of entrenchment and belief set K , the AGM theory will determine the form of K_A^* . The extended theory, in contrast, when applied to a conditional belief set, imposes almost no constraints on the conditionals deemed acceptable in the new belief set. This has a tremendous impact on the analysis of *revision sequences*. Should an agent revise its belief set K by A , its conditionals $A > B$ determine the form of K_A^* . However, if a revision function imposes no constraints on subsequently accepted conditionals, there can be no logical constraints on subsequent revision of K_A^* .² The fact that K_A^* was reached by revision of K becomes meaningless.

In this paper, we extend the AGM model of revision to account for such sequences of belief change. We take the epistemic state of an agent to consist of both objective beliefs and (unnested) conditional beliefs. These conditional beliefs can be viewed as capturing the relative degrees of entrenchment of objective beliefs, or equivalently, as completely specifying the agent's revision policies. Acceptance of a new belief forces a change in the agent's objective belief set in accordance with the AGM theory. This necessarily requires some change in the agent's revision policies, or conditional beliefs. We describe how this change can be effected semantically by minimally changing the agent's judgements of the relative entrenchment of objective beliefs. This in turn has the effect of changing as few conditional beliefs as possible consistent with the dictates of the AGM theory (when restricted to the objective component of the belief set). Thus, the original conditional belief set uniquely determines the revised conditional belief set. In particular, from a set of simple, unnested conditionals, the effects

¹Of course, Gärdenfors points out that, given the other constraints of the AGM theory (even weakened somewhat), such a condition is untenable. This celebrated *triviality result* has been devoted considerable attention. But very compelling arguments (e.g., [23, 27]) have been put forth suggesting that the postulates of the AGM theory should not hold in the case of nonobjective beliefs. This is the point of view adopted here, as we elaborate below.

²At least, there can be no constraints other than those imposed by the objective part of K_A^* .

of arbitrary sequences of revisions can be computed. We will demonstrate an especially interesting consequence of this result, namely, that any right-nested conditional arbitrary depth can be effectively reduced to an unnested conditional with the same acceptance conditions.

2 A Semantic Model for AGM Revision

We assume an agent to have a deductively closed set of beliefs K taken from some underlying language. For concreteness, we will assume this language to be that of classical propositional logic L_{CPL} , generated by some set of variables \mathbf{P} , and with an associated consequence operation Cn . The identically true and false propositions are denoted \top and \perp , respectively. We say K is finitely specifiable if $K = Cn(S)$ for some finite set of premises S . The *expansion* of K by new information A is the belief set $K_A^+ = Cn(K \cup \{A\})$. The *revision* of K by A is denoted K_A^* . The process of revision requires some care, for A may contradict elements of K . Alchourrón, Gärdenfors and Makinson [2] propose a method for logically delimiting the scope of acceptable revisions. To this end, the AGM postulates below, are maintained to hold for any reasonable notion of revision [15].

(R1) K_A^* is a belief set (i.e. deductively closed).

(R2) $A \in K_A^*$.

(R3) $K_A^* \subseteq K_A^+$.

(R4) If $\neg A \notin K$ then $K_A^+ \subseteq K_A^*$.

(R5) $K_A^* = Cn(\perp)$ iff $\models \neg A$.

(R6) If $\models A \equiv B$ then $K_A^* = K_B^*$.

(R7) $K_{A \wedge B}^* \subseteq (K_A^*)_B^+$.

(R8) If $\neg B \notin K_A^*$ then $(K_A^*)_B^+ \subseteq K_{A \wedge B}^*$.

An alternative model of revision is based on the notion of *epistemic entrenchment* [15, 16]. Given a belief set K , we can characterize the revision of K by ordering beliefs according to our willingness to give them up when some contradictory information requires such. If one of two beliefs must be retracted in order to accommodate some new fact, the less entrenched belief will be relinquished, while the more entrenched persists. Gärdenfors [15] presents five postulates for such an ordering and shows that these orderings determine exactly the space of revision functions satisfying the AGM

postulates. We let $B \leq_E A$ denote the fact that A is at least as entrenched as B in theory K . A complete set of sentences of this form is sufficient to specify a revision function. We note that the dual of an entrenchment ordering is a plausibility ordering on sentences. A sentence A is more plausible than B just when $\neg A$ is less entrenched than $\neg B$, and means that A would be more readily accepted than B if the opportunity arose. Grove [19] studied this relationship and its connection to the AGM theory.

2.1 A Semantic Model

Our extension of the AGM theory will be based on a slightly different semantic representation of revision functions. Grove [19] observed that an entrenchment relation over sentences can be represented using a plausibility ordering over models. Our semantics captures this intuition, though it is (cosmetically) distinguished by the use of an explicit ranking function rather than Grove's system of spheres. In what follows, we describe our basic semantic models and the truth conditions for a number of connectives. While we do not give details here, we note that this system can be axiomatized in using a simple modal logic CO^* , and that each of the connectives and semantic restrictions described below can be represented in this logic. We refer to [8, 9, 5] for details of this type.

We imagine an agent to have some objective belief set $K \subseteq \mathbf{L}_{CPL}$ as well as a collection of revision policies to guide changes of this belief set. A *revision model* is a triple $M = \langle W, \leq, \varphi \rangle$ intended to capture both aspects of the agent's epistemic state. W is a set of worlds and $\varphi : \mathbf{P} \rightarrow 2^W$ is a valuation mapping, assigning atoms to the subsets of worlds at which they hold. The truth of an arbitrary proposition at any world is defined in the usual way. We write $M \models_w A$ to indicate that A holds at world w , and denote by $\|A\|$ the set of such *A-worlds* (M is usually understood). We also use this notation for sets of sentences S , $\|S\|$ denoting those worlds satisfying each element of S , and somewhat loosely refer to worlds falsifying some element of S as $\neg S$ -worlds.

Finally, \leq is a binary relation over W . We insist that \leq be transitive and connected.³ The interpretation of \leq is as follows: $v \leq w$ iff v is as at least as *plausible* a state of affairs as w . Plausibility is a pragmatic measure that reflects the degree to which one would accept w as a possible state of affairs. If v is more plausible than w , loosely speaking v is "more consistent" with the agent's beliefs than w . Since \leq is a "total preorder," W is partitioned into \leq -equivalence classes, or *clusters* of equally plausible worlds. These cluster are themselves totally ordered by \leq . Thus, \leq can be

³ \leq is (totally) connected if $w \leq v$ or $v \leq w$ for any $v, w \in W$ (this implies reflexivity). This restriction is relaxed in [9], where we develop a weaker logic based on a reflexive, transitive accessibility relation.

viewed as a qualitative ranking relation, assigning to each world a degree of plausibility. Figure 1 illustrates the structure of such an ordering.

We require that epistemically possible worlds be more plausible than epistemically impossible worlds, and that all epistemically possible worlds are equally plausible. In other words, K -worlds should be exactly those minimal in \leq :

$$w \leq v \text{ for all } v \in W \text{ iff } M \models_w K$$

Such models are called *K-revision models* and have as their minimal cluster the set $\|K\|$ (see Figure 1).⁴ Although this is sufficient structure to define a reasonable revision function, we adopt two more restrictions. First, we insist that K -revision models are “complete;” that is, every propositional valuation is associated with some element of W . In other words, the valuations induced by φ and associated with each $w \in W$ must cover the entire range of logical possibilities. This ensures that an agent accords every possibility some degree of plausibility. Second, we assume that each model satisfies the following (restricted) well-foundedness condition.

Definition 1 Let $M = \langle W, \leq, \varphi \rangle$ be a CO-model. For any $A \in \mathbf{L}_{CPL}$ we define

$$\min(M, A) = \{w \in W : M \models_w A, \text{ and } M \models_v A \text{ implies } w \leq v \text{ for all } v \in W\}$$

A model is *well-founded* iff, for all A , $\min(A, M) \neq \emptyset$.

We are now in a position to define a conditional connective and associated belief revision function (suitable for K). When a new belief A is accepted, the agent should adopt the set of most plausible A -worlds, $\min(M, A)$, as its set of newly epistemically possible situations. The conditional sentence $A > B$ asserts that “If the agent were to adopt belief in A , it would believe B .” The acceptance of such a conditional relative to M is specified as follows:⁵

$$M \models A > B \text{ iff } \min(M, A) \subseteq \|B\|$$

The Ramsey equivalence immediately provides us with the revision function induced by M (w.r.t.

⁴If K is finitely specifiable then this constraint can be expressed in language of CO. Intuitively, this corresponds to the notion of “only knowing” K : all and only K -worlds are epistemically possible [22, 5].

⁵The acceptance condition can be specified in a way that does not presume the well-foundedness restriction, and the equivalence to the AGM theory described below is unaffected. Thus, the logic CO* provides a mechanism for expressing with unary modal connectives the truth conditions proposed by Lewis [24] to avoid the *Limit Assumption* [9]. However, the representation result of Grove ensures that the class of well-founded models is sufficient to capture all AGM revision functions.




Figure 1: Truth conditions for the conditional $A > B$

K). For any $A \in \mathbf{L}_{CPL}$ we define

$$K_A^{*M} = \{B \in \mathbf{L}_{CPL} : M \models A > B\}. \quad (1)$$

We can show that $*$ satisfies the AGM postulates for belief revision and any AGM revision operator has an equivalent formulation as such a $*$.

Theorem 1 [9] *Let M be a K -revision model and $*^M$ the revision function determined by M . Then $*^M$ satisfies postulates (R1) through (R8).*

Theorem 2 [9] *Let $*$ be a revision function satisfying postulates (R1) through (R8). For any theory K there exists a K -revision model M such that $K_A^* = K_A^{*M}$ for all A .*

We call $A > B$ a *simple conditional* iff $A, B \in \mathbf{L}_{CPL}$. Such unnested conditionals are especially interesting since they alone determine the revision function $*^M$.

In Figure 1, we have a typical K -revision model with each large circle representing a cluster of mutually accessible (equally plausible) worlds, with arrows indicating accessibility between clusters. The minimal cluster consists of all K -worlds, and we have that $K \vdash \neg A$. The set of minimal A -worlds is indicated by the shaded region, and this set forms the set of “newly possible” worlds when K is revised by A . Thus $A > B$ should hold just when B is true at each world in the shaded region.

2.2 Extended Belief Sets

We will have occasion to use several other connectives whose semantics we introduce here. First, we define a modality for belief, reading $\mathbf{B}\alpha$ as “ α is believed:”

$$M \models \mathbf{B}\alpha \text{ iff } \min(M) \subseteq \|\alpha\| \text{ iff } M \models \top > \alpha$$

This holds just when α is true at each epistemically possible world, those minimal in the plausibility ordering \leq . We note that \mathbf{B} behaves according to the usual weak S5 interpretation of belief [5].

Statements regarding the relative entrenchment and plausibility of various propositions can also be expressed. We write $B \leq_E A$ to denote that A is at least as entrenched as B , meaning that A is at least as strongly held a belief as B . We have

$$M \models B \leq_E A \text{ iff } \min(M, \neg B) \leq \min(M, \neg A)$$

The dual notion of plausibility, where $B \leq_P A$ denotes the fact that A is at least as plausible as B , is captured by

$$M \models B \leq_P A \text{ iff } \min(M, A) \leq \min(M, B) \text{ iff } M \models \neg A \leq_E \neg B$$

For any revision model, we define the *objective belief set* associated with it to be those propositional sentences that are “believed” in the model, namely.

$$\{\alpha \in \mathbf{L}_{CPL} : M \models \mathbf{B}\alpha\}$$

We will sometimes refer to this as the propositional belief set or simply the belief set for M . Naturally, the belief set for any K -revision model is just K . We will be more interested here in “subjective beliefs” associated with a revision model, those sentences that are believed, but contain modal operators. Of particular concern are those conditionals that are believed by an agent. We therefore extend the notion of belief set to cover arbitrary sentences. We denote by \mathbf{L} the full language that includes the connectives $>$, \mathbf{B} , \leq_E , and any other connectives that we introduce.⁶ The *extended belief set* associated with M is

$$\{\alpha \in \mathbf{L} : M \models \mathbf{B}\alpha\}$$

⁶Technically, we can take this to be the modal language of CO*, in which all such connectives are defined.

