Full Acceptance Through Argumentation - a Preliminary Report

Renata Wassermann
Institute for Logic, Language and Computation
University of Amsterdam
Plantage Muidergracht 24
1018TV - Amsterdam, The Netherlands
e-mail: renata@wins.uva.nl

Abstract

When an agent receives new pieces of information, these may contradict his previous beliefs. The agent must decide how to solve this contradiction. Most frameworks dealing with the problem of belief revision attach higher priority to incoming information, i.e., they may give up some part of the old beliefs in order to accommodate the new piece of information and keep consistency. In this paper, we propose the use of argumentation theory to decide whether incoming information should be accepted or not.

1 Introduction

The problem of belief revision, i.e., of how the beliefs of an agent should change in the presence of new information, has been recently addressed by various authors. In most approaches, specially those following the AGM paradigm [Alchourrón *et al.*, 1985], the agents are idealized in that they are assumed to have perfect recall and to hold only consistent beliefs, which are furthermore assumed to be closed under logic consequence.

Incoming information is usually given the highest priority, so that if a contradiction arises, some of the previous beliefs have to be given up. In approaches to non-prioritized belief revision [Hansson, 1997], i.e., revision in which the new piece of information does not have the highest priority, the decision whether to accept or not new information is taken by extra-logical means such as selection functions or incision functions, but there is no real recipe of how to choose these functions. In this paper, we explore a different idea - using argumentation theory for deciding whether new information is acceptable.

In [Wassermann, 1999b] we have developed a framework for belief revision which takes into account the effects of both limited memory and limited capacities of inference. In this model, the belief state of an agent is represented by a structure that distinguishes different kinds of beliefs: beliefs that are explicit or basic, beliefs that are implicit or merely derived, and beliefs that are active, i.e., in use. Besides the beliefs of the agent, the structure also represented "provisional beliefs", i.e., sentences for which the agent has some evidence but is not yet sure whether to accept them or not. In [Hansson and Wassermann, 1999] and [Wassermann, 1999a], some ways of deciding which beliefs were active during a certain operation of belief change were explored. In this paper, we turn to another question left open by the model in [Wassermann, 1999b], namely how to decide whether a provisional belief should be accepted.

According to [Dung, 1995], a formula is believable "if it can be argued successfully against attacking arguments". Dung also says that reasoning abount one's own beliefs is like performing an internal argument. Our concept of provisional beliefs is based on Harman's idea of *tentative hypotheses*. In order to be fully accepted, a tentative hypothesis has to survive the best attempts to refute it [Harman, 1986]. In our case, "best attempts" are as good as the agent is capable given his limitations.

This is reflected in the framework for resource-bounded argumentation given in [Loui, 1998]. Loui describes a very general framework where there are a number of parties involved, some of which (the players) are allowed to make locutions, the others being advocates. Each of the players try to get the current opinion to be in his favour by presenting arguments. A vector represents the resources consumed at each move.

A protocol for disputation has to be defined and depends on the application. These are the real "rules of the game", which determine what is allowed as a move, who is allowed to take next move, how the moves affect the current opinion, and what the conditions for termination are. In [Loui, 1998], some protocols are presented, which can be chosen according to the intended application.

In the next section we will present the framework for resource-bounded belief revision introduced in [Wassermann, 1999b]. Then we will present the theory of argumentation that we use, based on [Loui, 1998]. In section 4 we present our proposal for using the theory sketched in section 3 to enrich the framework presented in section 2.

2 Belief States and Change Operations

In this section, we are going to briefly present the framework for resource-bounded belief revision introduced in

[Wassermann, 1999b]. We start by introducing some distinctions between different kinds of beliefs. The example below motivates the distinctions.

Consider the following situation: Mary is going out, and her mother tells her that she should take an umbrella. Besides beliefs about other subjects, Mary holds the belief that if she is going to be outside for a long time, then she should take the umbrella. She also believes that she will be outside the whole day. If her mother had not mentioned the umbrella, Mary would not have thought of it. Upon it being brought to her notice, she concludes she should indeed take the umbrella.

Following Harman [Harman, 1986], we will assume that there are beliefs that are explicitly represented. We call *implicit beliefs* those beliefs that can be inferred from the set of the agent's explicit beliefs, according to the agent's logical ability. We will not concentrate in one particular inference operator, but use *Inf* to denote what an agent can infer in one step. The set of implicit beliefs is given by what the agent would be able to infer if he was given unlimited time, i.e., the result of applying *Inf* an unlimited number of times.

