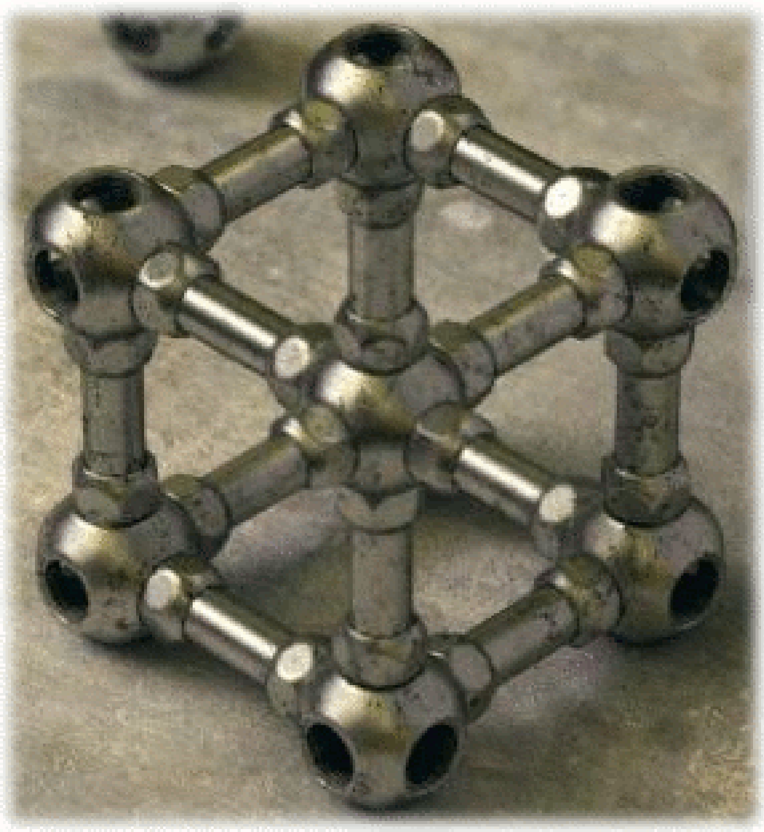

Horn Clause Contraction Functions: Belief Set and Belief Base Approaches [1]

Author(s):

James P. Delgrande

Renata Wassermann

[1]This work was supported by Fapesp Project LogProb, grant 2008/03995-5, São Paulo, Brazil.

Horn Clause Contraction Functions: Belief Set and Belief Base Approaches

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

Renata Wassermann
Department of Computer Science
University of São Paulo
05508-900 São Paulo,
Brazil
renata@ime.usp.br

Abstract

Standard approachs to belief change assume that the underlying logic contains classical propositional logic. Recently there has been interest in investigating approaches to belief change, specifically contraction, in which the underlying logic is not as expressive as full propositional logic. In this paper we consider approaches to belief contraction in Horn knowledge bases. We develop two broad approaches for Horn contraction, corresponding to the two major approaches in belief change, based on Horn *belief sets* and Horn *belief bases*. We argue that previous approaches, which have taken Horn remainder sets as a starting point, have undesirable properties, and moreover that not all desirable Horn contraction functions are captured by these approaches. This is shown in part by examining model-theoretic considerations involving Horn contraction. For Horn belief set contraction, we develop an account based in terms of *weak remainder sets*. Maxichoice and partial meet Horn contraction is specified, along with a consideration of package contraction. Following this we consider Horn belief base contraction, in which the underlying knowledge base is not necessarily closed under the Horn consequence relation. Again, approaches to maxichoice and partial meet belief set contraction are developed. In all cases, constructions of the specific operators and sets of postulates are provided, and representation results are obtained. As well, we show that problems arising with earlier work are resolved by these approaches.

Introduction

Belief change is the area of knowledge representation concerned with how a rational agent may alter its beliefs in the presence of new information. The best-known approach in this area is the so-called AGM paradigm (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988), named after the original developers. This work focussed on belief *contraction*, in which an agent may reduce its stock of beliefs, and belief *revision*, in which new information is consistently incorporated into its belief corpus. This work addresses belief change at the *knowledge level*, in which an agent's beliefs are characterised by *belief sets* or deductively closed sets of sentences, and in which the underlying logic includes classical propositional logic. A second major branch of belief change research concerns *belief bases* (Hansson 1999), wherein an agent's beliefs may not be deductively closed. However, again it is assumed that the underlying logic includes propositional logic.

In this paper we address 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, expressed in conjunctive normal form, a Horn clause will have at most one positive literal.) Specifically, in our approaches an agent's beliefs are represented by a Horn clause knowledge base, and the input is a conjunction of Horn clauses. This topic is interesting for several reasons. It sheds light on the theoretical underpinnings of belief change, in that it weakens the assumption that the underlying logic contains propositional logic. As well, Horn clauses have found extensive use in artificial intelligence and database theory, in areas such as logic programming, truth maintenance systems, and deductive databases. Further, as (Booth, Meyer, and Varzinczak 2009) points out, results here are also relevant to belief change in description logics, a topic that has also elicited recent interest. Last, belief change in Horn theories proves to be interesting in its own right.

Horn clause contraction has been addressed previously in (Delgrande 2008; Booth, Meyer, and Varzinczak 2009). As we discuss in the next section, this work centres on the notion of a *remainder set*, or maximal subset of a knowledge base that fails to imply a given formula. We show that remainder sets in the Horn case are too restricted and cannot give all feasible contraction operators. As well they yield contraction operators with undesirable properties.

We develop two broad approaches to Horn contraction, depending on whether the (Horn) knowledge base is regarded as a *belief set*, or deductively-closed set of formulas, or a *belief base*, i.e. an arbitrary set of Horn formulas. In the case of Horn belief sets we propose the notion of a *weak remainder set* that serves as a basis for generating all *maxichoice* contraction operators. Contraction is also considered in terms of the underlying model theory, a viewpoint that proves highly enlightening for studying Horn belief change. Given a specification for maxichoice contraction based on weak remainders, we go on to develop a specification for *partial meet* Horn contractions, and briefly consider package contraction.

For the second approach to Horn contraction, we address the case where a knowledge base is composed of an arbitrary set of Horn formulas. An advantage of this approach is that not only is contraction specified in a more realistic setting (i.e. closer to an implementation), it also allows for the nontrivial treatment of inconsistent knowledge bases. Belief base Horn contraction is considered both with respect to maxichoice and partial meet contraction operators. In all the contraction operators developed, we provide postulate sets along with constructions, and show representation results. Consequently we present a comprehensive exploration of the landscape of Horn contraction.

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 two sections give the approaches, based on belief sets and belief bases, respectively. The paper concludes with a discussion and a brief consideration of the prospects for belief revision in Horn clause theories. Proofs are given in an appendix.

Background: Belief Change

As mentioned previously, the *AGM approach* (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988) is the original and best-known approach to belief change. The goal is to describe belief change at the *knowledge level*, on an abstract level and 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 $\phi \in K$.

The central operators addressed are *contraction* in which an agent reduces its set of beliefs, and *revision* in which an agent consistently incorporates some new belief. In revision, since the new belief may be inconsistent with an agent's beliefs, some beliefs will need to be dropped in order to maintain a consistent set of beliefs. A third operator, belief *expansion* was also introduced: For belief set K and formula ϕ, the expansion of K by ϕ $K + \phi$ is the deductive closure of $K \cup \{\phi\}$.

