

Strong and Explicit Negation in Non-Monotonic Reasoning and Logic Programming

José Júlio Alferes^{1*}, Luís Moniz Pereira^{2*}, and Teodor C. Przymusiński^{3**}

¹ DM, U. Évora and CITIA, U. Nova de Lisboa 2825 Monte da Caparica, Portugal

² DCS and CITIA, U.Nova de Lisboa 2825 Monte da Caparica, Portugal

³ Dep. of Computer Science University of California Riverside, CA 92521, USA

1 Introduction

Logic programs, deductive databases, and more generally non-monotonic theories, use various forms of *default negation*, *not F*, whose major distinctive feature is that *not F* is assumed “by default”, i.e., it is assumed in the absence of sufficient evidence to the contrary. The meaning of “*sufficient evidence*” depends on the specific semantics used. For example, in Reiter’s *Closed World Assumption*, *CWA* [32], *not A* is assumed for atomic *A* if *A* is not provable, or, equivalently, if there is a minimal model in which *A* is false. On the other hand, in Minker’s *Generalized Closed World Assumption*, *GCWA* [18, 16], or in McCarthy’s *Circumscription*, *CIRC* [17], *not A* is assumed only if *A* is false in *all* minimal models. In Clark’s *Predicate Completion* semantics for logic programs [7], this form of negation is called *negation-as-failure* because *not A* is derivable whenever attempts to prove *A* finitely fail.

The more recent semantics proposed for logic programs and deductive databases, such as the *stable semantics* [14], *well-founded semantics* [34], *partial stable* or *stationary semantics* [27] and *static semantics* [30], propose even more sophisticated meanings for default negation, closely related to more general non-monotonic formalisms such as *Default Logic*, *DL* [33], *AutoEpistemic Logic*, *AEL* [19], and *AutoEpistemic logic of Beliefs*, *AEB* [29]. Under all of these semantics, however, default negation “*not*” significantly differs from classical negation “ \neg ”. For example, the formula $charged(x) \wedge \neg guilty(x) \supset acquitted(x)$ says that a person charged with a crime should be acquitted if he or she is actually proven to be not guilty. On the other hand, the formula $charged(x) \wedge not\ guilty(x) \supset acquitted(x)$ says that one should be acquitted of any charges by default unless sufficient evidence of that person’s guilt is demonstrated.

1.1 Default negation does not suffice

Although default negation proved to be quite useful in various domains and application frameworks, there are at least two important reasons why it is not the only form of negation that is needed (in addition to classical negation) in non-monotonic formalisms:

* Partially supported by JNICT-Portugal

** Partially supported by the National Science Foundation grant #IRI-9313061.

- While default negation *not A* of an atomic formula *A* is always assumed “by default”, we often need to be more careful before jumping to negative conclusions. For example, it would make little sense to say $\text{not innocent}(x) \supset \text{guilty}(x)$ to express the fact that being guilty is the opposite of being innocent because it would imply that people should be considered guilty “by default”. Similarly, we would not want to say: $\text{crossing}(x), \text{not train}(x) \supset \text{drive}(x)$ to express the fact that one may safely cross a railway track if (s)he first makes sure that there is no train approaching, because such a clause would imply that (s)he should cross the railway even if (s)he did not bother to look around at all.
- Although negation *not A* is always assumed “by default” for atomic formulae *A*, the same is not true for negated atoms $F = \neg A$ and thus default negation does not treat literals *A* and $\neg A$ symmetrically. For example, even though GCWA (or circumscription) applied to the theory $\{\text{charged}(\text{tom}); \text{charged}(x) \wedge \text{guilty}(x) \supset \text{convicted}(x)\}$ implies $\text{not convicted}(\text{tom})$ this is no longer the case if we perform a simple syntactic substitution of $\neg \text{innocent}$ for *guilty* thus obtaining $\{\text{charged}(\text{tom}); \text{charged}(x) \wedge \neg \text{innocent}(x) \supset \text{convicted}(x)\}$ because the new theory has minimal models in which $\text{convicted}(\text{tom})$ is true. As we can see, a simple syntactic substitution of $\neg \text{innocent}$ for *guilty* turns out to have a significant impact on the semantics of the resulting theory. As a consequence, default negation, *not F*, is heavily dependent on the syntactic form of the formula *F*, and, in particular, it is *not invariant* under the *equivalence* or *renaming* of predicates.