We note the following properties of belief sets and extended belief sets. Clearly, for any K -revision model we have $K \subseteq E$. Furthermore, extended sets are *stable sets of beliefs* in the sense of Stalnaker [33]:

Proposition 3 *Let E be the extended set determined by M . Then: a) $A \in E$ iff $\mathbf{B}A \in E$; b) $A \notin E$ iff $\neg\mathbf{B}A \in E$; and c) E is closed under tautological consequence.*

This is a reasonable requirement on the belief set of any rational agent with full powers of introspection. Given this introspective ability, we expect an agent to be able to completely determine the beliefs contained in any revised belief set K_A^* . While the Ramsey test ensures that simple conditionals $A > B$ will be in E if $B \in K_A^*$, we have an even stronger property: if a conditional fails the Ramsey test its negation is believed.

Proposition 4 *Let E be the extended set determined by M . Then for any simple conditional $A > B$, either $A > B \in E$ or $\neg(A > B) \in E$.*

We refer to this property as the *completeness of conditional beliefs*.

We note that the connectives \mathbf{B} and $>$ have “global” truth conditions. By this we mean that if, say, $A > B$ holds at some world in a model M , then it holds at all worlds in M . Its acceptance does not depend on the world at which it is being evaluated, but rather on the global (rather than local) structure of M . This has important consequences when we consider the meaningfulness of iterated belief and conditional statements in extended sets. For instance, we immediately have that $\mathbf{B}(A > B) \in E$ iff $A > B \in E$. Thus, extended sets provide a mechanism for incorporating both objective beliefs and belief in simple conditionals into a single collection. However, the global nature of acceptance conditions ensures that non-simple conditionals add nothing of interest to the language.

Let us dub any sentence such that all atomic variables lie within the scope of a modality \mathbf{B} or a conditional $>$ a *subjective sentence*. We concentrate on the case of conditionals whose consequent is subjective. It is easy to see that the truth conditions for any conditional $A > \alpha$, for any purely subjective sentence α , are uninteresting. As long as A is satisfiable, $A > \alpha$ holds iff α does. Thus, this model of the AGM postulates is suitable only for objective belief sets. For example, we have that $M \models A > (B > C)$ iff $\min(M, A) \subseteq \parallel B > C \parallel$ iff $M \models B > C$. This is certainly not a reasonable property of extended belief sets or revision models. In particular, if the Ramsey test is taken seriously this semantics is not suitable for iterated revision. However, it can be easily extended.

3 Minimal Conditional Revision

3.1 The Problem of Iterated Revision

Suppose an agent possesses a belief set K . Its revision policies and judgements of entrenchment are captured by the K -revision model M . If the agent is curious about the status of proposition B after an arbitrary propositional revision A , it may simply consider whether $B \in K_A^*$. In other words, for objective formulae B simply asking whether B is true at each world in $\min(M, A)$ is appropriate, and is sufficient (via the Ramsey test) to determine acceptance of the conditional $A > B$. While this characterization of K_A^* is clear, it is less obvious just what form the revised *extended* belief set, E_A^* , should take. As made clear above, the new *subjective* beliefs of the agent cannot be captured by the set of worlds $\min(M, A)$. If it were, an agent could drastically change its objective belief set K without altering its conditional beliefs, its nested beliefs, or its judgements of entrenchment.

It should be clear why this difficulty arises. The model M is a K -revision model, suitable for modeling the revision of belief set K , and for determining the truth of conditionals relative to K . When the agent's objective beliefs are revised, giving rise to belief set K_A^* , such a model is no longer adequate as a representation of the agent's epistemic state. It cannot retain the same conditionals given a different belief set. Recall, the Ramsey test provides *acceptance* conditions for conditionals, not *truth* conditions; and acceptance is relative to a given state of belief. Similarly, if the agent's beliefs change, so too must its "nested" beliefs and its judgements of entrenchment and plausibility. Thus, the K -revision model must be given up. The agent must adopt a revised model of its epistemic state, which we denote M_A^* . It is this model that determines the agent's revised extended belief set E_A^* .⁷

What are the natural requirements on the new model M_A^* ? Since K_A^* is uniquely determined by M we insist that $K_A^* \subseteq E_A^*$ and that K_A^* form the *entire* objective component of E_A^* . That is, M_A^* should be a K_A^* -revision model, or that K_A^* is *only known* in M_A^* . The minimal cluster of worlds in M_A^* should be exactly $\min(M, A)$. This is illustrated in Figure 2. Let us dub this constraint the *Basic Requirement* on revision functions as applied to models.

The Basic Requirement: If M is a K -revision model determining revision function $*$, then the revision model M_A^* must be such that $\min(M_A^*, \top) = \min(M, A)$.

In fact, from a purely logical perspective, this is probably all we want to say about M_A^* . If one changes an objective belief, it is impossible in general to predict what becomes of one's conditionals.

⁷Similar observations have been made by Levi [23], Rott [29] and Hansson [21].




Figure 2: General constraints on the mapping to a new revision model

This model of iterated revision is captured by Gärdenfors's [15] *belief revision systems*, although not in this semantic fashion. Such a system consists of a set \mathbf{K} of belief sets and a revision function $*$ that maps $\langle K, A \rangle$, where $K \in \mathbf{K}$ and $A \in \mathbf{L}_{CPL}$, to $K_A^* \in \mathbf{K}$. The function $*$ must satisfy the AGM postulates, but the behavior of $*$ on the extended belief set E associated (by the Ramsey test) with K is left unspecified. Only postulates (R7) and (R8) constrain how revision of K_A^* should take place, and these constraints are quite mild. This model provides no guidance as to what conditionals in E should be accepted or rejected in E_A^* . Put another way, no hints are provided on how to revise K_A^* , given its “origination” as a revision of K by A . The ordering of worlds in the nebulous region of Figure 2 (b) is completely unspecified.

Belief revision systems of the Gärdenfors type possess two unattractive features, one logical and one pragmatic. First, epistemic states are distinguished solely by their objective component. For any K , the extended belief set E associated with K (via the Ramsey test) is fixed, since K_A^* is fixed. There can be no two distinct extended belief sets that share the same objective part. In other words, there can only be one way of revising a belief set. This is certainly an unnatural restriction, one that is not imposed by our notion of K -revision models. There exists a multitude of K -revision models for any fixed K , reflecting the view that conditionals and judgements of entrenchment form an integral part of an agent's epistemic state. An agent can be in two different belief states, but have the same propositional beliefs in each and differ only on its accepted conditionals. This is a natural extension of revision systems, but it is also quite straightforward. The semantic “characterization” of such a system is completely captured by Figure 2. The general method of Rott [28] is based on the same considerations. He identifies the epistemic state of an agent with an ordering of entrenchment over

L_{CPL} , capturing exactly the same information as a revision model. Hansson [21] also associates potentially distinct revision functions with equivalent belief sets.⁸

The second criticism is of a more pragmatic nature and is the point we address in the rest of the paper. Everything above points to a mapping of M to M_A^* satisfying one condition, that the objective belief set associated with M_A^* be precisely the set K_A^* determined by M . As Figure 2 illustrates, just about all of the ordering information, capturing an agent's conditional beliefs and judgements of entrenchment, is (potentially) lost in this mapping. There is something unsatisfying about this model. The ordering relation \leq is intended to reflect the informational content or importance of beliefs. When certain beliefs must be given up, it seems reasonable to try to keep not only important objective beliefs, but as much of the *ordering* a possible. Revision should not generally change one's opinion of the relative importance of most sentences.

3.2 The Semantics of Minimal Conditional Revision

Instead of arbitrary mappings from M to M_A^* , we will propose a class of natural mappings that preserve as much ordering information as can be expected. This determines the class of *minimal conditional revision functions*, that tend to preserve the entrenchment information and conditional beliefs found in an extended belief set. It is important to note that the model we propose is *not* completely general, for it permits only a subset of those revision functions (on extended sets) allowed by the arbitrary mappings described above. However, it is a very natural subset, for it extends the notion of minimal change to subjective formulae. Furthermore, it is suitable for determining the result of propositional revision sequences, or the acceptance of right-nested conditionals, when the general model has little to offer.

The conditionals accepted by an agent are determined by its ordering of plausibility. If we insist that revision preserve as much of this ordering as possible, then, for the most part, the relative entrenchment and plausibility of sentences (hence conditional beliefs) will remain intact. Let $M = \langle W, \leq, \varphi \rangle$ be the revision model reflecting some extended belief set E . Given a propositional revision A of E (or the associated K), we must find a revision model $M_A^* = \langle W, \leq', \varphi \rangle$ such that \leq' reflects the minimal mutilation of \leq .

If $w \in \min(M, A)$, by the Basic Requirement w must be minimal in \leq' , and these must be the only minimal worlds in \leq' . For any such w the relationships $w \leq' v$ and $v \leq' w$ are completely determined by membership of v in $\min(M, A)$, independently of their relationship in \leq . Figure 2

⁸We discuss the relationship with these systems further in the concluding section.




Figure 3: The *minimal conditional revision* of a model.

illustrates this. For w, v not in $\min(M, A)$, this picture leaves $v \leq w$ completely unspecified. If \leq is to be left intact to the largest possible extent then the most compelling specification is to insist that $w \leq' v$ iff $w \leq v$. This has the effect of leaving \leq unaltered except as indisputably required by the Basic Requirement. Such a move, applied to the original model in Figure 2, is illustrated in Figure 3. We dub such a mapping on revision models the *minimal conditional revision operator*, and describe the revision function it induces on the associated belief and extended belief sets.

Definition 2 Let $M = \langle W, \leq, \varphi \rangle$ be a revision model. The *minimal conditional revision operator* (or MC-revision operator) $*$ maps M into M_A^* , for any $A \in \mathbf{L}_{CPL}$, where $M_A^* = \langle W, \leq', \varphi \rangle$, and: a) if $v \in \min(M, A)$ then $v \leq' w$ for all $w \in W$ and $w \leq' v$ iff $w \in \min(M, A)$; and b) if $v, w \notin \min(M, A)$ then $w \leq' v$ iff $w \leq v$.

Definition 3 Let E be the extended belief set associated with the revision model M . The *minimal conditional revision function* associated with M is $*$, defined for all $A \in \mathbf{L}_{CPL}$ as follows:

$$E_A^* = \{B \in \mathbf{L} : M_A^* \models \mathbf{B}B\}$$

Let $K \subseteq E$ be the objective component of E . The minimal conditional revision (or MC-revision) of K , denoted K_A^* , is the restriction of E_A^* to \mathbf{L}_{CPL} ; that is

$$K_A^* = \{B \in \mathbf{L}_{CPL} : M_A^* \models \mathbf{B}B\}$$

The MC-revision function is simply the AGM operator determined by M , when restricted to K .

Proposition 5 *Let M be a revision model for K , $*^M$ be the AGM revision function determined by M (as given by Definition 1), and $*$ be the MC-revision function determined by M . For any $A \in \mathbf{L}_{CPL}$, we have $K_A^{*^M} = K_A^*$.*

Notice that this enhances the AGM model of revision. The revised extended set E_A^* is defined using the updated revision model M_A^* and incorporates subjective beliefs, such as conditionals, statements of entrenchment and plausibility, and nested belief sentences.

If we are to extend the Ramsey test to include nested conditionals, the truth conditions for statements $A > \beta$ must be recast in this framework for subjective β . For M to satisfy $A > \beta$, we must have $\beta \in E_A^*$. If $\beta \in \mathbf{L}_{CPL}$, given Proposition 5, these truth conditions will be identical to those provided in Section 2. Thus, our new truth conditions for $>$ based on the Ramsey test will form a “conservative extension” of the old definition. However, for arbitrary $\beta \in \mathbf{L}_B$, especially sentences of the form $\beta = (B > C)$, we must test $B > C$ in M_A^* , not at $\min(M, A)$.

Definition 4 Let $M = \langle W, \leq, \varphi \rangle$ be a revision model, $A \in \mathbf{L}_{CPL}$ and $B \in \mathbf{L}$. M satisfies $A > B$ (written $M \models A > B$) iff $M_A^* \models \mathbf{B}B$.