Not all of the agents beliefs are available at the same time. We call *active beliefs* the set containing beliefs that are available for use and things about which the agent is not yet sure. These last are called *provisional beliefs*. Provisional beliefs are not real beliefs, since they are still under inspection. They are outside the set of explicit beliefs.

A belief state is a structure $\beta = \langle E, Inf, A\rangle$, where E is the set of the agent's explicit beliefs, *Inf* is the agents inference function and A is the set of the agent's active beliefs. The set of implicit beliefs is given by: $I = Inf^{*}(E) = Inf(E) \cup Inf^{2}(E) \cup Inf^{3}(E) \cup \ldots$..

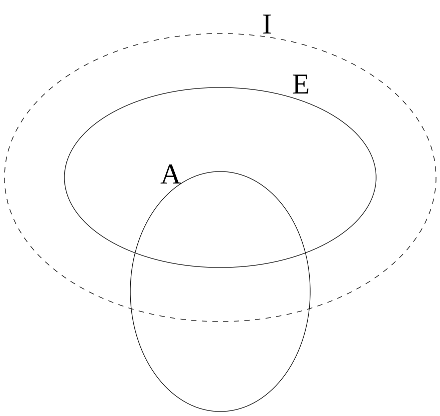

Figure 1: Structure of an agent's beliefs

At this point it may be useful to return to our small example to illustrate the difference between explicit and active beliefs. Mary's belief that if she is going to be outside for a long time, then she should take the umbrella is part of her explicit beliefs and so is her belief that she will be outside the whole day. These beliefs only become active when her mother mentions the umbrella. When Mary thinks of it, she infers that she should take the umbrella. This example shows an argument against representing belief states as logically closed sets. Mary did not hold the belief that she should take the umbrella until the time at which the inference was made. It also shows that not all beliefs are active at the same time. Let p stand for "Mary should take an umbrella" and q for "Mary will be outside for a long time". Before talking to her mother, Mary's explicit beliefs contain, among others, the beliefs q and $q \rightarrow p$. The implicit beliefs contain, among others, p. The set of active beliefs is empty (actually it could probably contain some remains of other reasoning, but this is not relevant for this argument). When the mother says that Mary should take an umbrella, p becomes a provisional active, but not explicit, belief. Mary does not necessarily believe everything her mother says immediately, so that she has to think about it. This is as if she were asking herself whether she should take the umbrella. The beliefs q and $q \rightarrow p$ become active, since they are relevant for deciding whether to accept p. When Mary eventually decides to accept p, this belief is made explicit and the set of active beliefs may get new elements according to new input.

Given our representation of belief state, the next step is to define operations that can be applied to belief states to modify them.

In AGM theory, three operations are defined on belief states: expansion, contraction, and revision. Expansion consists in adding a new belief to the belief state without checking the consistency of the resulting state, contraction consists in deleting a belief from a belief state in a way that the resulting state does not imply the deleted belief, and revision consists in adding a belief to a belief state in such a way that the resulting belief state is consistent. Traditionally, revision is seen as a sequence of a contraction and an expansion (in any order). But this is not a division into simpler steps, since contraction is (computationally) as complicated as revision. We want to decompose revision and contraction in simple operations that show what happens with an agent's belief state in each step, instead of only analyzing the initial and final states.

The set of active beliefs is based on the concept of a short-term memory. Beliefs that are active can be forgotten or stored as explicit (but inactive) beliefs. Since the set of active beliefs is assumed to be very limited in size, there must be a mechanism that, in cases of overflow, selects which beliefs will be forgotten or stored.

The first operation we define is similar to AGM expansion in the sense that it consists in simply adding new information to a set without checking for consistency. But the operation takes the limited size of the set into account.[1] When trying to add something to a set that is already at its maximum size, some elements of the set have to be given up. This can be seen as a kind of "forgetting".

If X is a set with maximum size m and α is an element we want to add to X, then:

$$X \cup^* \{\alpha\} = X' \cup \{\alpha\}, \text{ where } X' \subseteq X, |X'| < m.$$

Note that this operation reduces to a simple union as long as the set is not "full". Since the size m of the set is given as a parameter, the operation is more accurately denoted as $\cup^*_m$. When the set is already at its maximum size, something has to be discarded. If the set X is ordered (for example by the last time the beliefs were recalled), we can stipulate that the minimal elements of the set are the first to be dismissed, i.e., we want to ensure that if an element is dismissed, then there is no other element which is retained and that is less than the dismissed one in the order:

$$\forall y (y \in X \setminus X' \to \neg \exists x (x \in X' \wedge x < y)).$$

We now define six operations that can be applied on belief states to change the status of beliefs.

