Horn Clause Belief Change: Contraction Functions

James P. Delgrande
School of Computing Science,
Simon Fraser University,
Burnaby, B.C.,
Canada V5A 1S6.
jim@cs.sfu.ca

Abstract

The standard (AGM) approach to belief change assumes that the underlying logic is at least as strong as classical propositional logic. This paper investigates an account of belief change, specifically contraction, where the underlying logic is that governing Horn clauses. Thus this work sheds light on the theoretical underpinnings of belief change by weakening a fundamental assumption of the area. This topic is also of independent interest since Horn clauses have been used in areas such as deductive databases and logic programming. It proves to be the case that there are two distinct classes of contraction functions for Horn clauses: *e-contraction*, which applies to entailed formulas, and *i-contraction*, which applies to formulas leading to inconsistency. E-contraction is applicable in yet weaker systems where there may be no notion of negation (such as in definite clauses). I-contraction on the other hand has severe limitations, which makes it of limited use as a belief change operator. In both cases we explore the class of *maxichoice* functions which, we argue, is the appropriate approach for contraction in Horn clauses theories.

Introduction

The area of *belief change* in knowledge representation is concerned with how an agent may alter its corpus of beliefs in the presence of new information. For example, an agent may *contract*, or reduce its stock of beliefs, *revise* its beliefs in the face of new information, or may *merge* several pieces of information. Contraction and revision are in some sense the more basic operators (since they involve a single agent and single piece of information), and contraction is usually taken as the more fundamental operator, in terms of which revision may be defined. The best-known approach in this area is the so-called AGM approach (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988), named after the original developers.

This paper addresses belief change in the expressively weaker language of *Horn clauses*, where a Horn clause can be written as a rule in the form $a_1 \wedge a_2 \wedge \cdots \wedge a_n \rightarrow a$ for $n \geq 0$, and where a, a_i $(1 \leq i \leq n)$ are atoms. (Thus, written as a disjunction, a Horn clause can have at most one positive literal.) That is, we investigate belief change in which the agent's beliefs are represented in a Horn clause knowledge base and the input is a conjunction of Horn clauses, This topic is of interest for several reasons. First, it sheds light on the theoretical underpinnings of belief change, in that it weakens the assumption that the underlying logic is at least propositional logic. In particular, the notion of *consistency* plays an important role in contraction and revision. Given the significantly weaker notion of negation in Horn clauses, belief change with respect to Horn clauses needs to contend with a correspondingly weaker notion of consistency. Second, Horn clauses have found extensive use in AI and database theory, in areas such as logic programming, truth maintenance systems, and deductive databases.[1] Thus, belief change in such areas is a not unimportant question. Last, as discussed next, belief change in Horn clauses is interesting in its own right.

It proves to be the case that in a Horn clause knowledge base there are two distinct types of contraction, depending on how relative consistency between formulas is defined. In the first case, called *e-contraction*, a Horn clause knowledge base implies a formula, and we wish to weaken the knowledge base so that the formula is not implied. In the second case, *i-contraction*, adding a formula to a knowledge base would result in inconsistency, and we wish to weaken the knowledge base so that inconsistency would not result if the formula were added. Two approaches for *maxichoice contraction* are developed, corresponding to these two notions of consistency. In the AGM approach, maxichoice contraction has undesirable properties and in general is far too strong. In contrast, maxichoice e-contraction appears to be *the* appropriate approach for defining contraction: it has good properties, and the problems of maxichoice AGM contraction do not arise here. For (maxichoice) i-contraction, it is a different story, and problems analogous to maxichoice AGM contraction arise; moreover this problem appears to be fatal, in that moving to a notion of partial meet contraction does not solve the problem.

The next section introduces belief change and Horn clause reasoning. This is followed by material that is pertinent to Horn clause belief contraction. The following section gives the formal details: we provide a set of postulates for e-contraction, and a representation result in terms of *remainder sets*. Notably, the controversial *recovery* postulate does not hold. For i-contraction we similarly develop

[1] It should be pointed out that, while Horn clauses underlie logic programming, the approach presented here does not involve logic programming per se; in particular, negation as failure plays no role in the approach nor in the underlying reasoning system.

an account in terms of remainder sets, and discuss where problems arise. The paper concludes with a discussion on prospects for other representation results and for belief revision in Horn clause theories.

Background: Belief Change

A common approach in addressing belief change is to provide a set of *rationality postulates* for belief change functions. The *AGM approach* (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988) provides the best-known set of such postulates. The goal is to describe belief change at the *knowledge level*, that is on an abstract level, independent of how beliefs are represented and manipulated. Belief states are modelled by sets of sentences, called *belief sets*, closed under the logical consequence operator of a logic that includes classical propositional logic in a language $\mathcal{L}$. Thus a belief set K satisfies the constraint:

If K logically entails ϕ then $\phi \in K$.

For belief set K and formula ϕ, $K + \phi$ is the deductive closure of $K \cup \{\phi\}$, called the *expansion* of K by ϕ. $K_\bot$ is the inconsistent belief set (i.e. $K_\bot = \mathcal{L}$). Expansion captures the simplest form of belief change; it can be reasonably applied when new information is consistent with a belief set

Contraction represents the situation in which the reasoner loses information. Informally, the contraction of a belief set by a formula is another belief set in which that formula is not believed. Formally, a contraction function $\dot{-}$ is a function from $2^{\mathcal{L}} \times \mathcal{L}$ to $2^{\mathcal{L}}$ satisfying the following postulates.

($K\dot{-}1$) $K\dot{-}\phi$ is a belief set.

($K\dot{-}2$) $K\dot{-}\phi \subseteq K$.

($K\dot{-}3$) If $\phi \notin K$, then $K\dot{-}\phi = K$.

($K\dot{-}4$) If not $\vdash \phi$, then $\phi \notin K\dot{-}\phi$.

($K\dot{-}5$) If $\phi \in K$, then $K \subseteq (K\dot{-}\phi) + \phi$.

($K\dot{-}6$) If $\vdash \phi \equiv \psi$, then $K\dot{-}\phi = K\dot{-}\psi$.

($K\dot{-}7$) $K\dot{-}\phi \cap K\dot{-}\psi \subseteq K\dot{-}(\phi \wedge \psi)$.

($K\dot{-}8$) If $\psi \notin K\dot{-}(\psi \wedge \phi)$ then $K\dot{-}(\phi \wedge \psi) \subseteq K\dot{-}\psi$.

Thus, contraction is meaningful for believed sentences ($K\dot{-}3$) and yields a belief set ($K\dot{-}1$) in which the sentence for contraction ϕ is not believed (unless ϕ is a tautology) ($K\dot{-}4$). No new sentences are believed ($K\dot{-}2$). The fifth postulate, the so-called *recovery* postulate, states that nothing is lost if one contracts and expands by the same sentence. This postulate is controversial; see (Hansson 1999) for a discussion and plausible counterexamples. The sixth postulate asserts that contraction is independent of how a sentence is expressed. The last two postulates express relations between contracting by conjunctions and contracting by the constituent conjuncts. The first six postulates are referred to as the *basic postulates* while the last two are referred to as the *extended postulates*.

Revision represents the situation in which new information may be inconsistent with the reasoner's beliefs K, and needs to be incorporated in a consistent manner where possible. A revision function $*$ is a function from $2^{\mathcal{L}} \times \mathcal{L}$ to $2^{\mathcal{L}}$ satisfying a set of postulates analogous to those for contraction; given space limitations we omit the postulate set. Contraction is usually taken as being the more fundamental operator for belief change. Revision can be defined in terms of contraction by means of the *Levi Identity*:

$$K * \phi = (K\dot{-}\neg\phi) + \phi. \tag{1}$$

Thus, to revise by ϕ, make K consistent with ϕ then expand by ϕ. Contraction can be similarly defined in terms of revision by the *Harper identity*; we omit the details.

Various constructions have been proposed to characterise belief change. An *epistemic entrenchment ordering* is a total preorder on sentences of the language, reflecting how "entrenched" or strongly held each sentence is. It has been shown (e.g. (Gärdenfors 1988)) that there is a correspondence between entrenchment orderings and contraction functions satisfying the AGM postulates. A second construction is given by *a system of spheres* (Lewis 1971), comprising a total preorder on the set of interpretations of the language such that the models of the belief set are least in the ranking. (Grove 1988) shows a correspondence between a class of systems of spheres and AGM revision functions. A third construction (though chronologically the earliest) is given by *remainder sets*. Roughly, given a belief set K and formula ϕ, a remainder set of K with respect to ϕ is a maximal subset of K that fails to imply ϕ. A construction of contraction functions can be given in terms of a transitive ordering over remainder sets, where the maximum element(s) of the ordering determine the result of the contraction. Three classes of contraction functions can be defined in this way:

Maxichoice contraction: Contraction is defined to correspond to a single selected remainder set.