These operators are characterised by two means. On the one hand, a set of *rationality postulates* for a belief change functions may be provided; these postulates stipulate constraints that should govern any rational belief change function. On the other hand, specific constructions for a belief change function are given. *Representation results* are then given (or at least are highly desirable) showing that a set of rationality postulates exactly captures the operator given by a particular construction.

We review these notions for belief contraction. Informally, the contraction of a belief set by a formula is a 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 $K\dot{-}\phi = K\dot{-}\psi$.

Thus, contraction 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 ($K\dot{-}2$), and if the formula is not originally believed then contraction has no effect ($K\dot{-}3$). 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 for example (Hansson 1999). The sixth postulate asserts that contraction is independent of how a sentence is expressed.

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. The original construction was in terms of *remainder sets*, where a remainder set of K with respect to ϕ is a maximal subset of K that fails to imply ϕ. Formally:

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

1. $K' \subseteq K$
2. $K' \nvdash \phi$
3. *For any K'' s.t. $K' \subset K'' \subseteq K$, it holds that $K'' \vdash \phi$.*

$X \in K \downarrow \phi$ is a remainder set of K wrt ϕ.

Two classes of contraction functions are relevant for our concerns. In *maxichoice contraction*, contraction is defined to correspond to a single selected remainder set. In partial meet contraction, contraction corresponds to the intersection of some subset of the remainder sets.

Belief Change and Horn Clause Theories

Earlier work on belief change and Horn theories focussed on specific aspects of the problem, rather than a general characterisation of Horn clause belief change. For example, the complexity of specific approaches to revising knowledge bases is addressed in (Eiter and Gottlob 1992), including the case where the knowledge base and formula for revision are conjunctions of Horn clauses. Not unexpectedly, 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 ϕ, both Horn, the main problem addressed is whether the knowledge base, revised according to a given operator, can be expressed by a propositional formula whose size is polynomial with respect to the sizes of K and ϕ.

(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*. In particular, they give a criterion for the existence of a contraction operator satisfying the basic AGM postulates in terms of *decomposability*.

(Delgrande 2008) addresses maxichoice belief contraction in Horn clause theories, where contraction is defined in terms of remainder sets, using Definition 1, but expressed in terms of derivations among Horn clauses. (Booth, Meyer, and Varzinczak 2009) further develops this area, by considering other versions of contraction, all based on remainder sets: partial meet contraction, a generalisation of partial meet, and package contraction. These approaches are discussed in more detail once we have introduced appropriate notation and definitions.

Horn Clause Theories

Preliminary Considerations

We will deal with languages based on finite sets of atoms, or propositional letters $\mathbf{P} = \{a, b, c, \dots\}$, where $\mathbf{P}$ includes the distinguished atom $\perp$. $\mathcal{L}$ is the language of propositional logic over $\mathbf{P}$ and with the usual connectives $\neg$, $\wedge$, $\vee$, and $\rightarrow$.[1] $\mathcal{L}_{HC}$ is the restriction of $\mathcal{L}$ to *Horn formulas*, or conjunctions of Horn clauses. I.e. $\mathcal{L}_{HC}$ is given by:

1. Every $p \in \mathbf{P}$ is a Horn clause.
2. $a_1 \wedge a_2 \wedge \cdots \wedge a_n \rightarrow a$, where $n \geq 0$, and a, a_i $(1 \leq i \leq n)$ are atoms, is a Horn clause.
3. Every Horn clause is a Horn formula.
4. If ϕ and ψ are Horn formulas then so is $\phi \wedge \psi$.

For 1 above and (equivalently) the case $n = 0$ in 2, the Horn formula is a *fact*. For a rule r as in 2 above, $head(r)$ is a, and $body(r)$ is the set $\{a_1, a_2, \dots, a_n\}$. Allowing conjunctions of rules, as given in 4, 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.

Semantics: An *interpretation* of $\mathcal{L}$ is a function from $\mathbf{P}$ to $\{true, false\}$ such that $\perp$ is assigned $false$. Sentences of $\mathcal{L}$ are *true* or *false* in an interpretation according to the standard rules in propositional logic. An interpretation M is a *model* of a sentence ϕ (or set of sentences), written $M \models \phi$, just if M makes ϕ true. $Mod(\phi)$ is the set of models of formula (or set of formulas) ϕ; thus $Mod(\top)$ is the set of interpretations of $\mathcal{L}$. An interpretation is usually identified with the atoms true in that interpretation. Thus, for language $\mathcal{L} = \{p, q, r, s\}$ the interpretation given by $\{p, q\}$ is that in which p and q are true and r and s are false. For convenience, we also will express interpretations by juxtaposition of atoms. Thus the interpretations $\{\{p, q\}, \{p\}, \{\}\}$ will usually be written as $\{pq, p, \emptyset\}$.

All of these notions are inherited by the corresponding Horn formula language $\mathcal{L}_{HC}$. A key point is that Horn theories are characterised semantically by the fact that the models of a Horn theory are closed under intersections of positive atoms in an interpretation. That is, Horn theories satisfy the constraint:

If $M_1, M_2 \in Mod(H)$ then $M_1 \cap M_2 \in Mod(H)$.

This leads to the notion of the *characteristic models* (Khardon 1995) of a Horn theory: M is a characteristic model of theory H just if for every $M_1, M_2 \in Mod(H)$, $M_1 \cap M_2 = M$ implies that $M = M_1$ or $M = M_2$. E.g. $H = Cn(\{p \wedge q \rightarrow \perp, r\})$, has models $\{pr, qr, r\}$ and characteristic models $\{pr, qr\}$. Since $pr \cap qr = r$, r isn't a characteristic model of H.

Proof Theory: We assume a suitable inference relation $\vdash$ for classical propositional logic. The following axioms and rules give an inference relation for Horn formulas, where for simplicity, a and b, possibly subscripted, are taken as ranging over atoms.

Axioms: $\perp \rightarrow a$ $\quad$ $a \rightarrow a$

Rules:
1. From $a_1 \wedge \cdots \wedge a_n \rightarrow a$ and $b_1 \wedge \cdots \wedge b_n \rightarrow a_i$ infer $a_1 \wedge \cdots \wedge a_{i-1} \wedge b_1 \wedge \cdots \wedge b_n \wedge a_{i+1} \wedge \cdots \wedge a_n \rightarrow a$
2. From $a_1 \wedge \cdots \wedge a_n \rightarrow a$ infer $a_1 \wedge \cdots \wedge a_n \wedge b \rightarrow a$
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 simply states that the order of atoms in the body of a rule is irrelevant, as are repeated atoms. 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 $A \vdash \perp$. We use $\phi \leftrightarrow \psi$ to represent logical equivalence, that is $\phi \vdash \psi$ and $\psi \vdash \phi$.

Notation: We collect here for reference notation that is used in the paper. Lower-case Greek characters ϕ, ψ, $\dots$, possibly subscripted, denote arbitrary formulas of either $\mathcal{L}$ or $\mathcal{L}_{HC}$. Upper case Roman characters A, B, $\dots$, possibly subscripted, denote arbitrary sets of formulas. H (H_1, H', etc.) denotes Horn belief sets, so that $\phi \in H$ iff $H \vdash_{HC} \phi$.

$Cn(A)$ is the (classical, propositional) deductive closure of A where A is a formula or set of formulas of propositional logic. $Cn^h(A)$ is the deductive closure of a Horn formula or set of formulas A under Horn derivability. For set of formulas A, $Horn(A) = \{\phi \in A \mid \phi \text{ is a Horn formula}\}$. $|\phi|$ is the set of maximal, consistent Horn theories that contain ϕ

[1] To avoid clutter, and because no ambiguity results, we don't parameterize $\mathcal{L}$ by $\mathbf{P}$.

and $|\neg\phi|$ is the complement of $|\phi|$, i.e., the set of maximal, consistent Horn theories that do not contain ϕ.