1.2 Classical negation does not solve the problem

One can argue that the above features of default negation were intentionally chosen to distinguish default negation *not F* from classical negation $\neg F$ and therefore if *not F* does not work in some particular context then $\neg F$ should be used instead⁴. While the first part of this statement is certainly correct the second is not:

- Classical negation $\neg F$ satisfies the so called “*law of the excluded middle*”, $F \vee \neg F$, for every formula *F*. While it is a natural axiom in the domain of formal logic and mathematics, it is not suitable for *common-sense reasoning* where we are often faced with situations in which some things are true, some others are false and yet some others are not (perhaps, yet) determined to be either true or false. For example, an employer may be faced with some candidates who are clearly qualified for the job, and thus should be hired, some others who are clearly not qualified, and thus should be rejected, and yet some other candidates whose qualifications are

⁴ In fact, real classical negation $\neg F$ is not part of the language of standard logic programs and deductive databases and thus it is not immediately available without an appropriate extension of their languages and semantics, such as the one proposed in [30].

not clear and who therefore should be interviewed in order to determine their suitability for the job (see [13]). However, if we attempt to describe the first two statements using real classical negation “ \neg ” in the clauses: $\{qualified(x) \supset hire(x); \neg qualified(x) \supset reject(x)\}$ then, by virtue of the law of the excluded middle, we will immediately conclude that every candidate must either be hired or rejected without leaving any room for those who need to be further interviewed.

The either/or character of the law of the excluded middle presupposes that we have sufficient information to determine, at least in principle, whether any given property or its opposite is true. However, when dealing with *incomplete information* or with properties involving *gradual change* we often encounter “gray areas” in which neither the property nor its opposite seem to hold. Moreover, the law of the excluded middle leads to highly unintuitive results when used with reference to non-existing or non-applicable properties or objects, e.g., in statements like “the the pink elephant in my pocket is either old or not old” or “The chair is either happy or unhappy”.

Moreover, the law of the excluded middle is highly *non-constructive* and thus does not fit well within the somewhat vague “spirit” of logic programming and deductive databases which seems to rely heavily on “*directional reasoning*”, i.e., on first establishing the validity of premises of a clause before deducing its consequents and not the other way around. As a result of its non-constructive character, the law of the excluded middle is also computationally expensive and thus difficult to implement.

- As we pointed out before, classical literals A and $\neg A$ are not treated symmetrically by default negation. While, in the absence of evidence to the contrary, default negation *not* A of *positive* literals (atoms) A is always assumed, the same does not apply to negative literals $\neg A$. Often times, however, we would like to apply negation by default *not* symmetrically to both positive and negative literals. For instance, in the preceding example, given no information about *Tom*’s qualifications we would like to conclude that Tom was neither found qualified nor unqualified, i.e., that both *not* (*qualified*(*Tom*)) and *not* (\neg *qualified*(*Tom*)) hold. This would tell us that Tom has to be interviewed. This is not possible with classical negation “ \neg ” because the law of the excluded middle forces A to be true if $\neg A$ is not true and vice versa.

1.3 Symmetric negation

Based on the preceding discussion, we conclude that, in addition to the classical negation, “ \neg ”, and the default negation, “*not*”, we need a new negation, which we will denote by “ \sim ”, which has the following properties:

- As opposed to default negation, the negation $\sim A$ is not assumed by default;
- As opposed to classical negation, the negation “ \sim ” does not satisfy the law of the excluded middle, $F \vee \sim F$;
- The negation “ \sim ” is treated *symmetrically* by the default negation “*not*”, i.e., in the absence of evidence to the contrary, both *not* (A) and *not* ($\sim A$)

are assumed. More generally, the semantics of belief theories is invariant under the renaming of any predicates from A to $\sim A$ and vice versa.

We will call negation “ \sim ” with the above properties *symmetric negation*.