Definition 2.1 *Let $\langle E, Inf, A\rangle$ be a belief state and α a formula. We define the following operations on $\langle E, Inf, A\rangle$ (we will omit the second argument Inf since the operations defined do not affect it):*

1. *Observation ($+_o$): adds an external input to the set of active beliefs.*
   $\langle E, A\rangle +_o \alpha = \langle E, A \cup^* \{\alpha\}\rangle$
2. *Retrieval ($+_r$): retrieves an explicit belief into the set of active beliefs.*
   $$\langle E, A\rangle +_r \alpha = \begin{cases} \langle E, A \cup^* \{\alpha\}\rangle, & \text{if } \alpha \in E \\ \langle E, A\rangle & \text{otherwise} \end{cases}$$
3. *Acceptance ($+_a$): makes an active belief explicit.*[2]
   $$\langle E, A\rangle +_a \alpha = \begin{cases} \langle E \cup^* \{\alpha\}, A \setminus \{\alpha\}\rangle, & \text{if } \alpha \in A \\ \langle E, A\rangle & \text{otherwise} \end{cases}$$
4. *Inference ($+_i$): infers something from active beliefs.*
   $$\langle E, A\rangle +_i \alpha = \begin{cases} \langle E, A \cup^* \{\alpha\}\rangle, & \text{if } \alpha \in Inf(A) \\ \langle E, A\rangle & \text{otherwise} \end{cases}$$
5. *Doubting ($+_d$): a belief that was accepted is questioned, becoming provisional.*
   $$\langle E, A\rangle +_d \alpha = \begin{cases} \langle E \setminus \{\alpha\}, A\rangle, & \text{if } \alpha \in A \cap E \\ \langle E, A\rangle & \text{otherwise} \end{cases}$$
6. *Rejection ($+_c$): rejects an active belief.*
   $$\langle E, A\rangle +_c \alpha = \begin{cases} \langle E, A \setminus \{\alpha\}\rangle, & \text{if } \alpha \in A \\ \langle E, A\rangle & \text{otherwise} \end{cases}$$

[1] When we talk about the size of a set of formulas, we mean something like its complexity. The sets $\{p, q\}$ and $\{p \wedge q\}$ should have the same size. We could, for example, count the occurrence of atoms.

[2] Acceptance could also be defined without deleting the accepted belief from A, which seems to be more intuitive for human agents. The choice made here reflects our interest in artificial agents.

The six operations defined above can be combined to model more complex operations. As an example of such a composition, consider what happens when an agent gets new information via observation. The belief will first come into the set of active beliefs through the operation $+_o$ and then the agent may accept it ($+_a$). Another example is the case of an explicit belief that becomes active (retrieval: $+_r$), when it would be expected that some implicit beliefs will also become active, i.e., the retrieval operation will be followed by an inference ($+_i$).

3 Argumentation

In this section, we introduce the basic concepts of argumentation theory that we will need in this paper. This section is based on [Loui, 1998].

Argumentation has been investigated by researchers in the area of philosophy and artificial intelligence. Recently, it became clear that argumentation can be seen as a kind of non-monotonic reasoning. Arguments are not proof, but some kind of justification for a claim, usually defeasible. An argumentation process usually follows some protocol that defines what the possible moves are, how a move affects the current state of the disputation, who is allowed to move, etc. Once the parties involved in the disputation agree about the protocol, the outcome of an argumentation process following the protocol is considered fair.

Disputations are highly non-monotonic. The outcome depends on the particular way in which the argumentation process took place and if the process continues, the outcome may change. Nevertheless, the process is fair (provided the disputants agreed about the protocol) and the outcome is warranted.

An argument is usually a pair formed by a set of formulas and one special formula, the claim. The set of formulas serves as a justification for the claim. Arguments are related to each other is several ways. Arguments can interfere with each other, in case their claims (or subclaims) are inconsistent.

Loui [Loui, 1998] defines a very general framework for argumentation that has to be "filled in" in order to model particular kinds of disputation.