M (M_1, M', etc.) will denote (classical, propositional) interpretations over some understood language. $Mod(A)$ is the set of models of A. Arbitrary sets of interpretations will be denoted $\mathcal{M}$ ($\mathcal{M}'$ etc.). $Cl_\cap(\mathcal{M})$ is the intersection closure of a set of interpretations $\mathcal{M}$;[2] i.e. $\mathcal{M} \subseteq Cl_\cap(\mathcal{M})$ and $M_1, M_2 \in Cl_\cap(\mathcal{M})$ implies that $M_1 \cap M_2 \in Cl_\cap(\mathcal{M})$.

Since **P** is finite, a (Horn or propositional logic) belief set may be finitely represented, that is, for X a belief set, there is a formula ϕ such that $Cn(\phi) = X$. As well, we make use of the fact that there is a 1-1 correspondence between elements of $|\phi|$ and of $Mod(\phi)$.

Horn Clause Contraction and Remainder Sets

(Delgrande 2008) addressed maxichoice Horn belief set contraction based on (Horn) remainder sets, while (Booth, Meyer, and Varzinczak 2009) further investigated this area by considering partial meet contraction, a generalisation of partial-meet, based on the idea of *infra-remainder sets*, as well as package contraction, again based on remainder sets. Unfortunately, as we show next, it turns out that remainder sets (including the infra-remainder sets of (Booth, Meyer, and Varzinczak 2009)) are not sufficiently expressive for contraction; as well, contraction based on remainder sets can be shown to have undesirable properties.

The definition of *e-remainder sets* for Horn clause belief sets is the same as that for a remainder set (Definition 1) but with respect to Horn clauses and Horn derivability. For H a Horn belief set and $\phi \in \mathcal{L}_{HC}$, the set of *e-remainder sets* with respect to H and ϕ is denoted by $H \downarrow_e \phi$.

Observation 1 *If $H \downarrow_e \alpha_1 = H \downarrow_e \alpha_2$, then for any $H' \subseteq H$, $\alpha_1 \in Cn^h(H')$ iff $\alpha_2 \in Cn^h(H')$.*

Observation 2 *(Upper bound property) If $X \subseteq H$ and $\alpha \notin Cn^h(X)$, then there is some X' such that $X \subseteq X' \in H \downarrow_e \alpha$.*

(Booth, Meyer, and Varzinczak 2009) define *infra remainder sets* as follows:

Definition 2 *For belief sets H and X, $X \in H \Downarrow_e \phi$[3] iff there is some $X' \in H \downarrow_e \phi$ such that $(\bigcap H \downarrow_e \phi) \subseteq X \subseteq X'$. The elements of $H \Downarrow_e \phi$ are the* infra e-remainder sets *of H with respect to ϕ.*

All e-remainder sets are clearly infra e-remainder sets, as is the intersection of any set of e-remainder sets.

Example 1 *For $L = \{a, b, c\}$, let $H = Cn^h(a \wedge b)$, and where we consider candidates for $H \dot{-} (a \wedge b)$. There are three remainder sets, given by the Horn closures of $a \wedge (c \rightarrow b)$, $b \wedge (c \rightarrow a)$, and $(a \rightarrow b) \wedge (b \rightarrow a) \wedge (c \rightarrow a \wedge b)$). Any infra-remainder set must contain the closure of $(c \rightarrow a) \wedge (c \rightarrow b)$.*

[2] Recall that an interpretation is represented by the set of atoms true in the interpretation.

[3] (Booth, Meyer, and Varzinczak 2009) writes $X \in H \Downarrow_e \Phi$ where Φ is a *set* of Horn clauses.

counter-model	induced models	resulting KB	r.s.
ac	a	a	
a		$a \wedge (c \rightarrow b)$	$\surd$
bc	b	b	
b		$b \wedge (c \rightarrow a)$	$\surd$
c	$\emptyset$	$(a \rightarrow b) \wedge (b \rightarrow a)$	
$\emptyset$		$(a \rightarrow b) \wedge (b \rightarrow a) \wedge (c \rightarrow a \wedge b)$	$\surd$

Figure 1: Example: Candidates for Horn contraction

The fact that in the example any (infra-)remainder set contains $c \rightarrow a$ and $c \rightarrow b$. is not, on reflection, surprising: In the case of $c \rightarrow a$, since the original belief set contains a it also contains $c \rightarrow a$. A remainder set may not contain a, but due to the requirement of maximality, there is no reason to remove $c \rightarrow a$, and so $c \rightarrow a$ remains in any remainder set (or infra-remainder set). As we discuss below, this leads to some undesirable properties.

However, it is instructive to first consider remainder sets, and with them Horn contraction, from the point of view of the model theory. Assume that $H \models \phi$ and we wish to find a maximal belief set H' such that $H' \subset H$ and $H' \not\models \phi$. So H' will be a remainder set of H and ϕ. In classical AGM (maxichoice) contraction, from the semantic side one essentially adds a countermodel of ϕ to the models of H; this set characterises a candidate theory for maxichoice contraction. Consider the analogous process for Horn theories. Since we want a remainder set to be a Horn theory and the models of a Horn theory are closed under intersection, we would need to make sure that this constraint holds here. So, intuitively, to carry out maxichoice Horn contraction, we would add a countermodel of the formula for contraction, and close the result under intersections. However, critically, the theories resulting from this approach do not correspond to those obtained via remainder sets, and so do not correspond to maxichoice e-contraction as defined in (Delgrande 2008). To see this, consider again Example 1, and where the pertinent results are summarised in Figure 1.

We have that ac (viz. $\{a, \neg b, c\}$) is a countermodel of H; this is given in the first entry of the first row of the table. Since H has a model ab, the intersection of these models, $ab \cap ac = a$ must also be included; this is the item in the second column. The resulting belief set is characterised by the interpretations $Mod(H) \cup \{ac, a\} = \{abc, ab, ac, a\}$, which is the set of models of formula a, given in the third column. The result isn't a remainder set, since $Cn^h(a \wedge (c \rightarrow b))$ is a logically stronger belief set that fails to imply $a \wedge b$.

As previously noted, there are three remainder sets, as indicated in the last column. This result is problematic for both (Delgrande 2008) and (Booth, Meyer, and Varzinczak 2009). For example, in none of the approaches in these papers is it possible to obtain $H \dot{-}_e (a \wedge b) \leftrightarrow a$, nor is it possible to obtain $H \dot{-}_e (a \wedge b) \leftrightarrow (a \equiv b)$. But presumably these possibilities are desirable as *potential* contractions. To sharpen this point, in all of the approaches developed in the cited papers, it is not possible to have a contraction wherein $a \wedge \neg b \wedge c$ corresponds to a model of the contraction.

The diagnosis of the problem is clear. In the example, and for the countermodel given by $a \wedge \neg b \wedge c$, it is not possible to have a set of interpretations $\mathcal{M}$ satisfying:

1. $\mathcal{M}$ is closed under intersections
2. $Mod(H) \subseteq \mathcal{M}$
3. $\{a, \neg b, c\} \in \mathcal{M}$ (whence $\mathcal{M} \not\models a \wedge b$), and
4. $\mathcal{M}$ is a minimal set of interpretations satisfying 1, 2, 3, and $\mathcal{M} \not\models a \wedge b$.