Gelfond and Lifschitz [15] were the first to point out the need for such negation in logic programming and they also proposed a specific semantics for such negation for logic programs with the stable semantics. Somewhat unfortunately, they called their negation “*classical negation*”⁵. Subsequently, several researchers proposed different, often incompatible, forms of symmetric negation for various semantics of logic programs and deductive databases [11, 21, 23, 28, 30, 35]. To the best of our knowledge, however, no systematic study of symmetric negation in non-monotonic reasoning was ever attempted in the past⁶. In this paper we conduct such a systematic study of symmetric negation:

- We introduce and discuss two natural, yet different, definitions of symmetric negation: one is called *strong negation* and the other is called *explicit negation*. For logic programs with the stable semantics, both symmetric negations coincide with Gelfond-Lifschitz’ “classical negation”.
- We study properties of strong and explicit negation and their mutual relationship as well as their relationship to default negation “*not*” and classical negation “ \neg ”.
- Rather than to limit our discussion to some narrow class of non-monotonic theories, such as the class of logic programs with some specific semantics, we conduct our study so that it is applicable to a broad class of non-monotonic formalisms, which includes the well-known formalisms of circumscription, autoepistemic logic and all the major semantics recently proposed for logic programs (stable, well-founded, stationary, static and others).
- In order to achieve the desired level of generality, we define the notion of symmetric negation in the knowledge representation framework of *AutoEpistemic logic of Beliefs*, *AEB*, introduced by Przymusiński in [29], which was shown to isomorphically contain all of the above mentioned formalisms as special cases. As a result, we automatically provide the corresponding notions of symmetric negation for all formalisms embeddable into *AEB*.

The paper is organized as follows. In Section 2 we introduce strong negation, the first form of symmetric negation, and discuss its basic properties. Then we introduce explicit negation (the second form of symmetric negation), we establish its basic properties, compare it to strong negation, and in Section 4 we discuss applications of explicit negation to knowledge representation. We conclude with some final remarks. In Appendix we briefly recall the basic definition of the Autoepistemic Logic of Beliefs, *AEB*. A full account can be found in [29, 30, 31].

⁵ Independently, in [22] the authors also pointed out the need for such a negation in logic programs. They called it “strong” negation, due to the desired similarities with Nelson’s strong negation [20]. However their “strong” negation should not be confused with the strong negation defined in this paper.

⁶ For logic programs, an extensive study was carried out in [3].

2 Strong Negation

In the Introduction we concluded that in addition to default negation, *not F*, non-monotonic reasoning requires a new type of negation which is similar to classical negation $\neg F$, in the sense that it is not assumed by default, and yet does not satisfy the law of the excluded middle and allows a symmetric treatment of positive and negative information. In this section we define the so called *strong negation*, $\neg F$, and argue that it is a natural candidate for such a negation. We study properties of strong negation and its relationship to default negation and classical negation.

We introduce strong negation within the broad framework of the Autoepistemic Logic of Beliefs, *AEB*. As a result, the definition of strong negation applies to all non-monotonic formalisms embeddable into *AEB*, including circumscription, autoepistemic logic and all the major semantics recently proposed for logic programs. The definition is patterned after the original definition given in [15] and therefore, when applied to logic programs with stable semantics, strong negation coincides with Gelfond-Lifschitz' so called "classical negation" .

Strong negation is introduced to belief theories T in the Autoepistemic Logic of Beliefs, *AEB*, by:

- (1) augmenting the objective language \mathcal{L} with a set $\widehat{\mathcal{S}} = \{\neg A : A \in \mathcal{S}\}$ of new *objective* propositional symbols, called *strong negation atoms*, where \mathcal{S} is any fixed set of propositional symbols from \mathcal{L} . As a result we obtain an extended objective language $\widehat{\mathcal{L}}$ and an extended language of beliefs $\widehat{\mathcal{L}}_{AEB}$.
- (2) adding to the belief theory T the following *strong negation constraint*:

$$(S_A) \quad A \wedge \neg A \supset \perp \quad \text{or, equivalently,} \quad \neg A \supset \neg A,$$

for every strong negation atom $\neg A$ that occurs in T .

The strong negation constraint S_A states that A and $\neg A$ cannot be both true and thus ensures that the intended meaning of strong negation $\neg A$ is " $\neg A$ is the opposite of A ". For example, a proposition A may describe the property of being "*qualified*" while the proposition $\neg A$ describes the property of being "*unqualified*". The strong negation constraint states that a person cannot be *both* qualified and unqualified. We do not assume, however, that everybody is already known to be either qualified or unqualified.