We now have a conditional connective whose truth conditions are specified directly by the Ramsey test. In particular, conditionals with arbitrary consequents have meaningful acceptance conditions; and right-nested conditionals impose constraints on the process of iterated revision. We will see, however, that their acceptance can be verified using only conditions on the original revision model.

Notice that the truth of $A > B$ is unspecified for $A \notin \mathbf{L}_{CPL}$. MC-revision functions are suitable only for sequences of *objective* updates. The nesting of conditionals sanctioned in a meaningful way in this framework is *right-nesting*, for instance, $A > (B > C)$ where $A, B, C \in \mathbf{L}_{CPL}$. A sentence $(A > B) > C$ has an unspecified truth value for it asks if C is believed when a knowledge base is revised to include $A > B$. This framework does not specify how to revise a knowledge base with non-objective sentences, though we return to this problem in the concluding section.

3.3 Minimal Change of Conditional Beliefs

In this section, we investigate some of the properties of single changes to a revision model or belief set using the MC-revision function $*$. We assume throughout some revision model M capturing the extended belief set E and the belief set $K \subseteq E$. When M (or K or E) is revised by A the properties of K_A^* are obvious: $K_A^* = K_A^{*^M}$, where $*^M$ is the AGM operator determined by M . Of more interest is




Figure 4: MC-revision with a consistent revision.

the characterization of the new extended set E_A^* . Since minimal mutilation of the ordering relation \leq is intended to preserve as many conditional beliefs as possible, we must determine precisely which conditionals in E remain in E_A^* and which are sacrificed. Since we are interested in single revisions, we restrict our attention (for the time being) to *simple conditionals* of the form $A > B$ where both A and B are objective. One important property to note is the following:

Proposition 6 *Let M be a K -revision model where $A \in K$. Then $M_A^* = M$.*

Updating by a sentence already in a belief set not only causes no change in the belief set K , as required by the AGM postulates, but also leaves the revision model M intact. More generally, consider the two scenarios, illustrated by Figures 4 and 5, that might arise when M is revised by A . Figure 4 shows the situation where A is a consistent revision, $K \not\vdash \neg A$, while Figure 5 demonstrates the behavior of an inconsistent revision, $K \vdash \neg A$. We are interested in those simple conditionals $B > C$ that are true in the model M_A^* . In each of these two scenarios (in each figure) there are two cases to consider, $\neg B \notin K_A^*$ and $\neg B \in K_A^*$: either B is consistent with the new belief set K_A^* or it is not.

Consider the first situation where $\neg B \notin K_A^*$; that is, $M \models \neg(A > \neg B)$. This means that there is some B -world among the set $\min(M, A)$ of minimal A -worlds in M (the shaded region of the figures). Clearly then we have that $\min(M_A^*, B) = \min(M, A \wedge B)$; so $M_A^* \models B > C$ iff $M \models A \wedge B > C$. Therefore, whenever $K_A^* \not\vdash \neg B$, a conditional $B > C$ is in E_A^* iff $A \wedge B > C$ is in E . Notice that the status of $B > C$ in E has no bearing on its acceptance or rejection in E_A^* . This behavior is exactly in accordance with the AGM postulates (R7) and (R8). Any AGM revision function must ensure




Figure 5: MC-revision with an inconsistent revision.

that subsequent (iterated) consistent revisions are treated in the same manner as uniterated consistent revisions.

The second situation arises when $\neg B \in K_A^*$; that is, $M \models A > \neg B$. When K (or E) is revised by A , $\neg B$ is in the resulting belief set. This is true exactly when no B -world is contained in $\min(M, A)$ (again, the shaded region). Now, $M_A^* \models B > C$ just in case the set $\min(M_A^*, B)$ contains only C -worlds. Since $\neg B \in K_A^*$, the set $\min(M_A^*, B)$ is not contained in the minimal cluster of M_A^* . However, all worlds outside the minimal cluster stand in exactly the same relation as they do in M . Therefore $\min(M_A^*, B) = \min(M, B)$ and it follows that $M_A^* \models B > C$ iff $M \models B > C$. For conditionals $B > C$ whose antecedents are not made plausible by the acceptance of A (i.e., $K_A^* \not\models \neg B$), $B > C$ is in E_A^* iff $B > C$ is in E . Since nothing forces the conditional to be abandoned when A is accepted, it is retained.

We can summarize these considerations in the following theorem and equivalent corollaries.

Theorem 7 Let M be a revision model, let $*$ be the MC-revision operator and let $A, B, C \in \mathbf{L}_{CPL}$.

- a) If $M_A^* \models \mathbf{B}\neg B$ then $M_A^* \models B > C$ iff $M \models B > C$.
- b) If $M_A^* \not\models \mathbf{B}\neg B$ then $M_A^* \models B > C$ iff $M \models A \wedge B > C$.

Corollary 8 Let M be a K -revision model and let $*$ be the MC-revision function.

- a) If $K_A^* \vdash \neg B$ then $C \in (K_A^*)_B^*$ iff $C \in K_B^*$.
- b) If $K_A^* \not\vdash \neg B$ then $C \in (K_A^*)_B^*$ iff $C \in K_{A \wedge B}^*$.

Corollary 9 *Let M be an E -revision model and let $*$ be the MC-revision function.*

- a) *If $A > \neg B \in E$ then $B > C \in E_A^*$ (i.e., $A > (B > C) \in E$) iff $B > C \in E$.*
 b) *If $A > \neg B \notin E$ then $B > C \in E_A^*$ (i.e., $A > (B > C) \in E$) iff $A \wedge B > C \in E$.*

These results precisely characterize the conditionals that will be preserved in a revised extended belief set. Each shows that the sentences accepted in the new revision model or belief state can be determined by appeal to the original model or belief state. Theorem 7 shows that the conditional belief set captured by M_A^* can be determined by the conditional beliefs of M . Furthermore, it demonstrates that MC-revision preserves as much conditional information in the revised belief set as is consistent with the AGM postulates. The conditionals removed from E when constructing E_A^* are only those compelled by postulates (R7) and (R8). These are just those conditionals $B > C$ whose antecedent B is consistent with K_A^* . This is reflected in clause b) of the theorem. However, as indicated by clause a), the remaining set of conditionals (or negated conditionals) in E_A^* coincides precisely with the conditional information in the original extended set E . *Thus, no AGM revision function could preserve more conditional information than the MC-revision function.*

Corollary 8 shows that the sequence of two revisions applied to K can be reduced to a single revision, requiring no iterated revision, and that the test to establish which condition holds also requires no iterated revision. Similarly, Corollary 9 shows that the revised extended belief set E_A^* and the nested conditionals in E can be captured by the simple, unnested conditionals in E . These properties will play a vital role in our characterization of revision sequences in terms of single updates.

3.4 Revision Sequences

The objective belief set K_A^* formed by revising K with A relies only on one application of $*$ to K . Somewhat surprisingly, the simple conditionals in E_A^* can also be discovered by referring only to applications of $*$ to K . As the results above indicate, to ask if $B > C \in E_A^*$, one first asks if $\neg B$ is in K_A^* . If this is true, then asking if $B > C \in E_A^*$ is equivalent to asking if $C \in K_B^*$. If this is false, then asking if $B > C \in E_A^*$ is equivalent to asking if $C \in K_{A \wedge B}^*$. A “hypothetical” revision of K by B or by $A \wedge B$ is sufficient to determine the status of $B > C$ in E_A^* (or equivalently, the status of $A > (B > C)$ in E).

Suppose we have a revision sequence, A_1, \dots, A_n , to be applied to E , where $K \subseteq E$. We can clearly use a single revision to verify whether some simple conditional $B > C$ is in $E_{A_1}^*$, or if $M_{A_1}^* \models B > C$. However, if we replace B by A_2 we see that these conditionals determine precisely




Figure 6: The order dependence of MC-revision.

the beliefs obtained when $E_{A_1}^*$ is revised by A_2 , that is, $(E_{A_1}^*)_{A_2}^*$. If single revisions applied to K can establish the content of the belief set obtained by an iterated revision of two levels, there seems no reason single (uniterated) revisions of K cannot be used to decide the outcome of arbitrary revision sequences. We will examine the properties of such sequences and show how they may be reduced to single uniterated revisions, or unnested conditionals. This has important computational implications, for it means that queries regarding the beliefs of an agent after a sequence of revisions can be answered using only “virtual” revision of its belief set. In particular, it is not necessary to construct a new belief set and entrenchment ordering after each update. Information contained in the original ordering is sufficient to determine the result of an arbitrary sequence of revisions.

3.4.1 Order Dependence

The simplest true “sequence” of revisions consists of two elements A and B . An important property of the MC-revision function applied to revision sequences is its *order dependence*. In general, the sets $(K_A^*)_B^*$ and $(K_B^*)_A^*$ will differ. To see this, consider a revision model M such that $M \not\models A > \neg B$, but $M \models B > \neg A$ (see Figure 6). When M is revised by A , the dark shaded region of A -worlds become most plausible, and when M_A^* is then revised with B , the subregion of $\min(M, A)$ containing B -worlds becomes most plausible. By the results above, since $\neg B \notin K_A^*$ we have $(K_A^*)_B^* = K_{A \wedge B}^*$.

In contrast, when M is revised with B initially, the light shaded region $\min(M, B)$ becomes most plausible, and when M_B^* is revised with A , the dark region $\min(M, A)$ becomes most plausible, “leaving behind” the set $\min(M, B)$. Again, this is supported by our results of the previous section: since $\neg A \in K_B^*$ we have $(K_B^*)_A^* = K_A^*$. In this case, of course, $K_A^* \neq K_{A \wedge B}^*$, so $(K_B^*)_A^* \neq (K_A^*)_B^*$.

For this reason, it is important to keep in mind that we are dealing with revision *sequences* rather than simple sets of updates. Revision of a belief or extended belief set by some set of new facts will not be order insensitive. In the example above, applying A before B is the same as applying $A \wedge B$ (to K at least), while applying B before A is the same as simply revising by A .

This difference exists for two reasons: first, B is incompatible with A in the sense that $B > \neg A$ holds; second, A is less plausible than B . Because A is less plausible than B , revision by A causes more damage to K than revision by B . A is less expected, or conflicts with K to a greater degree. The belief set K_A^* can be thought of as a “radical shift” in belief from K (we draw a very loose analogy to Kuhn’s notion of *paradigm shift*). When update B is encountered, it must be reconciled with the radically different set K_A^* . If it is consistent with K_A^* then $(K_A^*)_B = K_{A \wedge B}^*$. In contrast, if K is revised with B first the results can be thought of as arising from a “routine” revision (routine in comparison to A) followed by the more “radical” revision A . Even though B has been incorporated in K_B^* , the radical shift to $(K_B^*)_A$ offers no protection for B . A radical shift has little respect for most routine information in a belief set, and B is as vulnerable as any other fact in K_B^* .

Given this interpretation, it is easy to ascertain just when the order of two revisions is irrelevant, that is, when $(K_B^*)_A = (K_A^*)_B$.⁹

Definition 5 Let M be a revision model for belief and extended belief sets K and E . Updates A and B are *mutually compatible* (with respect to M or E) iff $M \not\models A > \neg B$ and $M \not\models B > \neg A$.

Two revisions are mutually compatible just when each is a “routine” revision, relative to the other. This is equivalent to saying A and B are equally plausible: $A \leq_P B$ and $B \leq_P A$.

Proposition 10 If A and B are mutually compatible with respect to K -revision model M , then $(K_B^*)_A = (K_A^*)_B$.

Now, if A and B are incompatible (if one is more plausible than the other) the order of revision is critical. But there are circumstances where the order may still be reversed.

Proposition 11 Let M be a K -revision model such that $M \models A > \neg B$. Then $(K_B^*)_A = (K_A^*)_B$ iff $M \models B > A$.

Since $(K_A^*)_B = K_B^*$ whenever $A > \neg B$, it is easy to see that this holds. We simply observe that if

⁹The order of revision is “irrelevant” only with respect to the objective set K . Rarely will the revised models or extended sets be order insensitive. While $(K_B^*)_A = (K_A^*)_B$ may hold, it will not generally be the case that $(M_B^*)_A = (M_A^*)_B$. We will discuss this below for arbitrary sequences.