The solution also seems clear: From a semantic point of view, one wants the *characteristic models* of maxichoice candidates for $H \dot{-}_e \phi$ to consist of the *characteristic models* of H together with a single interpretation from $Mod(\top) \setminus Mod(\phi)$. The resulting theories, called *weak remainder sets*, would correspond to the theories given in the third column in Figure 1.

Before considering possible ways to (re)define e-contraction, we note also that contraction based on remainder sets alone has undesirable properties. First, it has been pointed out[4] that maxichoice e-contraction suffers from a triviality result analogous to that in AGM contraction. As well, for contraction (or package contraction) defined in terms of remainder sets, or intersections of remainder sets, or infra remainder sets, we have the result:

For p not mentioned in H, we have $(H \dot{-} \phi) + p \ \vdash \ \phi$.

The proof is straightforward; it is omitted due to space constraints, but is included in the full paper. Here is an illustrative example of this phenomenon (with apologies to (Hansson 1999)), in terms of package contraction:

1. You believe Cleopatra had a son and a daughter ($s \wedge d$).
2. You learn that the source of information was unreliable, so you remove this belief; i.e. you compute the package contraction $H \dot{-} \{s, d\}$.
3. You learn that it is raining outside (r).
4. You conclude that Cleopatra had a son and daughter ($s \wedge d$)

Horn Clause Belief Set Contraction

The previous section showed that basing Horn contractions solely on remainder sets (or infra-remainder sets) is problematic. We then suggested that an adequate version of contraction should be based on *weak remainder sets* where for belief set H and formula ϕ, there is a 1-1 correspondence between countermodels of ϕ and weak remainder sets. In this section we develop Horn contraction based on weak remainder sets. We first give two constructions for weak remainder sets, in terms of belief sets and in terms of sets of models, and show the constructions equivalent. We then characterise maxichoice Horn contraction in terms of weak remainder sets, showing via a representation result that the characterisations are equivalent. Following this we similarly characterise partial meet contraction, and briefly consider package contraction. We note that due to the added generality of weak remainder sets, the aforementioned triviality results do not hold in any of the approaches developed.

[4]David Makinson, personal communication

Definition 3 *Let H be a Horn belief set, and let ϕ be a Horn formula.*[5]

$H \Downarrow_e \phi$ is the set of sets of formulas s.t. $H' \in H \Downarrow_e \phi$ iff $H' = H \cap m$ for some $m \in |\neg\phi|$.

$H' \in H \Downarrow_e \phi$ is a weak remainder set *of H and ϕ.*

Example 2 *For $\mathbf{P} = \{a, b, c\}$, let $H = Cn^h(a \wedge b)$, $\phi = a \wedge b$. For $m = Cn^h(a \wedge \neg b \wedge c) \in |\neg\phi|$, we have that $H \cap m = Cn^h(a)$, since H and m are both closed under Horn consequence. (Note that full propositional closure gives $Cn(H \cap m) = Cn(a \wedge (b \vee c))$.)*

Definition 4 *Let H be a Horn belief set, and let ϕ be a Horn formula. $H \parallel_e \phi$ is the set of sets of formulas s.t. $H' \in H \parallel_e \phi$ iff there is $M \notin Mod(\phi)$ s.t. $Mod(H') = Cl_\cap(Mod(H) \cup \{M\})$.*

In our running example, $H \parallel_e \phi$ is given by the closure of the sets of formulas in column 3 in Figure 1.

Theorem 1 *For H a Horn belief set and ϕ a Horn formula:*
$H \Downarrow_e \phi \ = \ H \parallel_e \phi$.

Definition 5 *Let H be a Horn belief set. γ is a* selection function *for H if, for every $\phi \in \mathcal{L}_{HC}$,*

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

Definition 6 *Let γ be a selection function on H such that $\gamma(H \Downarrow_e \phi) = \{H'\}$ for some $H' \in H \Downarrow_e \phi$.*

The maxichoice Horn contraction *based on weak remainders is given by:*

$$H \dot{-}_w \phi = \gamma(H \Downarrow_e \phi)$$

We obtain the following representation result, relating the constructions to a postulate set characterising contraction:

Theorem 2 *Let H be a Horn belief set. Then $\dot{-}_w$ is an operator of maxichoice Horn contraction based on weak remainders iff $\dot{-}_w$ satisfies the following postulates.*

($H\dot{-}_w 1$) *$H \dot{-}_w \phi$ is a belief set.* *(closure)*

($H\dot{-}_w 2$) *If not $\vdash \phi$, then $\phi \notin H \dot{-}_w \phi$.* *(success)*

($H\dot{-}_w 3$) *$H \dot{-}_w \phi \subseteq H$.* *(inclusion)*

($H\dot{-}_w 4$) *If $\phi \notin H$, then $H \dot{-}_w \phi = H$.* *(vacuity)*

($H\dot{-}_w 5$) *If $\vdash \phi$ then $H \dot{-}_w \phi = H$* *(failure)*

($H\dot{-}_w 6$) *If $\phi \leftrightarrow \psi$, then $H \dot{-}_w \phi = H \dot{-}_w \psi$.* *(extensionality)*

($H\dot{-}_w 7$) *If $H \neq H \dot{-}_w \phi$ then $\exists \beta \in \mathcal{L}_{HC}$ s.t. $\{\phi, \beta\}$ is inconsistent, $H \dot{-}_w \phi \subseteq Cn^h(\{\beta\})$ and $\forall H'$ s.t $H \dot{-}_w \phi \subset H' \subseteq H$ we have $H' \not\subseteq Cn^h(\{\beta\})$.* *(maximality)*

Partial Meet Contraction: Partial meet contraction provides a general characterisation of, here, Horn contraction. The definition is analogous to that in AGM contraction, but based on weak remainder sets:

[5]Recall that a Horn formula is a conjunction of Horn clauses.

Definition 7 *Let γ be a selection function on H such that $\gamma(H \Downarrow_e \phi) \subseteq (H \Downarrow_e \phi)$.*

Then the partial meet Horn contraction *based on weak remainders is given by:*

$$H \dot{-}_{pm} \phi = \bigcap \gamma(H \Downarrow_e \phi)$$

A representation result involves a modification of the last postulate for maxichoice contraction:

Theorem 3 *Let H be a Horn belief set. Then $\dot{-}_w$ is an operator of partial meet Horn contraction based on weak remainders iff $\dot{-}_w$ satisfies the postulates $(H \dot{-}_w 1)$ – $(H \dot{-}_w 6)$ and:*

($H \dot{-}_{pm} 7$) *If $\beta \in H \setminus (H - \alpha)$, then there is some H' such that $H - \alpha \subseteq H'$, $\alpha \notin Cn^h(H')$ and $\alpha \in Cn^h(H' \cup \{\beta\})$ (weak relevance)*

Example 3 *For our running example, the partial meet given by the first and last weak remainder sets in Figure 1 is given by $Cn^h((b \to a) \wedge (c \to a))$. In terms of models, it is given by the models of $a \wedge b$, together with the two countermodels given by atoms ac and $\emptyset$, and closed under intersections.*

Package Contraction: Given its usefulness in Horn clauses knowledge bases, we briefly consider package contraction next. For belief set H and a set of formulas Φ, the package contraction $H \dot{-}_{pa} \Phi$ is a form of contraction in which no member of Φ is in $H \dot{-}_{pa} \Phi$.