Full meet contraction: Contraction is defined to be the intersection of all remainder sets.

Partial meet contraction: Contraction is defined to be the intersection of some remainder sets.

In the AGM approach, maxichoice and full meet contraction have undesirable properties. In maxichoice contraction, revision defined via the Levi identity has the property:

$$\text{If } K \vdash \neg\phi \text{ then } \forall\psi \text{ either } \psi \in K * \phi \text{ or } \neg\psi \in K * \phi \tag{2}$$

Thus in a revision defined in terms of maxichoice contraction, for any sentence ψ, either ψ or its negation is believed. While partial meet contraction avoids the triviality problems attending maxichoice and full meet contraction, it has been suggested that it violates the principle of informational economy that motivates AGM belief change (Levi 1991; Rott 2000).

Belief revision represents the situation wherein an agent receives information concerning some domain, perhaps concerning the current state of the world, or equally well, perhaps some point of time in the past. Belief *update* is a distinct operation that addresses the situation where, roughly, an agent learns of the results of a change in the domain; see (Katsuno and Mendelzon 1992) for a discussion of the difference between revision and updating, along with a development of update.

Last, there have been various specific approaches proposed for defining belief revision. The two best-known approaches are based on a notion of *distance* between the models of the knowledge base and of the sentence for revision. The approach of (Dalal 1988) uses a distance measure based on the number of atoms with differing truth values in two interpretations, while that of (Satoh 1988) is based on set containment. Analogous approaches have been proposed for defining update operators (Winslett 1988; Forbus 1989).

Belief Change and Horn Clause Theories

Work involving belief change and Horn clause theories has often focussed on particular approaches. The complexity of specific approaches to revising or updating knowledge bases is addressed in (Eiter and Gottlob 1992), covering both the case where the knowledge base and formula for revision is an arbitrary formula, and where they are conjunctions of Horn clauses. (As might be expected, results are generally better in the Horn case.) (Liberatore 2000) considers the problem of compact representation for revision in the Horn case. Basically, given a knowledge base K and formula A, both Horn, the central problem addressed is whether the knowledge base, revised according to a given operator, may be represented by a propositional formula whose size is polynomial with respect to the sizes of K and A.

(Langlois et al. 2008) approaches the study of revising Horn formulas by characterising the existence of a complement of a Horn consequence; such a complement corresponds to the result of a contraction operator. This work may be seen as a specific instance of a general framework developed in (Flouris, Plexousakis, and Antoniou 2004). In (Flouris, Plexousakis, and Antoniou 2004), belief change is studied under a broad notion of *logic*, where a logic is a set closed under a Tarskian consequence operator (i.e. that satisfies inclusion, $A \subseteq Cn(A)$; idempotency, $Cn(A) = Cn(Cn(A))$; and monotony, if $A \subseteq B$ then $Cn(A) \subseteq Cn(B)$). In particular, they give a criterion for the existence of a contraction operator satisfying the basic AGM postulates in terms of *decomposability*. (Restall and Slaney 1995) studies belief revision in the case of a logic weaker than classical propositional logic, that of a relevance logic.

In a somewhat different vein, there has been substantial work on *logic program updates*. Such work addresses a notion of belief change in the context of extended logic programs under the answer set semantics (Baral 2003). Roughly, such approaches consider the semantics of sequences of extended logic programs, $(P_1, \ldots, P_n)$, where the problem is to appropriately specify answer sets corresponding to a given sequence. Typically such approaches syntactically manipulate the rules in the constituent sets, and often a notion of preference among rules is employed. A detailed discussion of such approaches is beyond the scope of this paper; however see (Eiter et al. 2002) for an excellent survey.

Horn Clause Theories

Preliminary Considerations

We consider a finitary propositional language $\mathcal{L}$, over a set of atoms, or propositional letters, $\mathbf{P} = \{a, b, c, \ldots\}$ that includes the distinguished atom $\bot$, and truth-functional connectives $\land$ and $\rightarrow$. In propositional logic, a Horn clause is a clause (that is, a disjunction of literals) in which at most one literal is unnegated. We adopt an alternative notation in which a Horn clause is written as a rule r, in the form

$$a_1 \land a_2 \land \cdots \land a_n \rightarrow a \tag{3}$$

where $n \geq 0$, and a, a_i $(1 \leq i \leq n)$ are atoms. For rule r as above, $head(r)$ is a, and $body(r)$ is the set $\{a_1, a_2, \ldots, a_n\}$. If $n = 0$, then r is a *fact* and we write a in place of $\rightarrow a$. If a is $\bot$ then r is an *(integrity) constraint*. A formula of $\mathcal{L}$ is a conjunction of rules, optionally employing parentheses for readability. Thus the following are formulas:

$$a, \qquad p \land (p \land q \rightarrow r) \land (a \land b \rightarrow \bot)$$

Allowing conjunctions of rules adds nothing of interest to the expressibility of the language with respect to reasoning. However, it adds to the expressibility of contraction, as we are able to contract by more than a single Horn clause.

We adopt the following notation. Lower-case Greek characters (ϕ, ψ, ξ, $\ldots$), possibly subscripted, denote arbitrary formulas of $\mathcal{L}$. Upper case Roman characters $(A, B, K, \ldots)$, possibly subscripted, denote sets of formulas.

An *interpretation* of $\mathcal{L}$ is a function $\mathcal{I}$ from $\mathbf{P}$ to $\{T, F\}$ where $\mathcal{I}(\bot) = F$. M is the set of interpretations of $\mathcal{L}$. Sentences of $\mathcal{L}$ are true or false in an interpretation in the expected way: For $\phi \in \mathcal{L}$, ϕ is true in $\mathcal{I}$ just if:

1. ϕ is $a \in \mathbf{P}$ and $\mathcal{I}(a) = T$.
2. ϕ is $\phi_1 \land \phi_2$ and ϕ_1, ϕ_2 are true in $\mathcal{I}$.
3. ϕ is $a_1 \land a_2 \land \cdots \land a_n \rightarrow a$ and if every a_i $(1 \leq i \leq n)$ is true in $\mathcal{I}$ then so is a.

ϕ is false just if it isn't true according to the above. Hence, to be clear, we do not employ negation as failure. An interpretation $\mathcal{I}$ is a *model* of a sentence ϕ (or set of sentences), written $\mathcal{I} \models \phi$, just if $\mathcal{I}$ makes ϕ true. See (Brachman and Levesque 2004) for a discussion on the role of Horn clauses in the wider area of knowledge representation.

Horn Clause Theories and Belief Sets

The following axioms and rules give an inference relation for Horn clauses.[2]

Axioms:

1. $\bot \rightarrow a$
2. $a \rightarrow a$

Rules:

1. From $a_1 \land \cdots \land a_n \rightarrow a$ and $b_1 \land \cdots \land b_n \rightarrow a_i$ infer $a_1 \land \cdots \land a_{i-1} \land b_1 \land \cdots \land b_n \land a_{i+1} \land \cdots \land a_n \rightarrow a$

[2] Since the meaning is clear, and to keep notation to a minimum, the atoms in the axiom schemas and first two rules are used as variable schemas.

2. From $a_1 \wedge \cdots \wedge a_n \to a$ infer $a_1 \wedge \cdots \wedge a_n \wedge b \to a$

3. For rules r_1, r_2, if $body(r_1) = body(r_2)$ and $head(r_1) = head(r_2)$ then from r_1 infer r_2.

4.(a) From $\phi \wedge \psi$ infer ϕ and ψ

(b) From ϕ, ψ infer $\phi \wedge \psi$

Rule 3 allows us to derive, for example, $a \wedge b \to c$ from $b \wedge b \wedge a \to c$. A formula ψ can be derived from a set of formulas A, written $A \vdash_{HC} \psi$, just if ψ can be obtained from A by a finite number of applications of the above rules and axioms; for simplicity we drop the subscript and write $A \vdash \psi$. If $A = \{\phi\}$ is a singleton set then we just write $\phi \vdash \psi$. A set of formulas $A \subseteq \mathcal{L}$ is *inconsistent* just if $A \vdash \bot$. We use $\phi \leftrightarrow \psi$ to represent logical equivalence, that is $\phi \vdash \psi$ and $\psi \vdash \phi$. Though it is not germane to this paper, we have the following soundness and completeness result:

Theorem 1 *Let $A \subseteq \mathcal{L}$ and $\phi \in \mathcal{L}$. Then*
$A \vdash \phi$ iff: $\forall \mathcal{I} \in M$, if $\mathcal{I} \models A$ then $\mathcal{I} \models \phi$.