An argumentation process is a sequence of locutions, where each locution is a triple formed by one party, the argument and the resources consumed. The participants of the argumentation process do not have necessarily access to the same information. They may also have different shares of resources at their disposal. In our case, we will use argumentation processes where only two parties are involved, **pro** and **con**. A variable *current.opinion* stores the party which is winning the disputation at a certain point. The parties try to switch the current opinion in their favour by advancing locutions. Since we are modeling an internal argumentation process, where a single agent is involved and plays the roles of **pro** and **con**, we can assume that both parties have access to the same information.

4 Using Argumentation for Accepting Beliefs

In this section we present our proposal for using argumentation in order to decide whether a provisional belief should be accepted or not. In our case, the argumentation is an internal process where a single agent plays the role of **pro** and **con**, analyzing the arguments for and against a given provisional belief. Since we are dealing with resource-bounded agents, this internal argumentation will not always succeed in examining all reasons for accepting or rejecting a belief. By defining a protocol for this process, we have to take care that the outcome can be considered fair.

There are two ways in which a sentence can become a provisional belief:

1. New information may be acquired by an operation of observation, i.e., come from the outside world. This new piece of information has to be checked before being fully accepted. In this case, **con** tries to argue against it. If he fails, the provisional belief is accepted, since it has survived the best attempts to refute it. If **con** succeeds, the provisional belief is rejected.
2. A sentence that was previously accepted, an explicit belief, may become provisional if the agent gets evidence against it. In this case, inquiry is reopened ([Harman, 1986]) and **pro** tries to argue for the sentence. If he fails, the provisional belief is rejected. If **pro** succeeds, the provisional belief becomes fully accepted again.

In the framework presented in section 2, there are two clearly limited resources: the size of the set of active beliefs and the number of basic operations used in the disputation process. Since in our case a single agent is playing the roles of **pro** and **con**, the set of active beliefs is a shared resource, both **pro** and **con** have access to the whole set.

All the sentences in the arguments presented become active. The elements of the set of active beliefs are ordered according to the order in which they were introduced in the argumentation. When the set gets too big, the oldest elements are "forgotten". If the discarded elements were explicit beliefs that were retrieved, they remain in the set of explicit beliefs but become inactive. If they were only provisional beliefs, then they are irremediably forgotten and dismissed from the whole structure.

An argument for us will be a sequence of elements of the set of explicit beliefs which is a derivation for its claim according to a finite (small) number of applications of inference rules known by the agent. An argument *arg* of player p (=**pro** or **con**) is counterargued when the other player presents an argument against one of the elements of *arg* (its subclaims). An argument *arg* of player p is defeated if it is counterargued by *arg'* and p does not manage to counterargue *arg'* (either because there are no counterarguments or because the resources are exhausted).

When an argument is introduced by one of the players, the beliefs that are part of it are retrieved into the set of active beliefs. When an argument is counterargued, its claim becomes provisional. If an argument is defeated, its claim is rejected.

The protocol we will be using assumes that the resources are equally divided, i.e., if player p_1 has exhausted his share of resources but p_2 has not, then p_2 is still allowed one move. Except for this situation, the players alternate the moves. No repetition of counterargued (sub-)arguments is allowed.

Suppose a sentence α is observed. The current opinion is set to **pro** and **con** tries to find an argument for $\neg\alpha$. If he fails, then α is accepted, otherwise, current opinion is set to **con** and **pro** tries to either counterargue the last argument or present a new argument for α. If **pro** fails, then α is rejected. Otherwise, current opinion is set to **pro** and **con** tries to either counterargue the last argument or present a new argument for $\neg\alpha$. The process continues until resources are exhausted. The player favoured by current opinion wins.

5 Example

We will now see an example of application of the protocol described in section 4.

We first have to give some more details about the procedure. The claim to be verified, a provisional belief, remains active during the whole argumentation. It cannot be dismissed due to overflow in the set of active beliefs. The set of active beliefs is ordered by recency, i.e., beliefs that have been used first are the first to be forgotten in case of overflow. However, if an active belief is reused, it becomes more recent and changes place in the order. This agrees with cognitive models of memory, as for example in [Anderson, 1980].

The claim which is being verified and claims of arguments that have been counterargued cannot be used in new arguments.

The size of the set of active beliefs, one of the limited resources, is given by the number of atoms occurring in its formulas. Part of the history of the process is kept in the form of arguments advanced, so that there is no repetition. This set can also be limited in size like the set of active beliefs, but in the example we will ignore this fact.