We define a notion of Horn package contraction, and show that it is expressible in terms of maxichoice Horn contraction. Due to space limitations, we defer additional details to the full paper; as well, the full paper addresses a stronger version of package contraction where single countermodels of all members of Φ are added, where possible.

Definition 8 *Let H be a Horn belief set, and let Φ be a set of Horn formulas.*

$H \Downarrow_p \Phi$ is the set of sets of formulas s.t. $H' \in H \Downarrow_p \Phi$ iff

1. *$H' \subseteq H$, and*
2. *for every $\phi \in \Phi$ where $\phi \notin Cn^h(\top)$, $H' \subseteq m$ for some $m \in |\neg\phi|$, and*
3. *for every H'' where $H' \subset H'' \subseteq H$, we have $H'' \not\subseteq m$ for some $\phi \in \Phi$ where $m \in |\neg\phi|$.*

Condition 2 states that for every $\phi \in \Phi$ that is not a tautology, some countermodel of ϕ is in $Mod(H')$. The third condition states that H' is a largest subset of H that satisfies Condition 2. In the next definition, the notion of a selection function on H (Definition 5) is extended in the obvious fashion to apply to a set of Horn formulas.

Definition 9 *Let γ be a selection function on H such that $\gamma(H \Downarrow_p \Phi) = \{H'\}$ for some $H' \in H \Downarrow_p \Phi$.*

The (maxichoice) package Horn contraction *based on weak remainders is given by:*

$$H \dot{-}_{pa} \Phi = \gamma(H \Downarrow_p \Phi)$$

if $\emptyset \neq \Phi \cap H \not\subseteq Cn^h(\top)$; and H otherwise.

The following result relates elements of $H \Downarrow_p \Phi$ to weak remainders.

Theorem 4 *Let H be a Horn belief set and let $\Phi = \{\phi_1, \ldots, \phi_n\} \subset \mathcal{L}_{HC}$. We have that $H' \in H \Downarrow_p \Phi$ iff $H' = \bigcap_{i=1}^{n} H_i$ where $H_i \in H \Downarrow_e \phi_i$, $1 \leq i \leq n$.*

It follows from this that any maxichoice Horn contraction defines a package contraction, and vice versa.

Example 4 *Consider a variant on our running example where as before $\mathbf{P} = \{a, b, c\}$, $H = Cn^h(a \wedge b)$. Then among candidates for $H \dot{-}_{pa} \{a, b\}$ we have $Cn^h((c \to b) \wedge (a \to b))$ and $Cn^h(a \equiv b)$.*

Horn Clause Belief Base Contraction

In this section we turn our attention to contraction of *Horn belief bases*, sets of Horn formulas not necessarily closed under Horn consequence. We first note that the definition of e-remainders can be used directly for Horn belief bases. (Makinson 1987) has already defended maxichoice contraction on the grounds that it is not a construction that should be applied to belief sets, but only to belief bases. The operation of maxichoice defined in (Delgrande 2008), when applied to arbitrary sets of Horn formulas does not suffer from the same drawbacks as when applied to belief sets. Let us look at a slightly adapted version of Example 1:

Example 5 *For $L = \{a, b, c\}$, let $B_1 = \{a, b\}$, $B_2 = \{a, b, a \to b, b \to a\}$, and $B_3 = \{a, b, c \to a, c \to b\}$ and again we consider candidates for $B_i \dot{-} (a \wedge b)$. Although the three bases represent the same information at the belief set level (i.e, $Cn^h(B_1) = Cn^h(B_2) = Cn^h(B_3)$), the choice of which beliefs to represent explicitly leads to different results:*

$B_1 \downarrow_e a \wedge b = \{\{a\}, \{b\}\}$
$B_2 \downarrow_e a \wedge b = \{\{a, b \to a\}, \{b, a \to b\}, \{a \to b, b \to a\}\}$
$B_3 \downarrow_e a \wedge b = \{\{a, c \to a, c \to b\}, \{b, c \to a, c \to b\}\}$

Only in the last case we have that independently of the selection function, $B_3 \subseteq (B_3 \dot{-} a \wedge b) + c$.

For classical logic, maxichoice contraction satisfies the following postulate:

If $\beta \in B \setminus (B \dot{-} \alpha)$, then $\alpha \notin Cn(B \dot{-} \alpha)$ and $\alpha \in Cn(B \dot{-} \alpha \cup \{\beta\})$ (fullness)

We can prove the following result for Horn belief bases:

Theorem 5 *The operation $\dot{-}_e$ is an operator of maxichoice e-contraction on B if and only if for all sentences α:*

- *If $\alpha \notin Cn^h(\emptyset)$, then $\alpha \notin Cn^h(B \dot{-}_e \alpha)$ (success)*
- *$B \dot{-}_e \alpha \subseteq B$ (inclusion)*
- *If $\beta \in B \setminus (B \dot{-}_e \alpha)$, then $\alpha \notin Cn^h(B \dot{-}_e \alpha)$ and $\alpha \in Cn^h(B \dot{-}_e \alpha \cup \{\beta\})$ (fullness)*
- *If for all subsets B' of B $\alpha \in Cn^h(B')$ if and only if $\beta \in Cn^h(B')$, then $B \dot{-}_e \alpha = B \dot{-}_e \beta$ (uniformity)*

The proof is a simple generalization of Hansson's original proof for classical propositional logic to the Horn case.

Although for belief bases maxichoice and its characterising postulate fullness do not lead to trivialisation as is the case for belief sets, it is interesting to look at the general case of partial meet contraction. We can prove the following representation result:

Theorem 6 *Let B be an arbitrary set of Horn formulas. Then $\dot{-}$ is an operator of partial meet contraction on B if and only if for all Horn formulas α:*

- *If $\alpha \notin Cn^h(\emptyset)$, then $\alpha \notin Cn^h(B\dot{-}\alpha)$* *(success)*
- *$B\dot{-}\alpha \subseteq B$* *(inclusion)*
- *If $\beta \in B \setminus (B\dot{-}\alpha)$, then $\exists B'$ such that $B\dot{-}\alpha \subseteq B' \subseteq B$, $\alpha \notin Cn^h(B')$ and $\alpha \in Cn^h(B' \cup \{\beta\})$* *(relevance)*
- *If for all subsets B' of B $\alpha \in Cn^h(B')$ if and only if $\beta \in Cn^h(B')$, then $B\dot{-}\alpha = B\dot{-}\beta$* *(uniformity)*

The theorem follows from the results in (Hansson and Wassermann 2002), where it was shown that it holds for every underlying logic which is compact and monotonic. From this, we know that partial-meet e-contraction satisfies relevance, which is a weaker form of recovery. It states that if a belief is removed from the belief base, it was involved in some derivation of the contracted sentence.

Discussion of Related Work

This section summarising the technical differences between the different operations defined on Horn belief sets:

- *Every e-remainder is a weak remainder, but the converse is not true.*

This is clearly seen in Figure 1. For a Horn theory H and formula ϕ, the e-remainders are the maximal subsets of H that do not imply ϕ. The weak remainders are characterised by the models of H together with a single countermodel of ϕ, and then closed under intersection. In propositional logic these notions would coincide; here they do not.

As well, this means that weak remainders and partial meet are distinct notions, the latter corresponding to intersections of weak remainders.

- *Not all infra-remainders are weak-remainders.*

Looking again Figure 1, the set $Cn^h(\{c \to a, c \to b, a \to b\})$ is an infra-remainder but not a weak-remainder. It can however be obtained as the intersection of two remainders.