$Cn(\phi)$ is the set of deductive consequences of a formula or set of formulas ϕ, that is, $Cn(\phi) = \{\psi \mid \phi \vdash \psi\}$. For $K \subseteq \mathcal{L}$, K is a *belief set* just if it satisfies the constraint: $\phi \in K$ iff $K \vdash \phi$. The letters K, K_1, K' etc. will exclusively denote belief sets.

We obtain the following results that are consequences of our working with a finite language.

Theorem 2 *Let X be a belief set. Then there is a formula ϕ such that $Cn(\phi) = X$.*

Theorem 3 *Let X be a belief set and ϕ a formula such that $Cn(\phi) = X$, and let X' be a belief set and ϕ' a formula such that $Cn(\phi') = X'$. Then we have $\phi \vdash \phi'$, $\phi' \vdash \phi$ iff $X = X'$.*

Consistency and Horn Clause Belief Change

In classical logic we have the relation between consistency and derivability:

$$A \vdash \neg\phi \quad \text{iff} \quad A \cup \{\phi\} \vdash \bot.$$

Here, with a Horn clause language, this equivalence fails and we have two distinct relations between a set of formulas A and a formula ϕ:

$$A \vdash \phi \tag{4}$$
$$A \cup \{\phi\} \vdash \bot \tag{5}$$

The first case would have the meaning that A is inconsistent with the negation of ϕ except, of course, the negation of a general formula isn't defined. In the second case, as usual, ϕ is consistent with A if adding ϕ to A does not yield an inconsistency. These two notions of consistency (4)-(5) lead to two classes of belief contraction:

C1: $K \vdash \phi$ and we want a belief set such that $K \dot{-}_e \phi \nvdash \phi$.

C2: $\phi \nvdash \bot$, $K \cup \{\phi\} \vdash \bot$ and we want a belief set such that $K \dot{-}_i \phi \cup \{\phi\} \nvdash \bot$.

We refer to these types of (Horn clause) contraction operators as *entailment-based* contraction (or *e-contraction*) and *inconsistency-based* contraction (or *i-contraction*) respectively. In AGM belief change, e-contraction and i-contraction amount to the same thing; here, we have distinct operators. For e-contraction, one has a belief set that entails (hence the notation $\dot{-}_e$) a formula, and one wants a weaker belief set that doesn't entail the formula. Arguably this is the interpretation that one thinks of for "standard" belief contraction. E-contraction has the very appealing property that it does not rely on notions of consistency or negation. Hence it is appropriate in approaches that don't employ these notions. In fact, our results concerning $\dot{-}_e$ apply to definite clause theories,[3] wherein there is no such thing as an inconsistent set of definite clauses. As well of course, $K \dot{-}_e \phi$ is meaningful in the case where $K \vdash \bot$, and in fact $K \dot{-}_e \bot$ could be employed to remove an inconsistency.

For i-contraction, one has a belief set K that is inconsistent (hence $\dot{-}_i$) with a formula, and wants a weaker belief set such that adding the formula doesn't yield inconsistency. Hence, in the interesting case we do not contract by a formula in a belief set, but rather we contract so that a formula can be consistently added to a belief set. I-contraction would seem to naturally yields a revision operator, using the Levi Identity. Thus:

R: In the interesting case we have $K \cup \phi \vdash \bot$ and we want a consistent belief set such that $K *_i \phi \vdash \phi$.

Unfortunately, it will prove to be the case that i-contraction will prove unsuitable in other ways for defining revision. Note there is no revision operator similarly definable from e-contraction, since the negation of a formula is not defined in our language $\mathcal{L}$.

Horn Clause Contraction

Entailment-Based Contraction

For contraction, given a belief set K and a formula ϕ, the aim is to remove the least set of formulas from K such that a consistency condition involving ϕ holds. This notion is captured by that of a *remainder set*; see (Hansson 1999) for a thorough discussion. In the approach at hand there are two notions of remainder set, corresponding to the two notions of consistency as given in (4) and (5). This subsection presents results concerning Horn clause remainder sets suitable for e-contraction and then uses these results as the basis for specifying Horn clause contraction. The next subsection covers similar ground for i-contraction.

Remainder Sets Our main definition is the same as that in (Hansson 1999), but expressed with respect to Horn clauses.

Definition 1 *Let $K \subseteq \mathcal{L}$ and let $\phi \in \mathcal{L}$.*
$K \downarrow_e \phi$ is the set of sets of formulas s.t. $K' \in K \downarrow_e \phi$ iff

1. $K' \subseteq K$
2. $K' \nvdash \phi$
3. *$\nexists K''$ such that $K' \subset K'' \subseteq K$ and $K'' \nvdash \phi$.*

[3] For our rule-based notation, a definite clause theory is a Horn clause theory over alphabet $\mathbf{P} \setminus \bot$.

Each $K' \in K \downarrow_e \phi$ is a *remainder set* with respect to K and ϕ. Usually such a K' will just be referred to as a remainder set, since the underlying belief set and formula are clear.

We have the following results concerning remainder sets of Horn clause theories.

Theorem 4

1. *If* $X \in K \downarrow_e \phi$ *then* $X = Cn(X)$.
2. *If* $X \subseteq K$ *and* $\phi \notin Cn(X)$ *then* $\exists X'$ *such that* $X \subseteq X' \in K \downarrow_e \phi$.
3. $K \downarrow_e \phi = \emptyset$ *iff* $\vdash \phi$.
4. *If* $X \in K \downarrow_e \phi$ *and* $\psi \in K$, $\psi \notin X$ *then* $X \cup \{\psi\} \vdash \phi$.
5. $K \downarrow_e \phi = K \downarrow_e \psi$ *iff: for every* $X \subseteq K$, $X \vdash \phi$ *iff* $X \vdash \psi$.
6. *If* $\phi, \psi \in K$ *then* $(K \downarrow_e \phi) \cap (K \downarrow_e \psi) \subseteq K \downarrow_e (\phi \wedge \psi)$.
7. *If* $\phi, \psi \in K$ *then* $K \downarrow_e (\phi \wedge \psi) \subseteq (K \downarrow_e \phi) \cup (K \downarrow_e \psi)$.
8. *If* $X \in K \downarrow_e \phi$ *then there is* $X' \supseteq X$ *such that* $X' \in K \downarrow_e (\phi \wedge \psi)$.

The first part says that a remainder set is also a belief set. The second part says that an arbitrary set of formulas that fail to imply a formula can be expanded to a remainder set that fails to imply that formula. Part 3 says that remainder sets exist, except in the trivial case where a tautology is to be dropped, while the next part reflects the maximality of remainder sets. Part 5 is called *choice expansion* in (Hansson 1999). While similar in form to the AGM postulate ($K\dot{-}7$), there is no direct correspondence since ($K\dot{-}7$) specifies a relation between formulas, whereas this result is a relation between remainder sets. In Part 6, replacing $\subseteq$ by $=$ gives what (Hansson 1999) calls *choice distributivity*; here we have a somewhat weaker version. Part 7 gives a partial converse to the containment in Part 6, as does Part 8.

We note two results that hold for remainder sets that are based on classical propositional logic, but do not hold for closed Horn theories.

1. If $X \in K \downarrow_e \phi$ and $\psi \in K$ then $X \cup \{\phi\} \vdash \psi$.

   Counterexample:

   Let $K = Cn(a \rightarrow b)$, $\phi = a \wedge c \rightarrow b$ and $\psi = a \rightarrow b$.
   Then $K \downarrow_e \phi = \{X\}$ where $X = Cn(\{a \wedge c \wedge i \rightarrow b \mid i \in \mathbf{P} \setminus \{a, c\}\})$ We have that $a \rightarrow b \in K$ but $X \cup \{a \wedge c \rightarrow b\} \not\vdash a \rightarrow b$.

2. If $\psi \vdash \phi$ then $K \downarrow_e \phi \subseteq K \downarrow_e \psi$.

   Counterexample:

   Let $K = Cn(a \rightarrow b)$, and let $\psi = a \rightarrow b$ and $\phi = a \wedge c \rightarrow b$.
   Then $K \downarrow_e \psi = \{X_1\}$ where $X_1 = Cn(\{a \wedge i \rightarrow b \mid i \in \mathbf{P} \setminus \{a\}\})$ and $K \downarrow_e \psi = X_2$ where $X_2 = Cn(\{a \wedge c \wedge i \rightarrow b \mid i \in \mathbf{P} \setminus \{a, c\}\})$.
   We get that $K \downarrow_e \phi \not\subseteq K \downarrow_e \psi$, since $\{X_1\} \not\subseteq \{X_2\}$ (though of course $X_1 \subseteq X_2$).
   Note also that

   if $\psi \vdash \phi$ then $\bigcap_{X \in K \downarrow_e \phi} X \subseteq \bigcap_{X \in K \downarrow_e \psi} X$

   is a consequence of the last part of Theorem 4.