$B > A$ then $K_B^* = K_{A \wedge B}^*$ and $(K_B^*)^*_A = K_{A \wedge B}^*$. Of course, the only situation left is that where A and B are “super-incompatible”. In no instance will the order be irreversible.

Proposition 12 *Let M be a K -revision model such that $M \models A > \neg B$ and $M \models B > \neg A$. Then $(K_A^*)^*_B = K_B^*$, and $(K_B^*)^*_A = K_A^*$, and $K_A^* \neq K_B^*$.*

3.4.2 Reduction to Single Revisions

In this section, we examine the possibility of simplifying the revision process. Given a revision sequence A_1, \dots, A_n , we would like to determine the resulting belief set $((K_{A_1}^*)^*_{A_2} \dots)^*_{A_n}$ without having to perform each these n distinct revisions of different belief sets. In fact, we will show that any sequence of revisions can be “reduced” to a single revision. To be more precise we define a *characteristic* sentence for a revision sequence.

Definition 6 Let A_1, \dots, A_n be a revision sequence. We say this sequence is *characterized* by the sentence α iff $((K_{A_1}^*)^*_{A_2} \dots)^*_{A_n} = K_\alpha^*$.

Here we show that every revision sequence has such a characteristic sentence, and that this sentence can itself be determined by the simple unnested conditionals contained in the belief set (or, using the Ramsey test, by “hypothetical” unnested revisions of K).

While mutual compatibility is sufficient to ensure that revision ordering can be reversed, we are typically more concerned with processing updates in the order they are received. When A is processed before B , we have seen that $(K_A^*)^*_B = K_{A \wedge B}^*$ whenever $A > \neg B$ is false. The mutual compatibility of A and B is not important when revisions are processed in order. Rather the *forward compatibility* of B with A determines the content of $(K_A^*)^*_B$, and how it may be achieved with a single revision. If B is forward compatible with A , that is, if $\neg B \notin K_A^*$, then $(K_A^*)^*_B$ reduces to $K_{A \wedge B}^*$. If B is incompatible, then $(K_A^*)^*_B$ reduces to K_B^* . This can be extended in the obvious fashion to arbitrary sequences of revisions.¹⁰

Definition 7 Let M be a K -revision model determining MC-revision function $*$. The revision sequence A_1, \dots, A_n is *forward compatible* with respect to $*$ (or model M) iff $\neg A_{i+1} \notin ((K_{A_1}^*)^*_{A_2} \dots)^*_{A_i}$ for each $1 \leq i < n$.

¹⁰We concentrate on revision of models and belief sets, taking for granted the straightforward connections to extended sets and nested conditionals.

This can be restated as

$$M \not\models A_1 > (A_2 > \cdots (A_i > \neg A_{i+1}))$$

for each $i < n$. Loosely, we say that the sequence is forward compatible (or simply compatible) for K , when $*$, M or E is understood. Clearly, we have the following:

Proposition 13 *If a revision sequence A_1, \dots, A_n is forward compatible, so is each subsequence A_1, \dots, A_i for $i \leq n$.*

An obvious inductive argument, generalizing the case of the two-element sequence, gives us the following:

Theorem 14 *If A_1, \dots, A_n is forward compatible for K , then $((K_{A_1}^*)_{A_2}^* \cdots)_{A_n}^* = K_{A_1 \wedge \cdots \wedge A_n}^*$.*

Corollary 15 *A_1, \dots, A_n is forward compatible for K iff $\neg A_{i+1} \notin K_{A_1 \wedge \cdots \wedge A_i}^*$ for each $i < n$.*

Thus (by Corollary 15) testing for compatibility can be reduced to the application of single revisions to the belief set K , or testing simple conditionals $A_1 \wedge \cdots \wedge A_i > A_{i+1}$. Iterated revision or nested conditionals are not required to test for compatibility, nor (by Theorem 14) are they needed to compute the result of such a revision sequence applied to K . Computing compatible revision sequences is a straightforward extension of the case of two compatible revisions, and is reducible to a single revision, the conjunction of the elements.

In the *incompatible* instance, the two-element sequence was again reducible to a single revision: $(K_A^*)_B^* = K_B^*$ when $A > \neg B$. Accounting for an incompatible revision after a longer sequence of compatible revisions, however, is not so straightforward. Suppose we have a revision sequence, A_1, \dots, A_{n+1} , where A_1, \dots, A_n is compatible but the longer sequence is not. The analogy to the two-element case breaks down here, for in general $((K_{A_1}^*)_{A_2}^* \cdots)_{A_{n+1}}^* \neq K_{A_{n+1}}^*$. Unfortunately, earlier revisions leave a residual trace on the structure M , as shown in Figure 7. While $((K_{A_1}^*)_{A_2}^* \cdots)_{A_n}^* = K_{A_1 \wedge \cdots \wedge A_n}^*$, most certainly $((M_{A_1}^*)_{A_2}^* \cdots)_{A_n}^*$ is different from $M_{A_1 \wedge \cdots \wedge A_n}^*$. The history of the belief set K , the process by which it was formed, plays a vital role in future revisions.

In this model, the original structure of the K -revision model lies primarily above the dashed line. Revision by A_1 moved the minimal A_1 -worlds to the cluster just below the dashed line. Since A_2 is compatible with A_1 , revision by A_2 removed the A_2 -worlds from this new cluster (leaving behind $A_1 \wedge \neg A_2$ -worlds) and moved them to the second cluster below the line. This process was repeated up to A_n , resulting in a final minimal cluster of $A_1 \wedge \cdots \wedge A_n$ -worlds (shaded in Figure 7). The same minimal cluster would have been formed had K simply been revised by $A_1 \wedge \cdots \wedge A_n$, but the *sequence*




Figure 7: A sequence of compatible revisions.

of revisions has a drastically different effect on the structure of M , leaving a number of intermediate clusters in its wake. Simply revising by the conjunction $A_1 \wedge \dots \wedge A_n$ would have caused only the shaded cluster to form below the dashed line.

Referring still to Figure 7, if a subsequent revision A_{n+1} is not compatible with the sequence A_1, \dots, A_n , then $((K_{A_1}^*)_{A_2}^* \dots)_{A_n}^* \vdash \neg A_{n+1}$. That is, no A_{n+1} -worlds can be found in the shaded minimal cluster. When $((M_{A_1}^*)_{A_2}^* \dots)_{A_n}^*$ is revised with A_{n+1} , it need not be the case that $K_{A_{n+1}}^*$ results. If the minimal A_{n+1} -worlds are found in some cluster below the dashed line, that is, if $K_{A_1}^* \not\vdash \neg A_{n+1}$, then $((K_{A_1}^*)_{A_2}^* \dots)_{A_{n+1}}^*$ will not usually equal $K_{A_{n+1}}^*$.

So exactly where will minimal A_{n+1} -worlds be found in $((M_{A_1}^*)_{A_2}^* \dots)_{A_n}^*$ and what sentences will be in $((K_{A_1}^*)_{A_2}^* \dots)_{A_{n+1}}^*$? Suppose the sequence A_1, \dots, A_k, A_{n+1} is incompatible. This means there can be no A_{n+1} -worlds in the cluster formed when $((K_{A_1}^*)_{A_2}^* \dots)_{A_{k-1}}^*$ is revised by A_k , that is, the cluster of $A_1 \wedge \dots \wedge A_k$ -worlds. Of course, this implies that there can be no A_{n+1} -worlds in any lower clusters formed by the subsequent revisions A_{k-1} through A_n , since each of these is compatible and will only “select” worlds from this set of $\neg A_{n+1}$ -worlds. Conversely, if A_1, \dots, A_k, A_{n+1} is a compatible sequence, there must be some A_{n+1} -worlds among the the cluster of $A_1 \wedge \dots \wedge A_k$ -worlds representing $((K_{A_1}^*)_{A_2}^* \dots)_{A_k}^*$.

It now becomes clear that the minimal set of A_{n+1} -worlds must be located in the cluster of worlds “labeled” $A_1 \wedge \dots \wedge A_k \wedge \neg A_{k+1}$, where $k \leq n$ is maximal among the set of i such that A_1, \dots, A_i, A_{n+1}

is a compatible sequence of revisions. Since A_1, \dots, A_k is compatible, $((K_{A_1}^*)_{A_2}^* \dots)_{A_k}^* = K_{A_1 \wedge \dots \wedge A_k}^*$, and the minimal A_{n+1} -worlds in this set are captured by $((K_{A_1}^*)_{A_2}^* \dots)_{A_k}^*_{A_{n+1}}$. But this is equivalent to $K_{A_1 \wedge \dots \wedge A_k \wedge A_{n+1}}^*$ since A_{n+1} is compatible with the rest of the sequence.

Definition 8 Let A_1, \dots, A_n be a revision sequence. Update A_k ($k > 1$) is *incompatible* in this sequence iff $\neg A_k \in ((K_{A_1}^*)_{A_2}^* \dots)_{A_{k-1}}^*$. A sequence is incompatible iff it contains at least one incompatible member.

Proposition 16 Sequence A_1, \dots, A_n is incompatible iff it is not forward compatible.

If the belief set $((K_{A_1}^*)_{A_2}^* \dots)_{A_n}^*$ is characterized by a single revision α but A_{n+1} is incompatible, then the sentence representing $((K_{A_1}^*)_{A_2}^* \dots)_{A_{n+1}}^*$ is clearly not equivalent to $\alpha \wedge A_{n+1}$. The situation we have described above, where only the last revision in a sequence is incompatible, is easily characterized.

Theorem 17 Let A_1, \dots, A_{n+1} be an incompatible sequence such that A_1, \dots, A_n is compatible. Let k be the maximal element of

$$\{i \leq n : \neg A_{n+1} \notin ((K_{A_1}^*)_{A_2}^* \dots)_{A_i}^*\}$$

Then $((K_{A_1}^*)_{A_2}^* \dots)_{A_{n+1}}^* = K_{A_1 \wedge \dots \wedge A_k \wedge A_{n+1}}^*$.

Thus a sequence with one incompatible revision as its last element is reducible to a single revision. Notice that when the set of revisions compatible with A_{n+1} is empty, when this maximal element k does not exist, we have $((K_{A_1}^*)_{A_2}^* \dots)_{A_{n+1}}^* = K_{A_{n+1}}^*$. This is directly analogous to the two-element case, since there is no subsequence compatible with A_{n+1} . In Figure 7 this occurs exactly when there are no A_{n+1} -worlds below the dashed line; that is, when $\neg A_{n+1} \in K_{A_1}^*$.

It should be quite clear that subsequent compatible revisions, A_{n+2} and so on, should be treated as previously and simply “conjoined” to $((K_{A_1}^*)_{A_2}^* \dots)_{A_{n+1}}^*$.

Proposition 18 Let A_1, \dots, A_n be a revision sequence with one incompatible element A_k , and let $j < k$ be the maximal compatible revision for A_k (as in Theorem 17). Then $((K_{A_1}^*)_{A_2}^* \dots)_{A_n}^* = K_{A_1 \wedge \dots \wedge A_j \wedge A_k \wedge \dots \wedge A_n}^*$.

The final piece in the puzzle is the process by which subsequent incompatible revisions are achieved. Consider a revision sequence A_1, \dots, A_n where A_k is incompatible and has as its *most recent compatible revision* A_j (the element defined in Theorem 17). This situation is illustrated in Figure 8. Now suppose update A_{n+1} is incompatible, so that no A_{n+1} -worlds are located in the minimal




Figure 8: A sequence with an incompatible revision.

cluster. This occurs when $K_{A_1 \wedge \dots \wedge A_j \wedge A_k \wedge \dots \wedge A_n}^* \vdash \neg A_{n+1}$. Again, to find the minimal A_{n+1} -worlds in this structure we must look for the most recent compatible revision in the sequence A_1, \dots, A_n . If A_i is this update then $((K_{A_1}^*)_{A_2}^* \dots)_{A_{n+1}}^*$ is identical to $((((K_{A_1}^*)_{A_2}^* \dots)_{A_i}^*)_{A_{n+1}}^*)$. Furthermore, since A_1, \dots, A_n has only one incompatible revision, by Theorem 14 and Proposition 18, we have $((K_{A_1}^*)_{A_2}^* \dots)_{A_i}^* = K_\alpha^*$ for some sentence α . Thus $((K_{A_1}^*)_{A_2}^* \dots)_{A_{n+1}}^* = K_{\alpha \wedge A_i}^*$. It is interesting, however, to examine the various situations that arise with respect to the occurrence of this most recent compatible revision A_i .