Consider Example 3.2 in (Booth, Meyer, and Varzinczak 2009), where $H = Cn^h(\{p \to q, q \to r\})$ and one wants to contract by $p \to r$: In this case, the weak remainders coincide with the remainders. The set $\{p \wedge q \to r, p \wedge r \to q\}$ is an infra-remainder and cannot be obtained as the intersection of weak-remainders. The authors claim that this set is a desirable result of the contraction, but do not give any strong motivation.

- *Not all weak remainders are infra-remainders.*

Infra-remainders, by definition, must contain full-meet and be contained in some remainder. Weak remainders are contained in some remainder (or are a remainder) but do not always contain full meet, as can be seen in the table in Figure 1. Full-meet in that example would contain $\{c \to a, c \to b\}$ and there are two weak remainders ($Cn^h(a)$ and $Cn^h(b)$) which do not contain both formulas.

The last two items show that weak remainders and infra-remainders are independent concepts and their relation should be studied in more detail.

Another point that deserves attention can be seen again in the example from (Booth, Meyer, and Varzinczak 2009): For $H = Cn^h(\{p \to q, q \to r\})$, we have $H \downarrow_e p \to r = H \Downarrow_e p \to r = \{Cn^h(\{p \to q\}), Cn^h(\{q \to r, p \wedge r \to q\})\}$. There is an asymmetry here - while it is possible to obtain $Cn^h(\{p \to q\})$ as the result of contraction, e-remainders, weak remainders or infra-remainders do not allow for $Cn^h(\{q \to r\})$ as a possible outcome. This has motivated the study of Horn belief base contraction, where one may obtain $Cn^h(\{q \to r\})$, and where we think we may find other interesting alternatives.

Conclusion

In this paper we have (further) explored the landscape of belief contraction in Horn knowledge bases. Approaches to maxichoice and partial meet belief contraction for both belief sets and belief bases were presented. As well, package contraction in the case of Horn belief sets was also considered. In the case of belief set contraction, it proved to be the case that founding contraction on remainder sets (as is done in propositional logic) is problematic, in that the resulting approach is inexpressive and has undesirable properties. Based on an examination of model-theoretic considerations we developed an account of maxichoice Horn contraction in terms of *weak remainder sets*. This account captures the full range of maxichoice contraction, and hence partial meet contraction. We also developed approaches to Horn belief base contraction, in which the underlying knowledge base is not necessarily closed under the Horn consequence relation. Such approaches are valuable, in that the result of contraction reflects the syntactic expression of the knowledge base, which in turn may better reflect the knowledge base designer's intentions. In all cases, constructions of the contraction operators were specified, along with sets of characterising postulates, and representation results were provided, linking the constructions and postulate sets. last, we showed that problems arising with earlier work are resolved by these approaches.

Acknowledgements

The first author was partially supported by a Canadian NSERC Discovery Grant. The second author was partially supported by the Brazilian National Research Council (CNPq) and by São Paulo State Research Foundation (FAPESP). We thank Tommie Meyer, Marcio Ribeiro and David Makinson, along with the reviewers of the paper, for their helpful comments.

Proofs of the main results

Lemma 1 *Let T be a set of propositional formulas. Then $Cl_\cap(Mod(T)) = Mod(Horn(Cn(T)))$.*

Proof: We have that $Cl_\cap(Mod(T))$ is the least set of models such that $Mod(T) \subseteq Cl_\cap(Mod(T))$ and where $Cl_\cap(Mod(T))$ specifies a Horn theory. But this is just the *least upper Horn approximation* of T (Selman and Kautz 1996). But the least upper Horn approximation of T is given by

$T^h = \{\alpha \mid T \vdash \alpha$ where α is a Horn prime implicate of $T\}$.

We have that $Cn^h(T^h) = Horn(Cn(T))$ from which the result follows. ■

Proof of Theorem 1

1. $H \Downarrow_e \phi \subseteq H \parallel_e \phi$:

   Let $H' \in H \Downarrow_e \phi$; we show that $H' \in H \parallel_e \phi$.

   Since $H' \in H \Downarrow_e \phi$, by definition $H' = H \cap m$ for some $m \in |\neg\phi|$, and so $Mod(H') = Mod(H \cap m)$. H and m are Horn theories, thus $H \cap m$ is a Horn theory,

   Using the fact that for Horn theory T, $T = Horn(Cn(T))$, we have that $H \cap m = Horn(Cn(H \cap m))$ and so $Mod(H') = Mod(Horn(Cn(H \cap m))$.

   Applying Lemma 1 to $H \cap m$ we obtain that $Mod(Horn(Cn(H \cap m)) = Cl_\cap(Mod(Cn(H \cap m)))$. Now, $Cl_\cap(Mod(Cn(H \cap m))) = Cl_\cap(Mod(H \cap m)) = Cl_\cap((Mod(H) \cup Mod(m)))$.

   By definition, there is $M \in Mod(\top)$ such that $Mod(m) = \{M\}$. Putting the above together we get that $Mod(H') = Cl_\cap((Mod(H) \cup M))$, that is, $H' \in H \parallel_e \phi$.

2. $H \parallel_e \phi \subseteq H \Downarrow_e \phi$:

   This part follows immediately by essentially taking the preceding part in reverse order. ■

Proposition 1 *Maximality ($H \dot{-}_w 7$) is equivalent to the following property, which we will call ($H \dot{-}_w 7'$):*

If $H \neq H \dot{-} \phi$ then $\exists m \in |\neg\phi|$ s.t. $H \dot{-} \phi \subseteq m$ and $\forall H'$ s.t. $H \dot{-} \phi \subset H' \subseteq H$ we have $H' \not\subseteq m$.

Proof: Assume first that the property holds. Let β be the conjunction of literals appearing in m. So $Cn^h(\beta) = m$, and so ($H \dot{-}_w 7$) holds.
For the other direction, assume that ($H \dot{-}_w 7$) holds.

Claim: For given H and ϕ, if β satisfies the given conditions in ($H \dot{-}_w 7$) then for any $p \in \mathbf{P}$, so does one of $\beta \wedge p$ or $\beta \wedge (p \rightarrow \bot)$.

> Proof of Claim: It is clear that if $\{\phi, \beta\}$ is inconsistent then so is $\{\phi, \beta \wedge l\}$ for $l \in \{p, p \rightarrow \bot\}$; and if $H - \phi \subseteq Cn^h(\beta)$ then $H' \subseteq Cn^h(\beta \wedge l)$ for $l \in \{p, p \rightarrow \bot\}$.
>
> So we just need to show that for Horn theory H' where $H - \phi \subset H' \subseteq H$, either $H' \not\subseteq Cn^h(\beta \wedge p)$ or $H' \not\subseteq Cn^h(\beta \wedge (p \rightarrow \bot))$.
>
> Towards a contradiction, assume otherwise. Then $H' \subseteq Cn^h(\beta \wedge p)$ and $H' \subseteq Cn^h(\beta \wedge (p \rightarrow \bot))$. But since $Cn^h(\beta) = Cn^h(\beta \wedge \alpha) \cap Cn^h(\beta \wedge \neg\alpha)$, and consequently $H' \subseteq Cn^h(\beta)$, this contradicts that β satisfies ($H \dot{-}_w 7$) for H and ϕ.
>
> Hence our assumption was incorrect, and so $H' \not\subseteq Cn^h(\beta \wedge p)$ or $H' \not\subseteq Cn^h(\beta \wedge (p \rightarrow \bot))$.