We will use the following logic for the example:

1. atoms $a, b, c, ..., p$ standing for "*a*lbert comes to the party", "*b*etty comes to the party", "*c*harles comes to the party", ..., "*p*atrick comes to the party".
2. formulas $x \rightarrow y$ standing for "If x comes to the party, then y comes to the party" ; $x \rightarrow \neg y$ standing for "If x comes to the party, then y does not come to the party", etc.
3. inference rules modus ponens ($x, x \rightarrow y \Rightarrow y$) and inversion ($x \rightarrow y \Rightarrow \neg y \rightarrow \neg x$).

Depending on who likes whom and who dislikes whom, we know who is (or is not) going to come to the party given who is (or is not) coming. Moreover, we know of some people that are coming (*a*lbert, *f*erry, *h*arold, *k*ate, and *o*swald). Our initial set of explicit beliefs is:

$$E = \{a, a \to b, b \to c, c \to d, d \to g, f, f \to e, e \to \neg c, \neg c \to \neg p, h, h \to i, i \to j, j \to \neg e, k, k \to l, l \to m, m \to n, n \to \neg i, o, o \to p, p \to \neg l, \neg l \to \neg b\}.$$

We assume that the maximum size of the set of active beliefs is 20. We want to know whether *i*van is coming to the party:

- step 1: **con** tries to refute i, presenting an argument for $\neg i$. The formulas in the argument are retrieved from the set of explicit beliefs and stored as active beliefs. Inference is applied four times in order to get to the claim $\neg i$ from the argument.
  - **con** presents argument $\{k, k \to l, l \to m, m \to n, n \to \neg i\}$ for $\neg i$.
  - 9 basic operations: retrieval $\{k, k \to l, l \to m, m \to n, n \to \neg i\}$; inference $\{l, m, n, \neg i\}$
  - $A = \{i, k, k \to l, l, l \to m, m, m \to n, n, n \to \neg i, \neg i\}$; $|A|$=14
  - History: $\{\{k, k \to l, l \to m, m \to n, n \to \neg i\}\}$
  - current.opinion = **con**
- step 2: **pro** advances an argument against one of the subclaims of the previous argument. The previous argument is counterargued, but not yet defeated, since **con** may counterargue this present argument. The set of active beliefs grows to its maximum size, 20. The oldest active belief besides the claim (k) is dismissed to make space for the new activated beliefs.
  - **pro** presents counterargument $\{o, o \to p, p \to \neg l\}$ against l.
  - 5 basic operations: retrieval $\{o, o \to p, p \to \neg l\}$; inference $\{p, \neg l\}$.
  - $A = \{i, k \to l, l, l \to m, m, m \to n, n, n \to \neg i, \neg i, o, o \to p, p, p \to \neg l, \neg l\}$; $|A|$=20
  - History: $\{\{k, k \to l, l \to m, m \to n, n \to \neg i\}, \{o, o \to p, p \to \neg l\}\}$
  - current.opinion = **pro**
- step 3: **con** counterargues the previous argument. The oldest elements of the set of active beliefs (except i) are dismissed to make space for the new beliefs retrieved.
  - **con** presents counterargument $\{f, f \to e, e \to \neg c, \neg c \to \neg p\}$ against p.
  - 7 basic operations: retrieval $\{f, f \to e, e \to \neg c, \neg c \to \neg p\}$; inference $\{e, \neg c, \neg p\}$.
  - $A = \{i, \neg i, o, o \to p, p, p \to \neg l, \neg l, f, f \to e, e, e \to \neg c, \neg c, \neg c \to \neg p, \neg p\}\}$; $|A|$=19
  - History: $\{\{k, k \to l, l \to m, m \to n, n \to \neg i\}, \{o, o \to p, p \to \neg l\}, \{f, f \to e, e \to \neg c, \neg c \to \neg p\}\}$
  - current.opinion = **con**
- step 4: **pro** counterargues the previous argument. Again, some elements of the set of active beliefs must be dismissed.
  - **pro** presents counterargument $\{a, a \to b, b \to c\}$ against $\neg c$.
  - 5 basic operations: retrieval $\{a, a \to b, b \to c\}$; inference $\{b, c\}$
  - $A = \{i, \neg l, f, f \to e, e, e \to \neg c, \neg c, \neg c \to \neg p, \neg p, a, a \to b, b, b \to c, c\}$; $|A|$=19
  - History: $\{\{k, k \to l, l \to m, m \to n, n \to \neg i\}, \{o, o \to p, p \to \neg l\}, \{f, f \to e, e \to \neg c, \neg c \to \neg p\}, \{a, a \to b, b \to c\}\}$
  - current.opinion = **pro**
- step 5: **con**'s arguments were defeated, since he cannot counterargue the previous arguments advanced by **pro** anymore. **con** advances a new argument against i. Some of the beliefs used in this argument ($f, f \to e, e \to \neg j$) are already active so they do not need to be retrieved. They only change place in the set of active beliefs.
  - **con** presents counterargument $\{f, f \to e, e \to \neg j, \neg j \to \neg i\}$ against i.
  - 6 basic operations: retrieval $\{j \to \neg e, i \to j\}$; inferences $\{e \to \neg j, \neg j, \neg j \to \neg i, \neg i\}$
  - $A = \{i, b, b \to c, c, f, f \to e, e, j \to \neg e, e \to \neg j, \neg j, i \to j, \neg j \to \neg i, \neg i\}$; $|A|$=19
  - History: $\{\{k, k \to l, l \to m, m \to n, n \to \neg i\}, \{o, o \to p, p \to \neg l\}, \{f, f \to e, e \to \neg c, \neg c \to \neg p\}, \{a, a \to b, b \to c\}, \{f, f \to e, e \to \neg j, \neg j \to \neg i\}\}$
  - current.opinion = **con**
- step 6: **pro** cannot counterargue the previous argument, but presents instead a new argument for i. Since **con** does not have any other counterarguments or arguments for $\neg i$, **pro** wins the disputation.
  - **pro** presents argument $\{h, h \to i\}$ for i.
  - 3 basic operations: retrieval $\{h, h \to i\}$; inference $\{i\}$.
  - $A = \{i, c, f, f \to e, e, j \to \neg e, e \to \neg j, \neg j, i \to j, \neg j \to \neg i, \neg i, h, h \to i\}$; $|A|$=19
  - History: $\{\{k, k \to l, l \to m, m \to n, n \to \neg i\}, \{o, o \to p, p \to \neg l\}, \{f, f \to e, e \to \neg c, \neg c \to \neg p\}, \{a, a \to b, b \to c\}, \{f, f \to e, e \to \neg j, \neg j \to \neg i\}, \{h, h \to i\}\}$
  - current.opinion = **pro**