Defining Contraction via Remainder Sets We wish to define contraction in terms of remainder sets. Typically there are many remainder sets for any given set of formulas. For example, if $K = Cn(\{p, q, r\})$ then among the remainder sets of $K \downarrow_e (p \wedge q)$ will be a set containing p and r but not q, and another containing $p \rightarrow q$ and $q \rightarrow p$ but neither p nor q. The AGM approach assumes the existence of a *selection function* that chooses remainder set(s) to be used for the contraction. However, selection functions alone lack sufficient structure to capture AGM contraction, and so the standard approach also assumes that the selection function is defined in terms of a transitive preference ordering on remainder sets. The next four definitions are based on definitions in (Gärdenfors 1988; Hansson 1999).

Definition 2 *Let* $K \subseteq \mathcal{L}$. γ *is a* selection function *for* K *if, for every* $\phi \in \mathcal{L}$,

1. *If* $K \downarrow_e \phi \neq \emptyset$ *then* $\emptyset \neq \gamma(K \downarrow_e \phi) \subseteq K \downarrow_e \phi$.
2. *If* $K \downarrow_e \phi = \emptyset$ *then* $\gamma(K \downarrow_e \phi) = \{K\}$.

Definition 3 *Let* γ *be a selection function on* K.

γ *is* relational *if there is a binary relation* $\leq$ *over remainder sets of* K *such that for every* ϕ, *if* $K \downarrow_e \phi \neq \emptyset$ *then*

$$\gamma(K \downarrow_e \phi) = \{X \in K \downarrow_e \phi \mid \textit{for every } X' \in K \downarrow_e \phi \textit{ we have } X' \leq X\}.$$

γ *is* orderly *iff* γ *is relational and the relation* $\leq$ *is a total order (i.e.* $\leq$ *is reflexive, transitive, connected, and antisymmetric).*

Definition 4 *Let* γ *be a selection function on* K. *The* partial meet contraction *on a Horn clause belief set* K generated *by* γ *is given by:*[4] $K \dot{-}_e \phi = \gamma(K \downarrow_e \phi)$.

Definition 5 *Let* $K \subseteq \mathcal{L}$, γ *a selection function, and* $\dot{-}_e$ *the partial meet contraction function generated by* γ.

Then $\dot{-}_e$ *is a* maxichoice e-contraction *on* K *just if* γ *is orderly.*

This defines entailment-based maxichoice Horn clause belief contraction. As mentioned earlier, the corresponding AGM maxichoice operator is far too strong to be useful. However here, in an inferentially weaker system, maxichoice belief contraction is an appropriate (arguably *the* appropriate) contraction operator.

Example 1 *Let* $K = \mathcal{C}n(p \rightarrow q)$ *and* $\phi = p \wedge r \rightarrow q$.

Then $K \dot{-}_e \phi \not\vdash p \rightarrow q$ *(since* $p \rightarrow q \vdash p \wedge r \rightarrow q$*).*

Thus $K \dot{-}_e \phi$ *can be at most* $\mathcal{C}n(\{p \wedge r \wedge i \rightarrow q \mid i \in \mathbf{P} \setminus \{p, r\}\})$.

Example 2 *Let* $K = \mathcal{C}n(\{p \rightarrow q, q \rightarrow s, p \rightarrow r, r \rightarrow s\})$ *and* $\phi = p \rightarrow s$.

Then $\phi \in K$ *but* $\phi \notin K \dot{-}_e \phi$. *Since* $K \dot{-}_e \phi$ *is a belief set,* $K \dot{-}_e \phi$ *must exclude one of* $p \rightarrow q$, $q \rightarrow s$ *and one of* $p \rightarrow r$, $r \rightarrow s$.

Consider the following postulate set.

($K\dot{-}1$) $K \dot{-}_e \phi$ is a belief set. (closure)

[4] In fact $\dot{-}_e$ should be paramaterized by γ. Since we assume a single γ at any point, this notation is unambiguous and simpler.

($K\dot{-}2$) $K\dot{-}_e\,\phi \subseteq K$. (inclusion)

($K\dot{-}3$) If $\phi \notin K$, then $K\dot{-}_e\,\phi = K$. (vacuity)

($K\dot{-}4$) If not $\vdash \phi$, then $\phi \notin K\dot{-}_e\,\phi$. (success)

($K\dot{-}6$) If $\phi \leftrightarrow \psi$, then $K\dot{-}_e\,\phi = K\dot{-}_e\,\psi$. (extensionality)

($K\dot{-}T$) If $\vdash \phi$ then $K\dot{-}_e\,\phi = K$ (tautology)

($K\dot{-}8e$) If $\psi \notin K\dot{-}(\psi \wedge \phi)$ then $K\dot{-}(\phi \wedge \psi) = K\dot{-}\psi$. (conjunctive equality)

Postulates ($K\dot{-}1$) – ($K\dot{-}T$) are referred to as the *basic postulates* for e-contraction, as they are analogous to the AGM basic postulates ($K\dot{-}1$) – ($K\dot{-}6$). Similarly, ($K\dot{-}8e$) is referred to as the *extended postulate* for e-contraction.

($K\dot{-}8e$) is quite powerful, and reflects the fact that in the representation result following, we have a total order over remainder sets. We obtain the following consequences of our postulate set:

Theorem 5

1. $K\dot{-}_e\,\phi = K\dot{-}_e\,\phi \wedge \psi$ *or* $K\dot{-}_e\,\psi = K\dot{-}_e\,\phi \wedge \psi$
2. *If* $\phi \notin K$ *then* $K\dot{-}_e\,\phi \wedge \psi \not\subset K\dot{-}_e\,\phi$.
3. *If* $K \vdash \phi$, $K \vdash \psi$, $K\dot{-}_e\,\phi \not\vdash \psi$ *and* $K\dot{-}_e\,\psi \not\vdash \phi$ *then* $K\dot{-}_e\,\phi = K\dot{-}_e\,\psi$. *(orderliness)*
4. *If* $\phi \in K\dot{-}_e\,(\phi \wedge \psi)$ *then* $\phi \in K\dot{-}_e\,(\phi \wedge \psi \wedge \delta)$ *(conjunctive trisection)*

Since Theorem 5.1 clearly implies ($K\dot{-}7$), a maxichoice Horn clause operator satisfies all of the AGM postulates with the exception of the recovery postulate. Example 1 provides a counterexample to the recovery postulate. In this example, we have for $\phi = p \wedge r \rightarrow q$ that

$$\mathcal{C}n(p \rightarrow q)\ \dot{-}_e\ (p \wedge r \rightarrow q)\ \not\vdash\ p \rightarrow q.$$

We get $(K\dot{-}_e\,\phi) \cup \{\phi\} \not\vdash p \rightarrow q$, violating recovery. As another example showing the violation of recover, consider $K = Cn(a, b, c)$. There is no reason we could not have $K\dot{-}_e\,a = Cn(b)$, and quite clearly $K\dot{-}_e\,a \cup \{a\} \not\vdash c$.

The postulate ($K\dot{-}T$) (*tautology*, called *failure* in (Hansson 1999)) is derivable using the AGM postulates, but relies on the recovery postulate ($K\dot{-}5$) for its proof. Since we lack the recovery postulate, it is required here as a postulate, covering a special case, in its own right.

In the AGM approach, Theorem 5.1 entails the *fullness* postulate:

If $K \vdash \psi$ and $K\dot{-}_e\,\phi \not\vdash \psi$ then $K\dot{-}_e\,\phi \cup \{\psi\} \vdash \phi$.

However the derivation seems to require the recovery postulate, and so the derivation of fullness does not obtain. In the AGM approach fullness leads to the unfortunate result that in a revision function satisfying fullness, if $K \cup \{\phi\} \vdash \bot$ then $K * \phi$ is characterized by a single possible world or interpretation (and so (2) obtains). This triviality result does not obtain here as a result of the limited inferential capabilities of Horn clauses.