First, consider $i \geq k$. In this case, A_{n+1} is compatible with the previous incompatible revision A_k . The set of minimal A_{n+1} -worlds lies below the third dashed line, among those worlds representing $K_{A_1 \wedge \dots \wedge A_j \wedge A_k}^*$. In this circumstance we have $((K_{A_1}^*)_{A_2}^* \dots)_{A_{n+1}}^* = K_{A_1 \wedge \dots \wedge A_j \wedge A_k \wedge \dots \wedge A_i \wedge A_{n+1}}^*$.

Second, consider $j < i < k$. Clearly, A_{n+1} is incompatible with the incompatible revision A_k , but is compatible with the sequence A_1, \dots, A_j . This case is rather interesting for we cannot simply “backtrack” within our representative revision for A_1, \dots, A_n . Because $((K_{A_1}^*)_{A_2}^* \dots)_{A_n}^* = K_{A_1 \wedge \dots \wedge A_j \wedge A_k \wedge \dots \wedge A_n}^*$, one might think we could simply “back up” to the most recent compatible revision A_j in this representation and arrive at $K_{A_1 \wedge \dots \wedge A_j}^*$. However, this ignores the consistent revisions between A_j and A_k (between the second and third dashed lines) that were “left behind” when the incompatible revision A_k was incorporated. The minimal A_{n+1} -worlds lie in this region and this

must be taken into account. Indeed, since A_1, \dots, A_i is compatible we have $((K_{A_1}^*)_{A_2}^* \dots)_{A_{n+1}}^* = K_{A_1 \wedge \dots \wedge A_i \wedge A_{n+1}}^*$.

Finally, consider $i \leq j$. Since $((K_{A_1}^*)_{A_2}^* \dots)_{A_{n+1}}^* = K_{A_1 \wedge \dots \wedge A_i}^*$, this case is much like the one just mentioned. It is distinguished by the fact that the single revision $A_1 \wedge \dots \wedge A_i$ that represents it is a proper subsequence of the update $A_1 \wedge \dots \wedge A_j \wedge A_k \wedge \dots \wedge A_n$ representing $((K_{A_1}^*)_{A_2}^* \dots)_{A_n}^*$.

These considerations can be generalized to accommodate any number of incompatible revisions. Furthermore, they provide a constructive means (described inductively) of reducing any sequence of revisions of belief set K to a single revision of K , and demonstrate (through compatibility testing) how that revision can itself be determined using only single (non-iterated) revisions of K . Finally we shall see that, although the inductive description indicates a dependence of the characterization of $((K_{A_1}^*)_{A_2}^* \dots)_{A_n}^*$ on the characterization of $((K_{A_1}^*)_{A_2}^* \dots)_{A_i}^*$, for each $i < n$, it is only necessary to keep track of those sentences that characterize incompatible revisions.

Definition 9 Let A_1, \dots, A_n be a revision sequence with c incompatible updates. We use the strictly increasing function $\sigma : \{1, \dots, c\} \mapsto \{1, \dots, n\}$ to denote these incompatible elements: $A_{\sigma(1)}, \dots, A_{\sigma(c)}$. For each $1 < j \leq n$, the *maximal consistent incompatible revision* for A_j is A_k , where

$$k = \max\{\sigma(i) : \sigma(i) < j \text{ and } \neg A_j \notin ((K_{A_1}^*)_{A_2}^* \dots)_{A_{\sigma(i)}}^*\}$$

For each $1 < j \leq n$, the *most recent compatible revision* for A_j is A_k , where

$$k = \max\{i : i < j \text{ and } \neg A_j \notin ((K_{A_1}^*)_{A_2}^* \dots)_{A_i}^*\}$$

If either of these sets is empty, we take the most recent or maximal incompatible revision for A_j to be \top .

Lemma 19 Let A_1, \dots, A_n be a revision sequence such that each proper subsequence A_1, \dots, A_i is characterized by some sentence $s(A_i)$. Then A_1, \dots, A_n is characterized by $s(A_k) \wedge A_n$, where A_k is the most recent compatible revision for A_n . In other words, $((K_{A_1}^*)_{A_2}^* \dots)_{A_n}^* = K_{s(A_k) \wedge A_n}^*$.

This leads to the main result of this section.

Theorem 20 For any revision sequence A_1, \dots, A_n , there is some subset of these updates $S \subseteq \{A_1, \dots, A_n\}$ such that $((K_{A_1}^*)_{A_2}^* \dots)_{A_n}^* = K_A^*$ and $A = \bigwedge S$.

Corollary 21 For any revision sequence A_1, \dots, A_n , there is some subset of these updates $S \subseteq \{A_1, \dots, A_n\}$ such that $((E_{A_1}^*)_{A_2}^* \dots)_{A_{n-1}}^* \models A_n > B$ iff $E \models A > B$, and $A = \bigwedge S$.

This result is given its constructive character by Theorem 14, but it seems to suggest that one must keep track of a characteristic sentence $s(A_i)$ for each revision A_i . In fact, the critical sentences are only those corresponding to incompatible revisions, $s(A_{\sigma(i)})$. Every other characteristic sentence $s(A_i)$ is simply the conjunction of subsequent revisions to the *most recent* incompatible revision.

Theorem 22 *Let A_1, \dots, A_n be a revision sequence with c incompatible updates represented by σ . For each $1 \leq k \leq n$, if A_k is a compatible revision, then $((K_{A_1}^*)_{A_2}^* \cdots)_{A_k}^* = K_{s(A_k)}^*$, where: a) $s(A_k) = s(A_{\sigma(i)}) \wedge A_{\sigma(i)+1} \wedge \cdots \wedge A_k$; b) $\sigma(i) < k < \sigma(i+1)$, or $\sigma(i) < k$ if $i = c$; and c) $s(A_{\sigma(i)})$ characterizes subsequence $A_1, \dots, A_{\sigma(i)}$.*

We can provide a similar characterization of incompatible revisions, but these must be in terms of the *maximal consistent incompatible revision* rather than the most recent incompatible revision.

Proposition 23 *Let A_1, \dots, A_n be a revision sequence with c incompatible updates represented by σ . Let $A_{\sigma(i)}$ be the maximal consistent incompatible revision for A_k . If A_j is the most recent compatible revision for A_k , then $\sigma(i) \leq j < \sigma(i+1)$.*

Theorem 24 *Let A_1, \dots, A_n be a revision sequence with c incompatible updates represented by σ . For each $1 \leq k \leq n$, if A_k is an incompatible revision, then $((K_{A_1}^*)_{A_2}^* \cdots)_{A_k}^* = K_{s(A_k)}^*$, where: a) $s(A_k) = s(A_{\sigma(i)}) \wedge A_{\sigma(i)+1} \wedge \cdots \wedge A_j \wedge A_k$; b) $A_{\sigma(i)}$ is the maximal incompatible revision for A_k ; c) A_j is the most recent compatible revision for A_k ; and d) $s(A_{\sigma(i)})$ characterizes subsequence $A_1, \dots, A_{\sigma(i)}$.*

Taken together, these theorems show that one may implement a procedure that tests for membership of B in a multiply-revised belief set $((K_{A_1}^*)_{A_2}^* \cdots)_{A_n}^*$ using only an ‘‘oracle’’ that answers requests of the form ‘‘Is $\beta \in K_\alpha^*$?’’ for $\alpha, \beta \in \mathbf{L}_{CPL}$. Furthermore, the characteristic sentences $s(A_j)$ that need to be recorded are only those of the form $s(A_{\sigma(i)})$ where $A_{\sigma(i)}$ is some incompatible revision. The core of this algorithm is provided in Figure 9. It takes as input a revision sequence A_1, \dots, A_n and computes the characteristic sentence A such that $K_A^* = ((K_{A_1}^*)_{A_2}^* \cdots)_{A_n}^*$. The algorithm is incremental in the sense that a subsequent revision A_{n+1} requires only one further execution of the main ‘‘for loop’’. To ask whether $B \in ((K_{A_1}^*)_{A_2}^* \cdots)_{A_n}^*$, one simply computes the required characteristic sentence A and asks if $B \in K_A^*$.

3.5 Information Preservation

One desideratum of any model of the revision process is the minimization of information loss. In the AGM framework for single revisions, postulate (R4) ensures that no beliefs from set K are retracted

Input: revision sequence A_1, \dots, A_n ; revision function $*$ applicable to belief set K
 Output: characteristic update $A \leftarrow K_A^* = ((K_{A_1}^*)_{A_2}^* \dots)_{A_n}^*$
 Variables: L — list of updates
 IL — list of incompatible revisions: elements have form $\langle ind, S \rangle$, where A_{ind} is incompatible and subsequence A_1, \dots, A_{ind} is characterized by sentence S
 S — characteristic sentence A built up here

```

Initialize  $A \leftarrow \top$ ;  $L \leftarrow \top$ ;  $IL \leftarrow \langle 0, \top \rangle$ 
for  $i = 1$  to  $n$       ;; one loop for each update  $A_i$ 
   $L \leftarrow L + A_i$ 
  if  $\neg A_i \notin K_A^*$  then      ;;  $A_i$  is compatible
     $A \leftarrow A \wedge A_i$ 
  else      ;;  $A_i$  is incompatible
    for  $j = length(IL)$  to 1      ;; find maximal incompatible revision consistent with  $A_i$ 
       $S \leftarrow IL(j).S$ 
       $idx \leftarrow IL(j).ind$ 
      if  $\neg A_i \notin K_S^*$  then      ;; found maximal incompatible revision
        if  $x = length(IL)$  then      ;;  $top$  is max location of most recent compatible revision
           $top \leftarrow i - 1$ 
        else  $top \leftarrow IL(j).ind$ 
        end if
      break
    end if
  end for
  for  $k = idx + 1$  to  $top - 1$       ;; find most recent compatible revision
    if  $\neg A_i \in K_{S \wedge L(k)}^*$  then
      break
    else  $S \leftarrow S \wedge L(k)$ 
    end if
  end for
   $A \leftarrow S$ 
   $IL \leftarrow IL + \langle i, S \rangle$ 
end if      ;; end of incompatible revision
end for
  
```

Figure 9: Algorithm to compute characteristic update A for revision sequence A_1, \dots, A_n .

if none need to be. An arbitrarily long sequence of *consistent* revisions will simply make a belief set larger, resulting (in the limit) in a *complete* belief set. Eventually, we will believe either A or $\neg A$ for each objective sentence A . Once an inconsistent revision is processed, the AGM model asserts that, among those beliefs that could be given up to accommodate A , only those that are least entrenched are retracted. Unfortunately, no constraints are levied on the relative degrees of entrenchment of the members of the new belief set K_A^* . Much of the information associated with K , including relative degrees of entrenchment, is lost.

The MC-revision model is unique in the sense that this type of information is preserved. When a consistent revision A is processed, the minimal cluster of a revision model M is divided, resulting in two smaller clusters. A subsequent consistent revision breaks the minimal cluster of M_A^* , and so on. Eventually (supposing some finite language) after a suitable sequence of consistent revisions, we end up with a minimal cluster containing one possible world, representing a complete belief set.

The process is not altogether different for inconsistent revisions. Rather than dividing the minimal cluster of K -worlds, the cluster containing $\min(M, A)$ is divided. However, any information implicit in previous revisions that caused the formation of various clusters in M is preserved. All other clusters remain undisturbed. In no case will a revision cause any cluster to “grow” or “lose information.” Typically, a revision will cause the number of clusters to increase by one by splitting $\min(M, A)$ from a cluster, thus “shrinking” a cluster and ensuring information gain. Only in certain cases will no new clusters be formed, for example, when we revise K by A where $A \in K$. Thus, we have a continual “thinning” of our revision models.