We want to emphasize that rather than modifying the logic *AEB* itself by including strong negation constraints in the logical closure operator Cn_{AEB} we instead make strong negation constraints part of any belief theory in which they occur. As a result, theories which do not contain strong negation atoms are not affected by these constraints in any way. In the sequel, we will implicitly assume that the strong negation constraint S_A is part of any belief theory T which contains the atom $\neg A$.

For the sake of readability, instead of using somewhat cumbersome names, like $\neg good_teacher$, to denote the strong negation of the predicate *good_teacher*, we use more readable names, such as *bad_teacher*.

Furthermore, even though we only consider propositional languages, we often use variables as a shortcut for all ground instantiations of a given formula.

Example 1. ⁷ Consider a university that periodically evaluates faculty members based on their research and teaching performance. Faculty members who are known to be strong researchers and who are believed to be good teachers (as we all know, it is difficult to objectively evaluate teaching) receive positive evaluation. On the other hand, those who are believed not to be strong researchers as well as those that are believed to be poor teachers receive negative evaluation. Anyone who received a teaching award is considered to be a good teacher and anyone who published at least 10 reviewed papers during the evaluation period is considered a good researcher. The individuals whose evaluation status is not yet determined undergo some further (unspecified) review process. This leads us to the following belief theory in *AEB*:

$$\begin{aligned}
& \text{good_researcher}(x) \wedge \mathcal{B} \text{good_teacher}(x) \supset \text{good_evaluation}(x) \\
& \mathcal{B} \neg \text{good_researcher}(x) \supset \text{bad_evaluation}(x) \\
& \mathcal{B} \text{bad_teacher}(x) \supset \text{bad_evaluation}(x) \\
& \text{teaching_award}(x) \supset \text{good_teacher}(x) \\
& \text{many_publications}(x) \supset \text{good_researcher}(x).
\end{aligned}$$

The predicates *bad_teacher*, *bad_evaluation* and *bad_researcher* are intended to represent strong negation of the predicates *good_teacher*, *good_evaluation* and *good_researcher*, respectively, and, therefore, we need to add to our theory strong negation constraints (S_A) for each one of them:

$$\begin{aligned}
& \text{good_researcher}(x) \wedge \text{bad_researcher}(x) \supset \perp \\
& \text{good_teacher}(x) \wedge \text{bad_teacher}(x) \supset \perp \\
& \text{good_evaluation}(x) \wedge \text{bad_evaluation}(x) \supset \perp.
\end{aligned}$$

Suppose that both Ann and Tom published at least 10 papers. Suppose, further, that Ann is a good teacher, both Mary and Keith received teaching awards but Tom is considered to be a bad teacher. Moreover, Keith and Paul have a lot of joint publications so at least one of them must be a good researcher:

$$\begin{aligned}
& \text{many_publications}(\text{Ann}) & \text{many_publications}(\text{Tom}) & \text{good_teacher}(\text{Ann}) \\
& \text{teaching_award}(\text{Mary}) & \text{teaching_award}(\text{Keith}) & \text{bad_teacher}(\text{Tom}) \\
& \text{good_researcher}(\text{Keith}) \vee \text{good_researcher}(\text{Paul}).
\end{aligned}$$

The resulting belief theory has one consistent static expansion in which Ann receives a positive evaluation because she is a good researcher and is believed to be a good teacher. Tom receives a negative evaluation because even though he is a good researcher, he is also believed to be a bad teacher. Mary receives a negative evaluation as well because even though she is a good teacher, there is no evidence of good research work in her file (i.e., $\mathcal{B} \neg \text{good_researcher}(\text{Mary})$ holds).