We have the following representation result which, for convenience, we divide into two parts:

Theorem 6 *Let* $\dot{-}_e$ *be a maxichoice e-contraction function on Horn clause belief set* K. *Then* $\dot{-}_e$ *satisfies the postulates* ($K\dot{-}1$), ($K\dot{-}2$), ($K\dot{-}3$), ($K\dot{-}4$), ($K\dot{-}6$), ($K\dot{-}T$), ($K\dot{-}8e$).

Theorem 7 *Let* $\dot{-}_e$ *be an operator on Horn clause belief set* K *that satisfies the postulates* ($K\dot{-}1$), ($K\dot{-}2$), ($K\dot{-}3$), ($K\dot{-}4$), ($K\dot{-}6$), ($K\dot{-}T$), ($K\dot{-}8e$). *Then* $\dot{-}_e$ *is a maxichoice e-contraction operator on* K.

Inconsistency-Based Contraction

In this subsection, we develop inconsistency-based contraction, denoted $\dot{-}_i$. It proves be the case that there are some very significant drawbacks to this type of contraction. First, it is by no means clear that $\dot{-}_i$ is an interesting contraction operator, in that it appears to be very weak. Second is the fact that it falls prey to the same triviality result as for standard AGM maxichoice contraction. However, in this case, falling back to a partial meet operator is not an adequate solution. These points are developed next.

To begin, underlying the semantic specification of $\dot{-}_i$ we have again the notion of a *remainder set*, suitably modified:

Definition 6 *Let* $K \subseteq \mathcal{L}$ *and* $\phi \in \mathcal{L}$.
$K \downarrow_i \phi$ *is the set such that* $K' \in K \downarrow_i \phi$ *iff*

1. $K' \subseteq K$
2. $K' \cup \{\phi\} \not\vdash \bot$
3. $\nexists K''$ *such that* $K' \subset K'' \subseteq K$ *and* $K'' \cup \{\phi\} \not\vdash \bot$.

We obtain results similar to those of Theorem 4.

Theorem 8

1. *If* $X \in K \downarrow_i \phi$ *then* $X = Cn(X)$.
2. *If* $X \subseteq K$ *and* $X \cup \{\phi\} \not\vdash \bot$ *then* $\exists X'$ *such that* $X \subseteq X' \in K \downarrow_i \phi$.
3. $K \downarrow_i \phi = \emptyset$ *iff* $\phi \vdash \bot$.
4. *If* $X \in K \downarrow_i \phi$ *and* $\psi \in K$, $\psi \notin X$ *then* $X \cup \{\phi, \psi\} \vdash \bot$.
5. $K \downarrow_i \phi = K \downarrow_i \psi$ *iff for every* $K' \subseteq K$, $K' \cup \{\phi\} \not\vdash \bot$ *iff* $K' \cup \{\psi\} \not\vdash \bot$.
6. *If* $X \in K \downarrow_i \phi$ *and* $X \cup \{\phi, \psi\} \not\vdash \bot$ *then* $X \in K \downarrow_i (\phi \wedge \psi)$.
7. *If* $X \in K \downarrow_i (\phi \wedge \psi)$ *then* $\exists X'$ *such that* $X \subseteq X' \in K \downarrow_i \phi$.

We note two results that hold for remainder sets based on classical propositional logic, but do not hold for closed Horn theories under $\downarrow_i$.

1. If $X \in K \downarrow_i \phi$ and $X \cup \{\psi\} \not\vdash \bot$ then $X \in K \downarrow_i (\phi \wedge \psi)$.
   Counterexample:
   Let $K = Cn(a \rightarrow b)$, $\phi = b \rightarrow \bot$ and $\psi = a$.
   Then $K \downarrow_i \phi = \{K\}$. $K \cup \{\psi\} \not\vdash \bot$ but $K \notin K \downarrow_i (\phi \wedge \psi) = Cn(\{i \wedge a \rightarrow b \mid i \in P \setminus \{a\}\})$.
2. $K \downarrow_i (\phi \wedge \psi) \subseteq K \downarrow_i \phi \cup K \downarrow_i \psi$.
   Counterexample:
   Let $K = Cn(a, b)$, $\psi = a \rightarrow \bot$ and $\phi = b \rightarrow \bot$.
   Then $K \downarrow_i (\phi \wedge \psi) = \{X\} = \{Cn(\{i \rightarrow a, j \rightarrow b \mid i, j \in \mathbf{P}\})\}$
   $K \downarrow_i \phi = \{X_\phi\} = \{Cn(\{i \rightarrow a \mid i \in \mathbf{P}\} \cup \{b\})\}$
   $K \downarrow_i \psi = \{X_\psi\} = \{Cn(\{i \rightarrow b \mid i \in \mathbf{P}\} \cup \{a\})\}$
   Thus while we have $X \subseteq X_\phi \cup X_\psi$ we don't have $\{X\} \subseteq \{X_\phi\} \cup \{X_\psi\}$.

We define i-contraction in terms of remainder sets in a manner analogous to that for e-contraction. Definitions 2, 3, and 4 are repeated for $\dot{-}_i$, but referring to $\downarrow_i$ rather than $\downarrow_e$. Hence, in the same way as for e-contraction, we first define a total order on remainder sets as a first step to defining contraction in terms of such an ordering.

From this we give a definition for inconsistency-based maxichoice Horn clause contraction functions:

Definition 7 *Let $K \subseteq \mathcal{L}$, let γ be a selection function on i-remainder sets, and let $\dot{-}_i$ be the partial meet contraction function generated by γ.*

Then $\dot{-}_i$ is a maxichoice i-contraction *on K just if γ is orderly.*

Before considering what postulates hold (i.e. are sound) with respect to such an ordering, it is worthwhile considering those that do not. First, $(K\dot{-}7)$ does not hold. The following provides a counterexample:

Example 3 *Let $K = \mathcal{C}n(\{p \rightarrow q\})$, $\phi = p$, and $\psi = q \rightarrow \bot$. Then $K \dot{-}_i \phi = K \dot{-}_i \psi = K$, but $p \rightarrow q \notin K \dot{-}_i (\phi \wedge \psi)$.*

For $(K\dot{-}8)$, consider the following relation, which is clearly a consequence of $(K\dot{-}8)$:

$$K\dot{-}(\phi \wedge \psi) \subseteq K\dot{-}\psi \quad \text{or} \quad K\dot{-}(\phi \wedge \psi) \subseteq K\dot{-}\phi. \qquad (6)$$

The next example shows that this relation, and consequently $(K\dot{-}8)$ is not satisfied by a total (or indeed partial) ordering on remainder sets.

Example 4 *Let $K = \mathcal{C}n(\{a, b, c, d\})$, $\phi = a \wedge b \rightarrow \bot$, and $\psi = c \wedge d \rightarrow \bot$.*

Then there is an ordering on remainder sets so that $\{b, c, d\} \subset K\dot{-}\phi$ and $\{a, b, d\} \subset K\dot{-}\psi$. Intuitively, for a remainder set $X = K \dot{-}_i (\phi \wedge \psi)$ one would want either $a \notin X$ or $c \notin X$ in order to satisfy (6).

However there is $Y \in K \downarrow_i (\phi \wedge \psi)$, such that $a, c \subset Y$, and there is nothing preventing Y from being maximal among the elements of $K \downarrow_i (\phi \wedge \psi)$.

What these examples mean is that belief sets for $K\dot{-}\phi$ and $K\dot{-}\psi$ are essentially detached from those for $K\dot{-}(\phi \wedge \psi)$. Given that it is the interdependence among these elements in AGM revision that allows the representation results (specifically, enables the definition of an ordering on remainder sets), it is unlikely that such a general result is possible here.

What can be obtained is considered next. A postulate set for i-contraction is given as follows.

$(K\dot{-}1)$ $K \dot{-}_i \phi$ is a belief set. (closure)

$(K\dot{-}2)$ $K \dot{-}_i \phi \subseteq K$. (inclusion)

$(K\dot{-}3i)$ If $K \cup \{\phi\} \not\vdash \bot$ then $K \dot{-}_i \phi = K$. (vacuity for $\dot{-}_i$)

$(K\dot{-}4i)$ If $\phi \not\vdash \bot$ then $K \dot{-}_i \phi \cup \{\phi\} \not\vdash \bot$ (success for $\dot{-}_i$)

$(K\dot{-}6)$ If $\phi \leftrightarrow \psi$ then $K \dot{-}_i \phi = K \dot{-}_i \psi$. (extensionality)

$(K\dot{-}\bot)$ If $\phi \vdash \bot$ then $K \dot{-}_i \phi = K$ (falsity)

$(K\dot{-}7i)$ If $K \dot{-}_i \phi \cup \{\phi, \psi\} \not\vdash \bot$ then $K \dot{-}_i \phi = K \dot{-}_i \phi \wedge \psi$. (equality)

The basic postulate set is much the same as for e-contraction. The postulates that differ ($(K\dot{-}3i)$, $(K\dot{-}4i)$, and $(K\dot{-}\bot)$, do so on account of the different way e-contraction and i-contraction relate to the fundamental intuitions (C1 vs. C2) for contraction with respect to Horn clauses. $(K\dot{-}1)$ – $(K\dot{-}\bot)$ are the basic postulates for i-contraction. $(K\dot{-}7i)$ provides an "extended postulate" but it is clearly weak, and inadequate to yield a general representation result.