We have just shown that if β satisfies ($H \dot{-}_w 7$) for given H and ϕ, then so does one of $\beta \wedge p$ or $\beta \wedge (p \rightarrow \bot)$. An induction over (the finite set) $\mathbf{P}$ then establishes that if β satisfies ($H \dot{-}_w 7$) for given H and ϕ, then so does some β' where $\beta' \vdash p$ or $\beta' \vdash (p \rightarrow \bot)$ for every $p \in \mathbf{P}$. Hence β' is such that $Cn^h(\beta') \in |\neg\phi|$, and thus taking $m = Cn^h(\beta')$ satisfies the property. ■

Proof of Theorem 2:

1. Construction to Postulates:

   That the construction satisfies the first five postulates follows directly from the definitions of weak remainders and selection functions. To see that it satisfies ($H \dot{-}_w 6$) we only have to note that $\phi \leftrightarrow \psi$ implies that $H \Downarrow_e \phi = H \Downarrow_e \psi$ and since γ is a function, $H \dot{-}_w \phi = H \dot{-}_w \psi$.

   To see that the construction satisfies ($H \dot{-}_w 7$), suppose $H \neq H \dot{-}_w \phi$. This means that $H \Downarrow_e \phi \neq \emptyset$ and hence, there is $m \in |\neg\phi|$ such that $H \dot{-}_w \phi = H \cap m$. Let β be the conjunction of all literals appearing in m. Then, since $Cn^h(\beta) = m$, we have that $\{\phi, \beta\}$ is inconsistent, $H \dot{-}_w \phi \subseteq Cn^h(\{\beta\})$ and $\forall H'$ s.t $H \dot{-}_w \phi \subset H' \subseteq H$ we have $H' \not\subseteq Cn^h(\{\beta\})$.

2. Postulates to Construction:

   The proof uses ($H \dot{-}_w 7'$) rather than ($H \dot{-}_w 7$), as they were shown to be equivalent in Proposition 1 above.

   Let $\dot{-}_w$ be an operator that satisfies ($H \dot{-}_w 1$) – ($H \dot{-}_w 7'$). Let γ be defined by $\gamma(H \Downarrow_e \phi) = \{H \dot{-}_w \phi\}$.

   We have that γ is a function:

   > Assume that $H \Downarrow_e \phi = H \Downarrow_e \psi$; we need to show that $\gamma(H \Downarrow_e \phi) = \gamma(H \Downarrow_e \psi)$. If $\phi \notin H$, then $H \Downarrow_e \phi = \{H\}$ and since $H \Downarrow_e \phi = H \Downarrow_e \psi$, we have that $H \Downarrow_e \psi = H$ and hence $\psi \notin H$. Then, by ($H \dot{-}_w 4$), $H \dot{-}_w \phi = H \dot{-}_w \phi = H$ and by definition $\gamma(H \Downarrow_e \phi) = \gamma(H \Downarrow_e \psi)$.
   >
   > Now let us consider the case where $\phi, \psi \in H$. Since $H \Downarrow_e \phi = H \Downarrow_e \psi$ we have that $\{H \cap m \mid m \in |\neg\phi|\} = \{H \cap m \mid m \in |\neg\psi|\}$. It follows[6] that $|\neg\phi| = |\neg\psi|$, whence $|\phi| = |\psi|$ and so $\phi \leftrightarrow \psi$. From ($H \dot{-}_w 6$) we have $H \dot{-}_w \phi = H \dot{-}_w \psi$, and so $\gamma(H \Downarrow_e \phi) = \gamma(H \Downarrow_e \psi)$.

   If $\phi \notin H$, then from ($H \dot{-}_w 4$) we have that $H \dot{-}_w \phi = H$. Similarly, if $\vdash \phi$, then from ($H \dot{-}_w 5$) we again have that $H \dot{-}_w \phi = H$.

   Consequently assume that $\phi \in H$ and not $\vdash \phi$. We need to show that $H \dot{-}_w \phi \in H \Downarrow_e \phi$, that is, $H \dot{-}_w \phi = H \cap m$ for some $m \in |\neg\phi|$.

   Since not $\vdash \phi$, from ($H \dot{-}_w 2$) we have $\phi \notin H \dot{-}_w \phi$; since $\phi \in H$ we then have that $H \neq H \dot{-}_w \phi$.

   Since $H \neq H \dot{-}_w \phi$, from ($H \dot{-}_w 7'$) we get that there is $m \in |\neg\phi|$ such that $H \dot{-}_w \phi \subseteq m$.

   As well, ($H \dot{-}_w 3$) gives $H \dot{-}_w \phi \subseteq H$, and so this with $H \dot{-}_w \phi \subseteq m$ implies that $H \dot{-}_w \phi \subseteq (m \cap H)$.

   We need to show that $H \dot{-}_w \phi = (m \cap H)$. Towards a contradiction assume that $H \dot{-}_w \phi \neq (m \cap H)$, that is to say, $H \dot{-}_w \phi \subset (m \cap H)$.

[6] Suppose $\{H \cap m \mid m \in |\neg\phi|\} = \{H \cap m \mid m \in |\neg\psi|\}$ and $|\neg\phi| \neq |\neg\psi|$. Without loss of generality, suppose there is $m' \in |\neg\phi|$ such that $m' \notin |\neg\psi|$. Then m' is a maximal consistent theory that contains ψ. Since $\psi \in H$, we know that $\psi \in H \cap m'$. This means that $H \cap m' \in \{H \cap m \mid m \in |\neg\phi|\}$, but $H \cap m' \notin \{H \cap m \mid m \in |\neg\psi|\}$, as for any $m \in |\neg\psi|$ by definition $\psi \notin m$. This contradicts the initial hypothesis.

Let $\psi \in (m \cap H) \setminus (H \dot{-}_w \phi)$. Then

$$H \dot{-}_w \phi \subset Cn^h(H \dot{-}_w \phi \cup \{\psi\}) \subseteq m \cap H \subset H.$$

But, substituting $Cn^h(H \dot{-}_w \phi \cup \{\psi\})$ for H' in $(H \dot{-}_w 7)$ we get that $Cn^h(H \dot{-}_w \phi \cup \{\psi\}) \not\subseteq m$, contradiction.

Hence the assumption that $H \dot{-}_w \phi \neq (m \cap H)$ is incorrect; hence $H \dot{-}_w \phi = (m \cap H)$ where $(m \cap H) \in H \Downarrow_e \phi$, which was to be shown. ■

Proof of Theorem 3:

1. Construction to Postulates:

   $(H \dot{-}_w 1)$ follows from the fact that the intersection of Horn theories is a Horn theory. Postulates $(H \dot{-}_w 2)$ – $(H \dot{-}_w 6)$ follow immediately from the definitions of weak remainder, selection function and partial meet contraction.

   To see that the construction satisfies weak relevance, note that if $\beta \notin H \setminus H - \phi$, then there is some $X \in \gamma(H \Downarrow_e \phi)$ such that $\beta \notin X$. Since $\beta \in H$, then there is some $m \in |\neg\phi|$ such that $\beta \notin m$ and $X = H \cap m$. Take $H' = m$. Then $H - \phi \subseteq H'$, $\phi \notin Cn^h(H')$ and $\phi \in Cn^h(H' \cup \{\beta\}) = Cn^h(\bot)$.

2. Postulates to Construction:

   Let $\gamma(H \Downarrow_e \phi) = \{X \in H \Downarrow_e \phi \mid H - \phi \subseteq X\}$.

   We have to show that: (1) γ is a function; (2) γ is a selection function; and (3) $\bigcap \gamma(H \Downarrow_e \phi) = H - \phi$.