Proposition 25 *Let revision model $M_A^* = \langle W, \leq, \varphi \rangle$ be the MC-revision of M by A . If $\mathcal{C} \subseteq W$ is a cluster in M_A^* , then $\mathcal{C} \subseteq \mathcal{C}'$ for some cluster \mathcal{C}' in M . Furthermore, if $\min(M, A) \subset \mathcal{C}'$ for some cluster \mathcal{C}' in M , then the set of clusters in M_A^* consists of the set of clusters in M distinct from \mathcal{C}' together with $\min(M, A)$ and $\mathcal{C}' \perp \min(M, A)$. If $\min(M, A) = \mathcal{C}'$ for some cluster \mathcal{C}' in M , then the set of clusters in M_A^* is identical to that in M .*

Thus, a revision sequence A_1, \dots, A_n causes a non-decreasing change in “information” in a belief state. No belief set farther along in the revision sequence can be smaller than an earlier belief set.

Corollary 26 *Let M be a K -revision model and A_1, \dots, A_n a revision sequence. If $i \leq j$ then $((K_{A_1}^*)_{A_2}^* \cdots)_{A_j}^* \not\subseteq ((K_{A_1}^*)_{A_2}^* \cdots)_{A_i}^*$.*

This suggests that, as we process a revision sequence, our revision model becomes more and more *informationally complete* in the following sense.

Definition 10 A revision model M is *informationally complete* iff each cluster in M is a singleton set.

Proposition 27 *Let M be an informationally complete K -revision model. Then K_A^* is a complete theory for any $A \in \mathbf{L}_{CPL}$. Furthermore, M_A^* is informationally complete.*

Corollary 28 *If M is an informationally complete K -revision model then K is a complete theory.*

Proposition 25 ensures furthermore that M_A^* will be informationally complete whenever M is. It also shows that if we restrict our attention to languages with a finite number of propositional variables, we can eventually attain informational completeness.¹¹

Proposition 29 *Let M be a revision model and A_1, \dots, A_n a revision sequence such that every satisfiable sentence α is represented in this sequence; that is, $A_i \vdash \alpha$ for some A_i . Then $((M_{A_1}^*)_{A_2}^* \cdots)_{A_n}^*$ is informationally complete.*

An example of such a revision sequence, for a language with p propositional variables, would be the set of the 2^p complete (truth-functionally distinct) conjunctions of literals, each capturing a possible world or valuation, and each causing the corresponding world to be broken off into its own cluster. From that point on all revisions will result in complete belief sets since the set of most plausible A -worlds will have one element for each revision A .

4 Reasoning with Revision Sequences

The algorithm for reducing nested conditionals requires some method of establishing the truth of $B \in K_A^*$ for $A, B \in \mathbf{L}_{CPL}$. Typically, as is the case in most reasoning tasks, our premises do not provide us with complete knowledge, and we can only hope to derive as much as possible, leaving certain gaps in our knowledge. It is not reasonable to expect a completely specified revision function $*$, or equivalently, a complete set of conditionals containing one of $A > B$ or $\neg(A > B)$, among our set of premises. Naturally, the algorithm can then easily be modified to ask whether $A > B$ or $\neg(A > B)$ is *provable* from a given premise set.¹² If either is the case, the algorithm can proceed, having an answer to the query $B \in K_A^*$. If neither is derivable from the premises, then the algorithm

¹¹We assume that models contain no “duplicate” worlds having the same associated valuation. If so, our informational completeness in the following result is the type that ensures that K_A^* is complete for all A .

¹²This requires a logic in which conditionals can be expressed and reasoned with, e.g., [9].

must halt unsuccessfully, or proceed as if either could be the case.¹³ As should be expected when reasoning with incomplete information, the answer “unknown” must be returned for certain queries.

While the revision function $*$ is usually specified only partially by means of conditional premises (as well as direct statements of plausibility and entrenchment), there may be circumstances when the revision function is completely known. The problem changes from that of reasoning with incomplete information to that of specifying complete information in a reasonable manner. We cannot expect one to specify a complete conditional theory explicitly containing $A > B$ or $\neg(A > B)$ for each objective A and B , for we do not want to be forced to reason with infinite sets of premises. Even for finite languages with p atomic variables, where we require only a conditional or its negation for semantically distinct A and B , we are forced to reason with 2^{p+1} premises.¹⁴

There are cases, though, where a revision function can be captured finitely, and often with a manageable number of sentences. Often when a set of conditional premises is given, we have in mind a certain intended model (e.g., *preferred* entrenchment orderings are described in [26, 4, 28]), and these can often be represented finitely [4]. In general, a revision function can be compactly represented if there is a corresponding revision model that is “well-behaved” in the following sense.

Definition 11 A revision model $M = \langle W, \leq, \varphi \rangle$ is *finitely specifiable* iff $W = \cup_{i \leq n} \mathcal{C}_i$, where each \mathcal{C}_i is a cluster in M and $\mathcal{C}_i = \|S_i\|$ for some sentence $S_i \in \mathbf{L}_{CPL}$.

In other words, M is well-behaved if it consists of a finite number of clusters, each corresponding to some finite classical theory or sentence. It is easy to verify that if M is finitely specified by sentences S_1, \dots, S_n in the definition above, then these sentences are “mutually exclusive” and “exhaustive”; that is, $\vdash S_i \supset \neg S_j$ if $i \neq j$, and $\vdash S_1 \vee \dots \vee S_n$. We assume that S_1 characterizes the minimal (most plausible) cluster of M , S_2 the next most plausible and so on. We denote the model described above by $FSM(S_1, \dots, S_n)$.¹⁵

Definition 12 A revision function $*$ is *finitely specifiable* (for belief set K) iff there is some finitely specifiable K -revision model $M = FSM(S_1, \dots, S_n)$ such that $*^M = *$.

Given a finitely specifiable revision function for belief set K (in which case K is also finitely

¹³It often makes sense to continue, for a subsequent revision may be incompatible with previous revisions, and failing to reduce an earlier revision may have no effect on the effort to reduce an entire sequence.

¹⁴Alternatively, one could provide a complete set of entrenchment sentences, specifying the relative degrees of entrenchment of each pair of sentences: $A \leq_E B$ and/or $B \leq_E A$ for each $A, B \in \mathbf{L}_{CPL}$. This would allow the derivation of every simple conditional or its negation, and seems to be what Gärdenfors and Makinson [17, 25] have in mind in their presentation of *expectation inference* (see also [16]).

¹⁵Such a model can be described compactly in the logic CO^* as a function of the sentences S_1, \dots, S_n [6].

specified by S_1), we can use the sequence of sentences S_1, \dots, S_n to represent the model M , as well as determine the truth of every simple conditional sentence.

Theorem 30 *Let $S_1, \dots, S_n \in \mathbf{L}_{CPL}$ be such that $\vdash S_i \supset \neg S_j$ if $i \neq j$, and $\vdash S_1 \vee \dots \vee S_n$. For all $A, B \in \mathbf{L}_{CPL}$ either*

$$\begin{aligned} FSM(S_1, \dots, S_n) \models A > B; & \quad \text{or} \\ FSM(S_1, \dots, S_n) \models \neg(A > B) \end{aligned}$$

Fortunately, the ability to finitely specify a revision model is not disturbed by MC-revision. If $FSM(S_1, \dots, S_n)$ characterizes a revision model M , then M_A^* is formed by simply “dividing” the minimal cluster consistent with A in two clusters: $S_i \wedge A$ becomes most plausible; and $S_i \wedge \neg A$ replaces cluster S_i .

Theorem 31 *Let revision model $M = FSM(S_1, \dots, S_n)$ and let S_k be the minimal sentence in this set consistent with A ; that is, $S_k \not\vdash \neg A$ and $S_i \vdash \neg A$ if $i < k$. Then M_A^* is characterized by*

$$\begin{aligned} FSM(S_k \wedge A, S_1, \dots, S_{k-1}, S_k \wedge \neg A, S_{k+1}, \dots, S_n) & \quad \text{if } S_k \not\vdash A \\ FSM(S_k, S_1, \dots, S_{k-1}, S_{k+1}, \dots, S_n) & \quad \text{if } S_k \vdash A \end{aligned}$$

Thus we can completely specify a revision function with a compact set of premises and use this premise set to reduce nested conditional queries to simple conditionals and then establish the truth of these simple conditionals. Furthermore, we can explicitly revise a model and retain a compact representation. It also becomes clear that testing for the truth of a simple conditional in such a revision model is reducible to a simple propositional reasoning task.

Relative to propositional satisfiability tests, this “algorithm” is relatively efficient, for the minimal A -consistent cluster or sentence can be found using a linear search technique. Thus we need only perform $O(n)$ satisfiability tests to determine the truth of $A > B$. The complexity of these tests will depend on the size of the theories S_i and their structure. For instance, if the S_i are Horn theories, this test will be linear in the size of the theory.

5 Concluding Remarks

We have presented a model for iterated revision that captures sequences of objective revisions. The hallmark of MC-revision is the preservation of subjective information, such as conditional beliefs and entrenchment, throughout such a sequence. Because this information is retained across belief sets, simply knowing how a revision function behaves on single revisions is sufficient to characterize the

results of any sequence of revisions. If we adopt the Ramsey test for acceptance of conditionals, this demonstrates that right-nested conditionals are equivalent to simple unnested conditionals, and that this reduction can be performed using only knowledge of unnested conditionals.

The MC-revision model has other compelling properties as well. Because it preserves conditional information, the sequence of belief sets corresponding to a revision sequence can never decrease in propositional information content. It preserves nice properties of the revision models and functions being revised as well, for example, informational completeness and finite specifiability. In cases where a revision function is described simply (e.g., by a set of propositional sentences), the characterization of MC-revision is also easily computed.

5.1 Fixed Sentence Orderings

Several other proposals have been put forth that extend the AGM theory (or variants of it) to deal with iterated revision, but these all have a somewhat different nature. In particular, the question of minimizing the change in the conditional beliefs of an agent appears to have been unaddressed. Safe contraction [3], generalized epistemic entrenchment [29] and the probabilistically motivated system of Schlechta [30] each take a similar approach to the problem: each assumes the existence of a “global” ordering of entrenchment over all sentences in the language. For any belief set K the appropriate revision function is immediately available, and iteration of the process requires no additional apparatus. These models have the rather severe drawback that any objective belief set K is associated with a unique revision function. In particular, K uniquely determines the set of conditional beliefs. So while the ordering of entrenchment can change as K evolves, these changes are predetermined by the global ordering; the entrenchment ordering associated with K cannot depend on how K evolved. This stands in contrast with our model, where the ordering itself can be revised. Hansson [21] makes a similar observation, that many systems of iterated change fix the connection between objective and extended sets. Hansson proposes that instead a revision method be associated with *belief bases* rather belief sets. Thus, the same belief set may be revised in different ways if it is generated by different bases in each instance. While clearly the issue of base revision is crucial, it is not a problem we address here. In particular, the revision of conditional beliefs depends on existing conditional beliefs and, as in the AGM theory, there need not be strong ties to the underlying belief base.¹⁶

¹⁶Of course, such ties can be added. Work on base contraction [20, 21, 13] can be viewed in this light.

5.2 Nonminimal Change of Conditionals

The proposal most similar in structure and spirit to ours is that of Spohn [31], and indeed inspired this investigation of MC-revision. Spohn's use of ordinal conditional functions to represent belief states is much like the possible worlds model we have used, except that "clusters" are given explicit ordinal rankings. Spohn's notion of conditionalization on such models assigns to a proposition A a new "degree of belief" or ordinal ranking (which should be thought of as the degree of entrenchment). This is analogous to revising by A , and is achieved by "shifting" the rank of all A -worlds so that the minimal A -worlds have an ordinal ranking that is lower (or more plausible) than that of any $\neg A$ -world by the specified degree of belief. The key distinction between Spohn's approach and ours is that MC-revision by A requires the shift of only the *minimal* A -worlds. Spohn's approach is more reminiscent of probabilistic conditionalization, and if we applied our truth conditions for conditionals to OCF update, we would see a much larger change in conditional beliefs in general.