We obtain instead a weaker result. Call a selection function γ *singleton* if for every ϕ, $\gamma(K \downarrow_e \phi) = \{X\}$ for some $X \in K \downarrow_e \phi$, and an i-contraction function *singleton* if it is defined in terms of a singleton selection function. Then:

Theorem 9 *Let $\dot{-}_i$ be a maxichoice i-contraction operator on Horn clause belief set K. Then $\dot{-}_i$ satisfies the postulates $(K\dot{-}1)$, $(K\dot{-}2)$, $(K\dot{-}3i)$, $(K\dot{-}4i)$, $(K\dot{-}6)$, $(K\dot{-}\bot)$, $(K\dot{-}7i)$.*

Theorem 10 *If $\dot{-}_i$ is an operator on Horn clause belief set K that satisfies the postulates $(K\dot{-}1)$, $(K\dot{-}2)$, $(K\dot{-}3i)$, $(K\dot{-}4i)$, $(K\dot{-}6)$, $(K\dot{-}\bot)$, then $\dot{-}_i$ is a singleton i-contraction operator on K.*

However, the most negative result is the following:

Theorem 11 *Let K be a belief set, $a \in \mathbf{P}$, and $a \rightarrow \bot \in K$. Then for every $b \in \mathbf{P}$, we have: $b \in (K \dot{-}_i a) + a$ or $b \rightarrow \bot \in (K \dot{-}_i a) + a$.*

Thus, if revision were defined in terms of i-contraction and expansion via the Levi identity, then in the above noted-case revision would result in a complete and consistent set of literals – i.e. all structure of K, given in terms of Horn clauses would be lost. The standard recourse in AGM contraction – moving to a partial meet contraction – is of little help here, since one would still end up with a set of literals, and again any prior structure in K would be lost.

Discussion

AGM Contraction and Horn Clause Contraction There are several ways in which Horn clause contraction differs from AGM contraction. First, of course, is the absence of the recovery postulate. Arguably, at least for Horn clauses, the recovery postulate is undesirable: For example, for e-contraction, if a belief set contains the Horn clause closure of $p \rightarrow q$ and one contracts by $p \wedge r \rightarrow q$ and then expands by the same, it seems quite unreasonable that $p \rightarrow q$ be obtained in the result.[5]

Second, in Horn clause contraction, maxichoice contraction appears to constitute the appropriate approach, in which the result of a contraction is given semantically by a remainder set. In AGM belief change, a single remainder set leads to change functions with undesirable properties,[6] and so as a solution to this problem, one typically defines belief contraction with respect to the intersection of select remainder sets. This resolves the problem of undesirable properties, but introduces another. As (Levi 1991) (as well as (Rott 2000))

[5]Unless, perhaps, one develops a non-Markovian approach wherein the full sequence of changes bears on the final outcome.

[6]Essentially the problem with AGM maxichoice contraction is that if $\alpha \in K$ then for every β either $\alpha \vee \beta \in K\dot{-}\alpha$ or $\alpha \vee \neg\beta \in K\dot{-}\alpha$. This leads to the undesirable "fullness" result for revision defined in terms of contraction by the Levi Identity.

points out, while a single remainder set meets the criterion of informational economy, the intersection of remainder sets does not. In e-contraction this dilemma does not arise, since remainder sets for Horn clauses are better behaved than their propositional logic counterparts.

The third difference is that with Horn clauses one has two classes of contraction operators. In the first case, a formula is entailed by a belief set and the desired outcome is a belief set in which the formula is not entailed. In the second case, the addition of a formula would lead to inconsistency and the desired outcome is a smaller belief set to which this formula can be safely added. As discussed, the first approach, e-contraction, is applicable for approaches that have no notion of consistency or negation. Hence it is appropriate for belief change with respect to definite clauses, where the only needed belief change operators are contraction and expansion. In contrast, the second approach, i-contraction, appears to be too weak to be useful as a contraction operator; this is a disappointing result, as it would otherwise be the appropriate contraction operator for defining revision via the Levi identity.

Future Work From a technical point of view, it would be interesting to investigate partial meet Horn clause contraction in the case of e-contraction. Although we have argued that for Horn clauses theories all one needs is maxichoice belief change, an investigation of partial meet contraction for Horn clauses would illuminate the role of relative inferential strength with respect to the underlying postulate set.

Another direction of interest would seem to be to investigate representation results with respect to epistemic entrenchment orderings and systems of spheres. There are however obstacles to such potential alternative representation results. In the case of epistemic entrenchment orderings for example, the standard definition of contraction refers to arbitrary disjunctions which, of course, we don't have access to in Horn clauses. And for systems of spheres, there is a problem that one loses the clean correspondence between sets of possible worlds and theories.

A third interesting issue concerns the formulation of revision operators in Horn clause theories. The negative results concerning the most promising approach, viz. i-contraction, indicates that a *base* approach to belief change, where a knowledge base is not necessarily logically closed, may be the appropriate means to address revision.

Last, of course, it is of interest to see how and where the approach may be applied, presumably in applications involving some form of logic program, or perhaps in deductive databases, or involving Horn-form integrity constraints.

Acknowledgements

I thank Tommie Meyer and Maurice Pagnucco for helpful l discussions on an earlier version of the paper. Part of this work was completed while the author was a visitor at NICTA, Sydney, Australia. This work was partially supported by a Canadian NSERC Discovery Grant.

References

Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic* 50(2):510–530.

Baral, C. 2003. *Knowledge Representation, Reasoning and Declarative Problem Solving*. Cambridge University Press.

Brachman, R., and Levesque, H. 2004. *Knowledge Representation and Reasoning*. Morgan Kaufmann Publishers.

Dalal, M. 1988. Investigations into theory of knowledge base revision. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, 449–479.

Eiter, T., and Gottlob, G. 1992. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence* 57(2-3):227–270.

Eiter, T.; Fink, M.; Sabbatini, G.; and Tompits, H. 2002. On properties of update sequences based on causal rejection. *Theory and Practice of Logic Programming* 2(6):711–767.

Flouris, G.; Plexousakis, D.; and Antoniou, G. 2004. Generalizing the AGM postulates: Preliminary results and applications. In *Proceedings of the 10th International Workshop on Non-Monotonic Reasoning (NMR-04)*, 171–179.

Forbus, K. 1989. Introducing actions into qualitative simulation. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1273–1278.

Gärdenfors, P. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books, MIT Press, Cambridge Massachusetts.

Grove, A. 1988. Two modellings for theory change. *Journal of Philosophical Logic* 17:157–170.

Hansson, S. O. 1999. *A Textbook of Belief Dynamics*. Applied Logic Series. Kluwer Academic Publishers.

Katsuno, H., and Mendelzon, A. 1992. On the difference between updating a knowledge base and revising it. In Gärdenfors, P., ed., *Belief Revision*, 183–203. Cambridge University Press.

Langlois, M.; Sloan, R.; Szörényi, B.; and Turán, G. 2008. Horn complements: Towards Horn-to-Horn belief revision. In *Proceedings of the AAAI National Conference on Artificial Intelligence*.

Levi, I. 1991. *The Fixation of Belief and its Undoing*. Cambridge University Press.

Lewis, D. 1971. Completeness and decidability of three logics of counterfactual conditionals. *Theoria* 37:74–85.

Liberatore, P. 2000. Compilability and compact representations of revision of Horn knowledge bases. *ACM Transactions on Computational Logic* 1(1):131–161.

Restall, G., and Slaney, J. 1995. Realistic belief revision. In *Proceedings of the First World Congress in the Fundamentals of Artificial Intelligence*, 367–378.

Rott, H. 2000. Two dogmas of belief revision. *Journal of Philosophy* 97:503–522.

Satoh, K. 1988. Nonmonotonic reasoning by minimal belief revision. In *Proceedings of the International Conference on Fifth Generation Computer Systems*, 455–462.

Winslett, M. 1988. Reasoning about action using a possible models approach. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, 89–93.