6 Conclusions

We have presented some ideas on how to use argumentation theory in order to decide which beliefs should be fully accepted. These ideas enrich the framework presented in [Wassermann, 1999b].

Although the protocol defined and the example are quite simple, they illustrate the internal process of "weighting" the arguments in favour and against a certain claim that takes place when an agent is confronted with information about which he is not sure.

Further work includes examining existing implemented argumentation systems in order to refine the protocol of the argumentation process. One such system is presented in [Simari and Loui, 1992], together with a mathematical treatment of the relations between arguments.

Acknowledgments: I would like to thank Daniela Carbogim and Frans Voorbraak for comments on an earlier draft. This work is supported by a grant from the Brazilian funding agency CAPES.

References

[Alchourrón *et al.*, 1985] Carlos Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change. *Journal of Symbolic Logic*, 50(2):510–530, 1985.

[Anderson, 1980] John R. Anderson. *Cognitive Psychology and its Implications*. W.H.Freeman, 1980.

[Dung, 1995] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-persons games. *Artificial Intelligence*, 77:321–357, 1995.

[Hansson and Wassermann, 1999] Sven Ove Hansson and Renata Wassermann. Local change. In preparation (a preliminary version appeared in the Fourth Symposium on Logical Formalizations of Commonsense Reasoning, London, 1998), 1999.

[Hansson, 1997] Sven Ove Hansson, editor. *Theoria - special issue on non-prioritized belief revision*, volume 63, 1997.

[Harman, 1986] Gilbert Harman. *Change in View - Principles of Reasoning*. MIT Press, 1986.

[Loui, 1998] Ronald P. Loui. Process and police: Resource-bounded non-demonstrative reasoning. *Computational Intelligence*, 1998.

[Simari and Loui, 1992] Guillermo R. Simari and Ronald P. Loui. A mathematical treatment of defeasible reasoning and its implementation. *Artificial Intelligence*, 53(2-3):125–157, 1992.

[Wassermann, 1999a] Renata Wassermann. On structured belief bases. In Hans Rott and Mary-Anne Williams, editors, *Frontiers in Belief Revision*. Kluwer, 1999. to appear.

[Wassermann, 1999b] Renata Wassermann. Resource-bounded belief revision. *Erkenntnis*, 1999. to appear.