Spohn's approach is certainly more sophisticated, and in many ways is more compelling than MC-revision. In particular, when evidence A is brought to light *all* A -worlds become somewhat more plausible. However, this κ -conditionalization cannot be applied in our setting, for it requires that degrees of entrenchment be given an explicit quantitative interpretation and that the strength of the evidence be quantified on the same scale.¹⁷ MC-revision is applicable in settings where only qualitative entrenchment information is available. Indeed, it is hard to imagine what possible shifts of nonminimal A -worlds are justifiable if only the relative plausibility of worlds is known. Furthermore, the MC-model requires the simple shifting of one partial cluster, a more computationally feasible mechanism for revising a ranking than shifting all A -worlds. Thus, MC-revision may prove to be a useful approximation method in quantitative settings as well. The utility of MC-revision is especially noticeable if new evidence tends not to contradict earlier evidence.

Darwiche and Pearl [12] propose two postulates extending the eight AGM postulates that govern iterated revision. These impose rather mild constraints on the conditionals (or ordering) that results from a revision; indeed, both Spohn's OCF model (or more accurately a suitable qualitative abstraction thereof) and MC-revision satisfy these postulates. Hence, their proposal cannot be used to *fix* a particular revision method (as the AGM postulates cannot). What is missing, and what will allow the exploration and specification of iterated revision methods, is a means to express the relative entrenchment of conditional sentences. This has received little attention in the literature, but several possibilities exist (e.g., using nested conditional sentences as in Section 4 without the specific

¹⁷Spohn's use of OCFs requires that the magnitudes of the rankings are meaningful.

constraints of MC-revision; see also [10, 11]). With this in place, a general theory such as that of Darwiche and Pearl can be used to construct revision functions.

Darwiche and Pearl also suggest that the constraints of MC-revision are too strong, for they give little “priority” to incoming evidence. An example of this is the following: we believe that Mary won’t come to the party, \overline{M} , and that Ted won’t either, \overline{T} . We necessarily believe the rule $\overline{M} \Rightarrow \overline{T}$. But Ted may or may not come should Mary decide to. We are first told that M (and now are uncommitted to T or \overline{T}) and then that T . Some time later we are informed that we were correct initially and Mary isn’t coming \overline{M} after all. Should we still believe that Ted is coming? This depends on the relative importance or weight of rules versus plain beliefs. If we accept the rule $\overline{M} \Rightarrow \overline{T}$ more strongly than the evidence T , then we should return to our original state of belief $(\overline{M}, \overline{T})$ as MC-revision suggests. This is certainly plausible given this scenario. However, if priority is to be given to the evidential report, then perhaps we should abandon our rule. One way MC-revision can be interpreted is as a system in which ultimate priority is given to existing rules. Darwiche and Pearl seem to come down on the side of evidential superiority (and argue that one should keep believing T) — a view that is just as extreme. Again, a general theory for specifying the relative importance of conditionals and propositional beliefs will allow one to adopt less extreme revision methods. With such a theory in place, MC-revision can still be viewed as a computationally attractive approximation method for such “ideal” revision.¹⁸

5.3 Nested Conditionals

It is also interesting to note that a nested conditional $A > (B > C)$ is often equivalent to $A \wedge B > C$ (whenever A and B are compatible). It has been suggested by a number of people that nested conditionals should be reduced to unnested conditionals with all antecedents conjoined to form a single antecedent, Adams [1] and Levi [23] among them. This reduction is sanctioned by the MC-revision model for nested conditionals with compatible antecedents. But as described in Section 4, this is not the case when incompatible antecedents are present. For instance, $A > (B > C)$ reduces simply to $B > C$ when A and B are incompatible. As suggested by Levi, this reduction may seem inappropriate in normal linguistic usage, for the nested conditional seems to imply that A should continue to hold when $B > C$ is evaluated with the Ramsey test, thus suggesting the reduction to $A \wedge B > C$. Indeed, the MC-revision model cannot account for this circumstance when incompatible

¹⁸The party example is a rewording of the example in [12] of observing a red bird. That example is somewhat misleading for it describes an instance of adding new beliefs (about new domain objects) rather than revising old beliefs.

antecedents are involved. In contrast, our model provides non-trivial acceptance conditions for nested conditionals such as $A > (\neg A > B)$.

Levi [23] is rather critical of the enterprise of determining truth or acceptance conditions for nested conditionals, or even allowing conditionals to be part of a belief set. He offers the opinion that an element of a belief set or *corpus* ought to be practicable as a standard of serious possibility. Since conditionals do not perform this function, they are not accorded the status of beliefs. While certainly their role differs from that of garden-variety propositional beliefs, they perform an indispensable function in the process of deliberation. Conditionals suggest hypothetical possibilities to an agent and aid an agent in changing its mind. If conditionals are not “beliefs,” they at least must be representations of an agent’s revision policies; and such policies themselves must be the objects of revision from time to time. Whether or not such policies are called beliefs, a model of conditional revision is crucial.

A few words on triviality are in order, as well. The Gärdenfors triviality result ensures that no meaningful notion of belief revision can be applied to belief sets that contain conditionals in accordance with the Ramsey test and satisfy all eight AGM postulates.¹⁹ It is clear that MC-revision is a revision function that can be applied to conditional belief sets in a nontrivial fashion, giving rise to new conditional belief sets.²⁰ Furthermore, by Proposition 5, MC-revision applied to K satisfies all AGM postulates. Notice, however, that triviality is avoided because the MC-revision function satisfies the postulates when *restricted* to K . Thus, the postulates *do* hold when we consider the “objective component” of the revision, but not when considered with respect to an extended set E . Upon reflection, however, it becomes clear that the postulates, in particular postulates pertaining to consistent revisions such as (R4), lose much of their appeal when applied to conditional belief sets. As argued convincingly by Levi [23] and Rott [27], among others, the proper escape from “triviality” requires relaxing the requirements of the AGM theory when discussing non-objective belief sets.

5.4 Extensions

There are a number of interesting avenues that remain to be explored. This model is restricted to right-nested conditionals, or propositional revisions. In general, we want to allow revision of a knowledge base with conditional information, or statements of entrenchment as well. A fully general model of this type is currently under investigation. Some preliminary results may be found in [10, 11]. Other models that allow arbitrary nesting include Hansson’s [21] base-set model.

¹⁹Indeed, only a few of the properties of AGM operators are required for this result.

²⁰One need only construct a K -revision model for a suitably rich K to show that MC-revision is nontrivial.

We are currently developing a model of revision and conditionalization for single revisions that adds to the basic revision model probabilistic degrees of belief [7]. If this model can be grafted onto the MC-revision model, “conditional objects” that make statements of conditional probability can be nested in a meaningful way, and given a natural semantics. We are also exploring the application of these ideas to the processes of *J*-conditionalization and *L*-conditionalization. These models were proposed by Goldszmidt and Pearl [18] to capture changes in belief that have degrees of certainty attached, these degrees corresponding to Spohn’s OCFs. Both extensions of MC-revision (either probabilistic or possibilistic) offer ways of attaching quantitative strength measurements to our conditionals and revising these conditionals.

Finally, this model reflects the bias of the AGM model to accepting without question the most recent update. The *primacy of the most recent information* is clearly not a principle that should be accepted in all circumstances. Sometimes things we learn are so radically incompatible with our knowledge that we reject them out of hand, and do not attempt to reconcile them with our current beliefs. Generalizing the AGM and MC-revision models in this way is a difficult task, but one that certainly deserves inquiry.

Acknowledgements

Thanks to Verónica Becher, Moisés Goldszmidt and Mirek Truszczyński for helpful discussions. An anonymous referee provided some very helpful suggestions and pointers to relevant literature. This research was supported by NSERC Research Grant OGP0121843.

References

- [1] Ernest W. Adams. *The Logic of Conditionals*. D.Reidel, Dordrecht, 1975.
- [2] Carlos Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [3] Carlos Alchourrón and David Makinson. On the logic of theory change: Safe contraction. *Studia Logica*, 44(4):405–422, 1985.
- [4] Craig Boutilier. Inaccessible worlds and irrelevance: Preliminary report. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, pages 413–418, Sydney, 1991.
- [5] Craig Boutilier. Epistemic entrenchment in autoepistemic logic. *Fundamenta Informaticae*, 17(1–2):5–30, 1992.
- [6] Craig Boutilier. Sequences of revisions: On the semantics of nested conditionals. Technical Report 92-24, University of British Columbia, Vancouver, 1992.

REFERENCES

36

- [7] Craig Boutilier. The probability of a possibility: Adding uncertainty to default rules. In *Proceedings of the Ninth Conference on Uncertainty in AI*, pages 461–468, Washington, D.C., 1993.
- [8] Craig Boutilier. Conditional logics of normality: A modal approach. *Artificial Intelligence*, 68:87–154, 1994.
- [9] Craig Boutilier. Unifying default reasoning and belief revision in a modal framework. *Artificial Intelligence*, 68:33–85, 1994.
- [10] Craig Boutilier and Moisés Goldszmidt. Revision by conditional beliefs. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 649–654, Washington, D.C., 1993.
- [11] Craig Boutilier and Moisés Goldszmidt. A theory of conditional belief revision. In G. Crocco, L. Fariñas del Cerro, and A. Herzig, editors, *Conditional Logics in Artificial Intelligence*. Oxford University Press, London, 1994. (to appear).
- [12] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. In *Proceedings of the Fifth Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 5–23, Pacific Grove, 1994.
- [13] André Fuhrmann. Theory contraction through base contraction. *Journal of Philosophical Logic*, 20:175–203, 1991.
- [14] Peter Gärdenfors. Belief revisions and the Ramsey test for conditionals. *The Philosophical Review*, 95:81–93, 1986.
- [15] Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, 1988.
- [16] Peter Gärdenfors and David Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the Third Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 83–95, Pacific Grove, 1988.
- [17] Peter Gärdenfors and David Makinson. Nonmonotonic inference based on expectations. *Artificial Intelligence*, 65:197–245, 1994.
- [18] Moisés Goldszmidt and Judea Pearl. Reasoning with qualitative probabilities can be tractable. In *Proceedings of the Eighth Conference on Uncertainty in AI*, pages 112–120, Stanford, 1992.
- [19] Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.
- [20] Sven Ove Hansson. In defense of base contraction. *Synthese*, 91:239–245, 1992.
- [21] Sven Ove Hansson. In defense of the Ramsey test. *Journal of Philosophy*, 89:522–540, 1992.
- [22] Hector J. Levesque. All I know: A study in autoepistemic logic. *Artificial Intelligence*, 42:263–309, 1990.
- [23] Isaac Levi. Iteration of conditionals and the Ramsey test. *Synthese*, 76(1):49–81, 1988.
- [24] David Lewis. *Counterfactuals*. Blackwell, Oxford, 1973.
- [25] David Makinson and Peter Gärdenfors. Relations between the logic of theory change and nonmonotonic logic. In Andre Fuhrmann and Michael Morreau, editors, *The Logic of Theory Change*, pages 185–205. Springer-Verlag, Berlin, 1990.
- [26] Judea Pearl. System Z: A natural ordering of defaults with tractable applications to default reasoning. In M. Vardi, editor, *Proceedings of Theoretical Aspects of Reasoning about Knowledge*, pages 121–135. Morgan Kaufmann, San Mateo, 1990.

- [27] Hans Rott. Conditionals and theory change: Revisions, expansions, and additions. *Synthese*, 81(1):91–113, 1989.
- [28] Hans Rott. A nonmonotonic conditional logic for belief revision. In A. Fuhrmann and M. Morreau, editors, *The Logic of Theory Change*, pages 135–181. Springer-Verlag, Berlin, 1990.
- [29] Hans Rott. Belief change using generalized epistemic entrenchment. *Journal of Logic, Language and Information*, 1(1):45–78, 1992.
- [30] Karl Schlechta. Theory revision and probability. *Notre Dame Journal of Formal Logic*, 307(2):307–318, 1991.
- [31] Wolfgang 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. D. Reidel, Dordrecht, 1987.
- [32] Robert C. Stalnaker. A theory of conditionals. In W.L. Harper, R. Stalnaker, and G. Pearce, editors, *Ifs*, pages 41–55. D. Reidel, Dordrecht, 1968. 1981.
- [33] Robert C. Stalnaker. *Inquiry*. MIT Press, Cambridge, 1984.