Proofs of main theorems

Proof 4:

1. If $X = K$ then $X = Cn(X)$ by definition; hence assume that $X \neq K$. Assume that $X \in K \downarrow_e \phi$ but $X \neq Cn(X)$. Let $\delta \in Cn(X) \setminus X$. Then $X \subset X \cup \{\delta\} \subseteq K$ and $X \cup \{\delta\} \not\vdash \phi$, contradicting the assertion than X is a remainder set. Hence $X \neq Cn(X)$ cannot be the case and so $X = Cn(X)$.

2. Let $X \subseteq K$ and $\phi \notin Cn(X)$. If $X = K$ then the result follows trivially. So assume that $X \subset K$. We have that $X \not\vdash \phi$. Given that we have a finite language, it is easy to show that a maximal $X' \supseteq X$ such that $X' \not\vdash \phi$ satisfies $X' = Cn(X')$, and so $X' \in K \downarrow_e \phi$.

3. If $K \downarrow_e \phi = \emptyset$ then there is no $K' \subseteq K$ such that $K' \not\vdash \phi$. Specifically we have that it is not the case that for $K' = \emptyset$ that $K' \not\vdash \phi$, i.e. $\vdash \phi$.

   Similarly for the converse, if $\vdash \phi$ then there is no set of formulas K' than fail to imply ϕ.

4. Assume that $X \in K \downarrow_e \phi$ and $\psi \in K$, $\psi \notin X$. If $X \cup \{\psi\} \not\vdash \phi$, then $Cn(X \cup \{\psi\})$ contradicts the maximality condition for X being a remainder set.

5. Left-to-right: Assume that $K \downarrow_e \phi = K \downarrow_e \psi$ but for some $X \subseteq K$, we have (without loss of generality) $X \vdash \phi$ but $X \not\vdash \psi$. So by Part 2 of the present theorem there is X' where $X \subseteq X'$ and $X' \in K \downarrow_e \psi$. By assumption we also have $X' \in K \downarrow_e \phi$. But $X \vdash \phi$, $X \subseteq X'$, hence $X' \vdash \phi$, contradiction.

   Right-to-left: Assume that for every $X \subseteq K$, $X \vdash \phi$ iff $X \vdash \psi$, but that $K \downarrow_e \phi \neq K \downarrow_e \psi$. Without loss of generality assume that there is $X \in K \downarrow_e \phi$ and $X \notin K \downarrow_e \psi$. Thus $X \not\vdash \phi$ and so by assumption (since $X \subseteq K$) we have $X \not\vdash \psi$. Since $X \notin K \downarrow_e \psi$, by Part 2 of the present theorem there is X' where $X' \supset X$ and $X' \not\vdash \psi$. Hence $X' \in K \downarrow_e \psi$. Since $X' \not\vdash \psi$, by our initial assumption we have $X' \not\vdash \phi$, contradicting the fact that $X' \supset X \in K \downarrow_e \phi$.

6. We show that for $\phi, \psi \in K$ that $X \in K \downarrow_e \phi$ and $X \in K \downarrow_e \psi$ implies that $X \in K \downarrow_e (\phi \wedge \psi)$.

   $X \in K \downarrow_e \phi$ implies that $X \not\vdash \phi$ and for any X', $X \subset X' \subseteq K$ implies that $X' \vdash \phi$.

   $X \in K \downarrow_e \psi$ implies that $X \not\vdash \psi$ and for any X', $X \subset X' \subseteq K$ implies that $X' \vdash \psi$.

   Therefore for $X \in K \downarrow_e \phi$ and $X \in K \downarrow_e \psi$ we get that $X \not\vdash \phi \wedge \psi$ but for every $X' \supset X$ where $X' \subseteq K$ we have $X' \vdash \phi$, $X' \vdash \psi$ whence $X' \vdash \phi \wedge \psi$. Thus $X \in K \downarrow_e (\phi \wedge \psi)$.

7. Let $\phi, \psi \in K$ and $X \in K \downarrow_e (\phi \wedge \psi)$. We show that $X \in K \downarrow_e \phi$ or $X \in K \downarrow_e \psi$.

   Since $X \not\vdash \phi \wedge \psi$ by assumption, we must have $X \not\vdash \phi$ or $X \not\vdash \phi$. Assume without loss of generality that $X \not\vdash \phi$.

   For every $X' \supset X$ and $X' \subseteq K$ we have that $X' \vdash \phi \wedge \psi$ (since $X \in K \downarrow_e (\phi \wedge \psi)$). Thus for every such X' we have $X' \vdash \phi$ and so by definition $X \in K \downarrow_e \phi$.

8. Let $X \in K \downarrow_e \phi$. If $X \vdash \psi$ then since X is a maximal subset of K that fails to imply ϕ then X is also a maximal subset of K that fails to imply $\phi \wedge \psi$. Hence $X \in K \downarrow_e (\phi \wedge \psi)$.

   If $X \not\vdash \psi$ then $X \not\vdash \phi \wedge \psi$. By Part 2 of the theorem there is X' such that $X \subseteq X' \in K \downarrow_e (\phi \wedge \psi)$. ■

Proof 5:

1. From $(K \dot{-} 4)$ we have that $K \dot{-}_e (\phi \wedge \psi) \not\vdash \phi \wedge \psi$

   Hence $K \dot{-}_e (\phi \wedge \psi) \not\vdash \phi$ or $K \dot{-}_e (\phi \wedge \psi) \not\vdash \psi$.

   Via $(K \dot{-} 8e)$ we get $K \dot{-}_e \psi = K \dot{-}_e \phi \wedge \phi$ or $K \dot{-}_e \psi = K \dot{-}_e \phi \wedge \psi$.

2. Assume that $\phi \notin K$. We have that $\phi \notin K \dot{-}_e \phi \wedge \psi$ or $\phi \in K \dot{-}_e \phi \wedge \psi$. In the former case our result follows via $(K \dot{-} 8e)$.

   In the latter case, if $K \dot{-}_e \phi \wedge \psi \subset K \dot{-}_e \phi$ then $\phi \in K \dot{-}_e \phi \wedge \psi$ would imply that $\phi \in K \dot{-}_e \phi$, contradicting $(K \dot{-} 4)$.

3. We show that orderliness is implied by $K \dot{-} 7e$, $K \dot{-} 8e$, and $K \dot{-} F$.

   Assume that $K \vdash \phi$, $K \vdash \psi$, $K \dot{-}_e \phi \not\vdash \psi$ and $K \dot{-}_e \psi \not\vdash \phi$. From Theorem 5.1 we have that either $K \dot{-}_e \phi = K \dot{-}_e (\phi \wedge \psi)$ or $K \dot{-}_e \psi = K \dot{-}_e (\phi \wedge \psi)$. Assume without loss of generality that $K \dot{-}_e \phi = K \dot{-}_e (\phi \wedge \psi)$. Since $K \dot{-}_e \phi \not\vdash \psi$ and $K \dot{-}_e \phi = K \dot{-}_e (\phi \wedge \psi)$, we have that $K \dot{-}_e (\phi \wedge \psi) \not\vdash \psi$.

   $K \dot{-} 8e$ then yields that $K \dot{-}_e (\phi \wedge \psi) = K \dot{-}_e \psi$. Since we already have that $K \dot{-}_e \phi = K \dot{-}_e (\phi \wedge \psi)$, we obtain that $K \dot{-}_e \phi = K \dot{-}_e \psi$.

4. Omitted ■

Proof 6:

Let $\dot{-}_e$ be a maxichoice e-contraction operator defined on Horn clause belief set K, and let $\mathcal{K}$ be the set of Horn clause belief sets contained in K.

Let γ be the selection function associated with $\dot{-}_e$, and let $\leq$ be the total order associated with γ.

$K \dot{-}_e \phi$ is defined by: If $K \vdash \phi$ and $\not\vdash \phi$, then $K \dot{-}_e \phi = K'$ where $K' \not\vdash \phi$ and for every K'' where $K'' \not\vdash \phi$ we have $K'' \leq K'$. If $K \not\vdash \phi$ or $\vdash \phi$ then $K \dot{-}_e \phi = K$.

$(K \dot{-} 1)$ is satisfied since $\dot{-}_e$ is well defined, and $\leq$ is defined on belief sets. $(K \dot{-} 2)$ and $(K \dot{-} 3)$ follow directly from the definition of $\dot{-}_e$, as does $(K \dot{-} 4)$. $(K \dot{-} 6)$ is a consequence of the fact that we're working with belief sets: for belief set K and $\phi \leftrightarrow \psi$ we have $\phi \in K$ iff $\psi \in K$.

$(K \dot{-} T)$ is a direct consequence of the definition of $\dot{-}_e$.