   (1) Let $H \Downarrow_e \phi_1 = H \Downarrow_e \phi_2$. We must show that $\gamma(H \Downarrow_e \phi_1) = \gamma(H \Downarrow_e \phi_2)$. As in the proof for maxichoice, $H \Downarrow_e \phi_1 = H \Downarrow_e \phi_2$ implies that $\vdash \phi_1 \leftrightarrow \phi_2$ and then, by postulate $(H \dot{-}_w 6)$, $H - \phi_1 = H - \phi_2$. By the construction of γ, $\gamma(H \Downarrow_e \phi_1) = \gamma(H \Downarrow_e \phi_2)$.

   (2) From the construction of γ, we know that $\gamma(H \Downarrow_e \phi) \subseteq H \Downarrow_e \phi$. So we have to show that if $H \Downarrow_e \phi \neq \emptyset$, then $\gamma(H \Downarrow_e \phi) \neq \emptyset$ and otherwise $\gamma(H \Downarrow_e \phi) = \{H\}$.

   (i) If $H \Downarrow_e \phi \neq \emptyset$, then $H \neq \emptyset$ and $|\neg\phi| \neq \emptyset$. By (H-2) and (H-1), $\phi \notin Cn(H - \phi)$. Then there is $m \in |\neg\phi|$ such that $H - \phi \subseteq m$. By (H-3), $H - \phi \subseteq H$, hence, $H - \phi \subseteq H \cap m \in \gamma(H \Downarrow \phi)$.

   (ii) If $H \Downarrow_e \phi = \emptyset$, then $\vdash \phi$ and by (H-5), $H - \phi = H$.

   (3) We know that $H - \phi \subseteq \bigcap \gamma(H \Downarrow_e \phi)$. Suppose there is $\beta \in \bigcap \gamma(H \Downarrow_e \phi)$ such that $\beta \notin H - \phi$. Since $\bigcap \gamma(H \Downarrow_e \phi) \subseteq H$, $\beta \in H \setminus (H - \alpha)$ and by weak relevance we know that there is some H' such that $H - \phi \subseteq H'$, $\phi \notin Cn^h(H')$ and $\phi \in Cn^h(H' \cup \{\beta\})$. Then there is $m \in |\neg\phi|$ such that $H' \subseteq m$ and $\beta \notin m$. Take $X = H \cap m$. Then $X \in H \Downarrow \phi$ and from (H-3) we have that $H - \phi \subseteq X$ and hence, $X \in \gamma(H \Downarrow_e \phi)$. But $\beta \notin X$, which leads to a contradiction. ■

Proof of Theorem 4: Let H be a Horn belief set and $\Phi = \{\phi_1, \ldots, \phi_n\} \subset \mathcal{L}_{HC}$.

1. $\Longleftarrow$:

   For each $1 \leq i \leq n$, let $H_i \in H \Downarrow_e \phi_i$. Then for $H' = \bigcap_{i=1}^n H_i$, to show that $H' \in H \Downarrow_p \Phi$ we show that H' satisfies the three conditions in Definition 8:

   1. Clearly $\bigcap_{i=1}^n H_i \subseteq H$, since we have that $H_i \subseteq H$ for $1 \leq i \leq n$.
   2. Consider $\phi \in \Phi$ where $\phi \notin Cn^h(\top)$. We have that $H' \subseteq H_i$ and $H_i \in H \Downarrow_e \phi_i$. From Definition 3 we have that $H_i = H \cap m$ for some $m \in |\neg\phi|$; hence $H' \subseteq m$ for that $m \in |\neg\phi|$.
   3. Let H'' be such that $H' \subset H'' \subseteq H$. If there is no such H'' then the third condition is satisfied vacuously.

      Let $\psi \in H''$ and $\psi \notin H'$. Thus $\psi \notin \bigcap_{i=1}^n H_i$, and so for some j, $1 \leq j \leq n$, we have that $\psi \notin H_j$.

      We have that $H_j = H \cap m$ for some $m \in |\neg\phi_j|$. Since $\psi \in H''$, so $\psi \in H$; hence $\psi \notin m$.

      Since $\psi \in H''$, $\psi \notin m$, we have that $H'' \not\subseteq m$.

      Since H'' was arbitrarily chosen, this shows that the third condition is satisfied.

   Thus $H' = \bigcap_{i=1}^n H_i$ satisfies the three conditions of Definition 8. Hence $H' \in H \Downarrow_p \Phi$.

2. $\Longrightarrow$:

   Let $H' \in H \Downarrow_p \Phi$.

   From Definition 8 we have:

   1. $H' \subseteq H$;
   2. for every $\phi \in \Phi$ where $\phi \notin Cn^h(\top)$, $H' \subseteq m$ for some $m \in |\neg\phi|$; and
   3. for every H'' where $H' \subset H'' \subseteq H$, we have $H'' \not\subseteq m$ for some $\phi \in \Phi$ where $m \in |\neg\phi|$.

   From 1, 2 we obtain that for each i, $1 \leq i \leq n$, that there is $m_i \in |\neg\phi_i|$ such that:

   $$H' \;\subseteq\; H \cap \bigcap_{i=1}^n m_i \;=\; \bigcap_{i=1}^n (H \cap m_i)$$

   For each i, $1 \leq i \leq n$, we have $H \cap m_i \in H \Downarrow_e \phi_i$ by Definition 3.

   Assume toward a contradiction that $H' \subset \bigcap_{i=1}^n (H \cap m_i)$ and let $\psi \notin H'$ but $\psi \in \bigcap_{i=1}^n (H \cap m_i)$.

   But then $\bigcap_{i=1}^n (H \cap m_i) \subseteq m_i$ for $1 \leq i \leq n$. Thus setting $H'' = \bigcap_{i=1}^n (H \cap m_i)$ contradicts the third condition of Definition 8. This in turn contradicts the fact that $H' \in H \Downarrow_p \Phi$.

   We conclude that $H' \not\subset \bigcap_{i=1}^n (H \cap m_i)$; thus $H' = \bigcap_{i=1}^n (H \cap m_i)$. ■

Proof Sketch of Theorem 5:
The proof for the classical case can be found in (Hansson 1999). For the Horn case, we just need Observations 1 and 2. ■

Proof Sketch of Theorem 6:
The proof is almost identical to the proof of Theorem 5.2.8 in (Wassermann 2000), given Observations 1 and 2. ■

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.

Booth, R.; Meyer, T.; and Varzinczak, I. J. 2009. Next steps in propositional Horn contraction. In *Proceedings of the International Joint Conference on Artificial Intelligence*.

Delgrande, J. 2008. Horn clause belief change: Contraction functions. In Brewka, G., and Lang, J., eds., *Proceedings of the Eleventh International Conference on the Principles of Knowledge Representation and Reasoning*, 156–165. Sydney, Australia: AAAI Press.

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

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.

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

Hansson, S. O., and Wassermann, R. 2002. Local change. *Studia Logica* 70(1):49–76.

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

Khardon, R. 1995. Translating between Horn representations and their characteristic models. *Journal of Artificial Intelligence Research* 3:349–372.

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*.

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

Makinson, D. 1987. On the status of the postulate of recovery in the logic of theory change. *Journal of Philosophical Logic* 16:383–394.

Selman, B., and Kautz, H. 1996. Knowledge compilation and theory approximation. *Journal of the ACM* 43(2):193–224.

Wassermann, R. 2000. *Resource-Bounded Belief Revision*. Ph.D. Dissertation, Institute for Logic, Language and Computation, University of Amsterdam.