A Proofs of Main Results

The truth of many of the propositions in the paper is rather obvious and their proofs are excluded. Certain results are described in the body of the paper and have their proofs sketched there. These proofs are also excluded.

Theorem 7 *Let M be a revision model, let $*$ be the MC-revision operator and let $A, B, C \in \mathbf{L}_{CPL}$.*

- a) *If $M_A^* \models \mathbf{B}\neg B$ then $M_A^* \models B > C$ iff $M \models B > C$.*
- b) *If $M_A^* \not\models \mathbf{B}\neg B$ then $M_A^* \models B > C$ iff $M \models A \wedge B > C$.*

Proof The proof of this theorem is sketched, for the most part, in the text preceding its statement in the body of the paper. ■

Theorem 17 *Let A_1, \dots, A_{n+1} be an incompatible sequence such that A_1, \dots, A_n is compatible. Let k be the maximal element of*

$$\{i \leq n : \neg A_{n+1} \notin ((K_{A_1}^*)_{A_2}^* \cdots)_{A_i}^*\}$$

Then $((K_{A_1}^)_{A_2}^* \cdots)_{A_{n+1}}^* = K_{A_1 \wedge \dots \wedge A_k \wedge A_{n+1}}^*$.*

Proof Let M be the revision model for K and $*$ the MC-revision function. Let A_k be the last element of A_1, \dots, A_{n+1} such that $\neg A_{n+1} \notin ((K_{A_1}^*)_{A_2}^* \cdots)_{A_k}^*$.

a) If there is no such element then $K_{A_1}^* \vdash \neg A_{n+1}$. The minimal cluster of $M_{A_1}^*$ is formed by the set $\min(M, A_1)$. Since the relative ordering of all other worlds is unaffected by this revision, $\min(M_{A_1}^*, A_{n+1}) = \min(M, A_{n+1})$. Now, each subsequent revision is compatible, so $\min(((M_{A_1}^*)_{A_2}^* \cdots)_{A_{i-1}}^*, A_i)$ is contained in the minimal cluster of $((M_{A_1}^*)_{A_2}^* \cdots)_{A_{i-1}}^*$ for each $i \leq n$. An obvious inductive argument shows that (since no A_{n+1} -worlds are contained in the first minimal cluster)

$$\min(((M_{A_1}^*)_{A_2}^* \cdots)_{A_n}^*, A_{n+1}) = \min(M, A_{n+1})$$

Hence $((K_{A_1}^*)_{A_2}^* \cdots)_{A_{n+1}}^* = K_{A_{n+1}}^*$.

b) If such a k exists the, by Proposition 13 and Theorem 14,

$$((K_{A_1}^*)_{A_2}^* \cdots)_{A_k}^* = K_{A_1 \wedge \cdots \wedge A_k}^*$$

Furthermore, since $((K_{A_1}^*)_{A_2}^* \cdots)_{A_k}^* \not\vdash \neg A_{n+1}$, the set $\min(((M_{A_1}^*)_{A_2}^* \cdots)_{A_k}^*, A_{n+1})$ lies within the minimal cluster of $((M_{A_1}^*)_{A_2}^* \cdots)_{A_k}^*$. Since sequence A_1, \dots, A_n is compatible, $\min(((M_{A_1}^*)_{A_2}^* \cdots)_{A_k}^*, A_{k+1})$ must also lie within the minimal cluster of $((M_{A_1}^*)_{A_2}^* \cdots)_{A_k}^*$. However, by the maximality of k , this set must be disjoint from the set $\min(((M_{A_1}^*)_{A_2}^* \cdots)_{A_k}^*, A_{n+1})$. Thus,

$$\min(((M_{A_1}^*)_{A_2}^* \cdots)_{A_{k+1}}^*, A_{n+1}) = \min(((M_{A_1}^*)_{A_2}^* \cdots)_{A_k}^*, A_{n+1})$$

Since all subsequent revisions up to A_{n+1} are compatible, as in case a) above,

$$\min(((M_{A_1}^*)_{A_2}^* \cdots)_{A_n}^*, A_{n+1}) = \min(((M_{A_1}^*)_{A_2}^* \cdots)_{A_k}^*, A_{n+1})$$

But A_{n+1} is compatible with A_1, \dots, A_k , so

$$((K_{A_1}^*)_{A_2}^* \cdots)_{A_{n+1}}^* = K_{A_1 \wedge \cdots \wedge A_k \wedge A_{n+1}}^*$$

■

Lemma 19 *Let A_1, \dots, A_n be a revision sequence such that each proper subsequence A_1, \dots, A_i is characterized by some sentence $s(A_i)$. Then A_1, \dots, A_n is characterized by $s(A_k) \wedge A_n$, where A_k is the most recent compatible revision for A_n . In other words, $((K_{A_1}^*)_{A_2}^* \cdots)_{A_n}^* = K_{s(A_k) \wedge A_n}^*$.*

Proof Let A_k be the most recent compatible revision for A_n . Since A_n is compatible with the subsequence A_1, \dots, A_k , we have that $\min(((M_{A_1}^*)_{A_2}^* \cdots)_{A_k}^*, A_n)$ lies within the minimal cluster of $((M_{A_1}^*)_{A_2}^* \cdots)_{A_k}^*$. However, each subsequent revision A_{k+1}, \dots, A_{n-1} is incompatible with A_n so there can be no A_n -worlds in the set $\min(((M_{A_1}^*)_{A_2}^* \cdots)_{A_{i-1}}^*, A_i)$ for $k < i < n$. Since only the relative status of these worlds is changed by these subsequent revisions, an obvious inductive argument (on the number $n \perp k$ of subsequent revisions) shows that $\min(((M_{A_1}^*)_{A_2}^* \cdots)_{A_{n-1}}^*, A_n) = \min(((M_{A_1}^*)_{A_2}^* \cdots)_{A_k}^*, A_n)$. Since A_1, \dots, A_k is characterized by $s(A_k)$, the minimal cluster of $((M_{A_1}^*)_{A_2}^* \cdots)_{A_k}^*$ is $\min(M, s(A_k))$. Clearly then the minimal cluster of $((M_{A_1}^*)_{A_2}^* \cdots)_{A_n}^*$ is simply $\min(M, s(A_k \wedge A_n))$. Therefore A_1, \dots, A_n is characterized by $s(A_k) \wedge A_n$. ■

Theorem 20 *For any revision sequence A_1, \dots, A_n , there is some subset of these updates $S \subseteq \{A_1, \dots, A_n\}$ such that $((K_{A_1}^*)_{A_2}^* \cdots)_{A_n}^* = K_A^*$ and $A = \wedge S$.*

Proof This result can be shown using a simple inductive argument on n , the number of updates in the sequence. If $n = 1$ then the theorem obviously is true, for A_1 characterizes itself. Now suppose each subsequence A_1, \dots, A_i , $i < n$, is characterized by some sentence $s(A_i) = \wedge S$, where $S \subseteq \{A_1, \dots, A_i\}$. By Lemma 19, the sequence A_1, \dots, A_n is characterized by the sentence $s(A_k) \wedge A_n$ for some $k < n$, where A_k is the most recent compatible revision for A_n (or it is characterized by A_n if no such k exists). By the inductive hypothesis, $s(A_n) = \wedge S$ where $S \subseteq \{A_1, \dots, A_k\} \cup \{A_n\} \subseteq \{A_1, \dots, A_n\}$. (If $k = 0$ we simply observe that $S = \{A_n\} \subseteq \{A_1, \dots, A_n\}$.) ■

Theorem 22 *Let A_1, \dots, A_n be a revision sequence with c incompatible updates represented by σ . For each $1 \leq k \leq n$, if A_k is a compatible revision, then $((K_{A_1}^*)_{A_2}^* \cdots)_{A_k}^* = K_{s(A_k)}^*$, where: a) $s(A_k) = s(A_{\sigma(i)}) \wedge A_{\sigma(i)+1} \wedge \cdots \wedge A_k$; b) $\sigma(i) < k < \sigma(i+1)$, or $\sigma(i) < k$ if $i = c$; and c) $s(A_{\sigma(i)})$ characterizes subsequence $A_1, \dots, A_{\sigma(i)}$.*

Proof Let A_k be a compatible revision in the sequence A_1, \dots, A_n . If no incompatible revision follows A_k , then $k > \sigma(i)$, where $i = c$. Otherwise $\sigma(i) < k < \sigma(i+1)$ for some $i < c$. In either case $A_{\sigma(i)}$ is the maximal incompatible revision in the subsequence A_1, \dots, A_k . Let $s(A_{\sigma(i)})$ characterize $A_1, \dots, A_{\sigma(i)}$. Since each revision $A_{\sigma(i)+1}, \dots, A_k$ is compatible, $A_{\sigma(i)}$ is the most recent compatible revision for A_k , and by $k \perp \sigma(i)$ applications of Lemma 19, we have that A_1, \dots, A_k is characterized by $s(A_k) = s(A_{\sigma(i)}) \wedge A_{\sigma(i)+1} \wedge \cdots \wedge A_k$. ■

Theorem 24 Let A_1, \dots, A_n be a revision sequence with c incompatible updates represented by σ . For each $1 \leq k \leq n$, if A_k is an incompatible revision, then $((K_{A_1}^*)_{A_2}^* \dots)_{A_k}^* = K_{s(A_k)}^*$, where: a) $s(A_k) = s(A_{\sigma(i)}) \wedge A_{\sigma(i)+1} \wedge \dots \wedge A_j \wedge A_k$; b) $A_{\sigma(i)}$ is the maximal incompatible revision for A_k ; c) A_j is the most recent compatible revision for A_k ; and d) $s(A_{\sigma(i)})$ characterizes subsequence $A_1, \dots, A_{\sigma(i)}$.

Proof Let A_j be the most recent compatible revision for A_k . By Theorem 22 and Proposition 23, A_1, \dots, A_j is characterized by $s(A_j) = s(A_{\sigma(i)}) \wedge A_{\sigma(i)+1} \wedge \dots \wedge A_j$. By Lemma 19, A_1, \dots, A_k is characterized by $s(A_j) \wedge A_k$. ■

Theorem 30 Let $S_1, \dots, S_n \in \mathbf{L}_{CPL}$ be such that $\vdash S_i \supset \neg S_j$ if $i \neq j$, and $\vdash S_1 \vee \dots \vee S_n$. For all $A, B \in \mathbf{L}_{CPL}$ either

$$\begin{aligned} FSM(S_1, \dots, S_n) \vdash_{CO^*} A > B; & \quad \text{or} \\ FSM(S_1, \dots, S_n) \vdash_{CO^*} \neg(A > B) \end{aligned}$$

Proof This follows immediately since the theory $FSM(S_1, \dots, S_n)$ is “categorical” in the sense that there is only one CO^* -structure satisfying it (modulo “duplicate worlds”, which can have no influence on the truth of any sentence in the model). ■

Theorem 31 Let revision model M be characterized by $FSM(S_1, \dots, S_n)$ and let S_k be the minimal sentence in this set consistent with A ; that is, $S_k \not\vdash \neg A$ and $S_i \vdash \neg A$ if $i < k$. Then M_A^* is characterized by

$$\begin{aligned} FSM(S_k \wedge A, S_1, \dots, S_{k-1}, S_k \wedge \neg A, S_{k+1}, \dots, S_n) & \quad \text{if } S_k \not\vdash A \\ FSM(S_k, S_1, \dots, S_{k-1}, S_{k+1}, \dots, S_n) & \quad \text{if } S_k \vdash A \end{aligned}$$

Proof Clearly the minimal cluster in M containing A -worlds is that cluster specified by the sentence S_k . Thus the set $\min(M, A)$ consists exactly of those worlds in M satisfying $S_k \wedge A$. In M_A^* , this set forms the minimal cluster and all other clusters remain in the same relative order. However, the cluster that was specified by S_k is now reduced to those worlds satisfying $S_k \wedge \neg A$. If $S_k \vdash A$, this is still true; but the sentence $S_k \wedge A$ is equivalent to S_k , and the cluster in M_A^* satisfying $S_k \wedge \neg A$ is empty. ■