For $(K \dot{-} 8e)$, assume that $\psi \notin K \dot{-} (\psi \wedge \phi)$. Assume towards a contradiction that $K \dot{-} (\phi \wedge \psi) \neq K \dot{-} \psi$. Let

$K' = K \dot{-} (\phi \wedge \psi)$ and $K'' = K \dot{-} \psi$. We must have $K' < K''$ or $K'' < K'$ since $\leq$ is a total order. The former case is not possible, since $K'' \not\vdash \phi \wedge \psi$, contradicting the maximality of K' in $\leq$ with respect to remainder sets that fail to imply $\phi \wedge \psi$. Similarly the latter case, $K'' < K'$, is not possible: we have $K' \not\vdash \psi$ since $K' = K \dot{-} (\phi \wedge \psi)$ and we assumed at the outset that $\psi \notin K \dot{-} (\psi \wedge \phi)$. But $K' \not\vdash \psi$ contradicts the maximality of K'' in $\leq$ with respect to remainder sets that fail to imply ψ. ■

Proof 7: Let $\dot{-}_e$ be an operator on Horn clause belief set K that satisfies the postulates for e-contraction.

Define $\gamma(K \downarrow_e \phi) = \{K \dot{-}_e \phi\}$.

We need to show first that

1. γ is well-defined;
2. γ is a maxichoice selection function; and
3. for every formula ϕ, $\cap\gamma(K \downarrow_e \phi) = K \dot{-}_e \phi$.

These parts follow exactly as in (Hansson 1999, p. 128) and are not repeated here.

For the extended postulate ($K \dot{-} 8e$), for remainder sets $X, Y \subseteq K$, define $X \leq Y$ iff

1. $Y = K$, or
2. (a) there is $\phi \in K$ such that $Y = K \dot{-}_e \phi$, and
   (b) for every ϕ where $\phi \notin X$, $\phi \notin Y$, we have $X = K \dot{-}_e \phi$ only if $X = Y$.

We show that $\leq$ is a total order over the set

$$\{X \mid \exists\phi \text{ such that } X = K \dot{-}_e \phi\}.$$

1. Reflexivity: This is trivially satisfied.
2. Antisymmetry: Assume that $X \leq Y$ and $Y \leq X$; we show that $X = Y$.

   Since $X \leq Y$, according to condition 2(a) there is a formula, ϕ, such that $Y = K \dot{-}_e \phi$. As well, condition 2(b) stipulates that for every δ such that $\delta \notin X$ and $\delta \notin Y$, $X = K \dot{-}_e \delta$ implies that $X = Y$.

   Since $Y \leq X$, according to condition 2(a) there is a formula, ψ, such that $X = K \dot{-}_e \psi$. As well, condition 2(b) stipulates that for every δ such that $\delta \notin Y$ and $\delta \notin X$, $Y = K \dot{-}_e \delta$ implies that $X = Y$.

   From Theorem 5.1 we have $K \dot{-}_e \phi = K \dot{-}_e \phi \wedge \psi$ or $K \dot{-}_e \psi = K \dot{-}_e \phi \wedge \psi$.

   Assume first that $Y = K \dot{-}_e \phi = K \dot{-}_e \phi \wedge \psi$.

   Since $Y \leq X$ and $\phi \wedge \psi \notin X$, $\phi \wedge \psi \notin Y$, condition 2(b) requires that $X = Y$.

   In the same manner, for the case $X = K \dot{-}_e \phi = K \dot{-}_e \phi \wedge \psi$ we obtain that $X = Y$.
3. Connectedness: We need to show for every X, Y that $X \leq Y$ or $Y \leq X$. Assume towards a contradiction that there are X, Y where $X < Y$ and $Y < X$.

   Since $X < Y$ we have $X \leq Y$ and $X \neq Y$. From condition 2(b) we have that for every δ such that $\delta \notin Y$ and $\delta \notin X$, if $X = K \dot{-}_e \delta$ then $X = Y$. Since $X \neq Y$ by assumption, this means that for every δ such that $\delta \notin Y$ and $\delta \notin X$, then $X \neq K \dot{-}_e \delta$.

   Since $Y < X$, an analogous argument gives that for every δ such that $\delta \notin Y$ and $\delta \notin X$, then $Y \neq K \dot{-}_e \delta$.

   By condition 2(a) there are formulas ϕ, ψ such that $X = K \dot{-}_e \phi$ and $Y = K \dot{-}_e \psi$. It follows that $\phi \wedge \psi \notin X$, $\phi \wedge \psi \notin Y$,

   Since $X < Y$, this means that $X \neq K \dot{-}_e (\phi \wedge \psi)$; since $Y < X$, this means that $Y \neq K \dot{-}_e (\phi \wedge \psi)$.

   But Theorem 5.1 requires that $X = K \dot{-}_e \phi = K \dot{-}_e (\phi \wedge \psi)$ or $Y = K \dot{-}_e \psi = K \dot{-}_e (\phi \wedge \psi)$, contradiction.
4. Transitivity: Assume that $X \leq Y$, and $Y \leq Z$; we need to show that $X \leq Z$.

   If $Z = K$ then our result follows from condition 1 of the definition of $\leq$.

   If $Y = Z$ then our result follows directly from conditions 2(a) and 2(b).

   If $X = Y$ then our result follows directly from conditions 2(a) and 2(b).

   Consequently, assume that $Z \neq K$, $Y \neq Z$, $X \neq Y$.

   In order to show that $X \leq Z$, by the preceding assumption, condition 1 doesn't apply; so we need to show conditions 2(a) and 2(b). Condition 2(a) follows from the fact that $Y \leq Z$.

   To show condition 2(b) assume that there is $\phi \in K$ such that $\phi \notin X$ and $\phi \notin Z$. As well, towards a contradiction assume that $X = K \dot{-}_e \phi$.

   Since $Y \leq Z$ there is ψ such that $Z = K \dot{-}_e \psi$.

   Since $X \leq Y$ there is δ such that $Y = K \dot{-}_e \delta$.

   Theorem 5.1 gives that $K \dot{-}_e \psi = K \dot{-}_e (\psi \wedge \delta)$ or $K \dot{-}_e \delta = K \dot{-}_e (\psi \wedge \delta)$. Condition 2(b) implies that in fact $K \dot{-}_e \psi = K \dot{-}_e (\psi \wedge \delta)$.

   Turning to $X \leq Y$, we have from before that for ϕ, δ as above, $X = K \dot{-}_e \phi$ and $Y = K \dot{-}_e \delta$. Applying Theorem 5.1 we have that $K \dot{-}_e \delta = K \dot{-}_e (\phi \wedge \delta)$ or $K \dot{-}_e \phi = K \dot{-}_e (\phi \wedge \delta)$. Condition 2(b) implies that $K \dot{-}_e \delta = K \dot{-}_e (\phi \wedge \delta)$.

   Applying the same argument to $Z = K \dot{-}_e \psi = K \dot{-}_e (\psi \wedge \delta)$ and $Y = K \dot{-}_e \delta = K \dot{-}_e (\phi \wedge \delta)$, and applying Theorem 5.1 gives that $K \dot{-}_e (\psi \wedge \delta) = K \dot{-}_e (\phi \wedge \psi \wedge \delta)$ or $K \dot{-}_e (\phi \wedge \delta) = K \dot{-}_e (\phi \wedge \psi \wedge \delta)$ and condition 2(b) requires that in fact $K \dot{-}_e (\psi \wedge \delta) = K \dot{-}_e (\phi \wedge \psi \wedge \delta)$, hence $Z = K \dot{-}_e (\phi \wedge \psi \wedge \delta)$

   By assumption we have that $\phi \notin Z$, and so $\phi \notin K \dot{-}_e (\phi \wedge \psi \wedge \delta)$.

   Postulate ($K \dot{-}_e 8e$) then gives that $K \dot{-}_e (\phi \wedge \psi \wedge \delta) = K \dot{-}_e \phi$. But $K \dot{-}_e \phi = X$ and $X \neq Z$, contradiction.

   Hence our original assumption, that $X = K \dot{-}_e \phi$ is false, and it follows that condition 2(b) is satisfied.

Last, we observe that for X, Y such that $\exists\phi$ such that $Y = K \dot{-}_e \phi$ and $\not\exists\phi$ such that $X = K \dot{-}_e \phi$ we have that $X \leq Y$ but not $Y \leq X$. It then follows from the definitions of γ and $\leq$ that for every ϕ, $K \dot{-}_e \phi = X$ where $X \not\vdash \phi$ and for every $Y > X$ we have $Y \vdash \phi$. ■