Keith's and Paul's status is not yet clear and thus they will have to be further reviewed. Indeed, none of them individually has been shown to be a strong researcher so they are not eligible for positive evaluation. At the same time none of them is believed to be a bad teacher and there is some evidence supporting each one of them being a good researcher (i.e., neither $\mathcal{B} \neg \text{good_researcher}(\text{Paul})$

⁷ The correctness of the examples discussed in this section has been verified by using the *interpreter for static semantics* developed by Stefan Brass based on the results obtained in [6]. The interpreter is available from <ftp://ftp.informatik.uni-hannover.de/software/static/static.html> via FTP and WWW.

nor $\mathcal{B}\neg\text{good_researcher}(\text{Keith})$ is true) and therefore they are not subject to negative evaluation, either. Observe, however, that if we later find out that Paul is in fact a lousy researcher, $\text{bad_researcher}(\text{Paul})$, then Paul will receive a negative evaluation and we will conclude (by virtue of the strong negation constraint) that Keith is a good researcher and thus Keith will receive a positive evaluation. \square

Even though strong negation $\neg F$ is defined so far only for atomic objective formulae F , we can easily extend it to all objective formulae F of the extended language $\widehat{\mathcal{L}}$. The following definition is recursive and F and G are assumed to be objective formulae:

$$\neg(\neg F) \equiv F \quad \neg(\neg\neg F) \equiv \neg F \quad \neg(F \vee G) \equiv \neg F \wedge \neg G \quad \neg(F \wedge G) \equiv \neg F \vee \neg G.$$

The above extension preserves the basic property of strong negation, namely, the fact that it is stronger than classical negation:

Proposition 1 Strong Negation is Stronger than Classical Negation.

Suppose that T is a belief theory and F is an objective formula such that T includes the strong negation constraint S_A , for every atom A that occurs in F . Then: $T \models (\neg F \supset \neg F)$.

From the above Proposition one easily derives an important relationship between strong negation and beliefs:

Proposition 2 Strong Negation vs. Beliefs. *Suppose that T is a belief theory. All static autoepistemic expansions T° of T are closed under the inference rule:*

$$\frac{\neg F}{\mathcal{B}\neg F},$$

where F is any objective formula such that T includes the strong negation constraint S_A , for every atom A that occurs in F .

Observe that for any objective positive formula F , both F and $\neg F$ are assumed false by default, i.e., both $\mathcal{B}\neg F$ and $\mathcal{B}\neg(\neg F)$ are true, because, in the absence of any other information, F and $\neg F$ are false in all minimal models.

However, one can easily ensure that one of the formulae, F or $\neg F$, is *true by default*, and thus only the other one is minimized, by assuming one of the following *default axioms*: $\mathcal{B}\neg F \supset \neg F$ or $\mathcal{B}\neg(\neg F) \supset F$ which say that if we disbelieve F (respectively, $\neg F$) then $\neg F$ (respectively, F) can be assumed to be true. For example, the first default axiom causes $\neg F$ to be true by default thus forcing strong negation to behave in a way that is similar (but not identical) to classical negation. This will prove useful when applying strong negation to theory and interpretation update [5].

One can also *prevent* any minimization of F and $\neg F$ by assuming the *law of the excluded middle*, $F \vee \neg F$, for this particular formula F , which causes precisely one of F or $\neg F$ to be true at all times.

As we mentioned in the Introduction, non-monotonic formalisms based on some form of *predicate minimization*, such as *CWA*, circumscription and most

semantics of logic programs, do not treat atomic formulae A and their classical negations $\neg A$ symmetrically and thus they are *not invariant* under a simple renaming substitution replacing atoms by their negations and vice versa.

The Autoepistemic Logic of Beliefs, a superset of such formalisms, naturally suffers from the same problem. For example, the belief theory $\mathcal{B}\neg A \supset C$ has a unique static expansion in which C is true, and, yet, after substituting A' for $\neg A$, the resulting theory $\mathcal{B}A' \supset C$ has a unique expansion which no longer contains C . However, since strong negation atoms $\neg A$ are introduced as regular objective atoms and since renaming of atoms has no effect on the semantics, we immediately conclude that beliefs in AEB are invariant under renaming of any predicates from A to $\neg A$ and vice versa. The next proposition illustrates the most basic difference between classical and strong negation.

Proposition 3 Strong Negation is Symmetric. *Static semantics of belief theories in AEB is invariant under the renaming of any predicates from A to $\neg A$ and vice versa.*

More precisely, suppose that S is a subset of the set of all objective predicates from \mathcal{L} and let Ψ be the operator simultaneously replacing all occurrences of the atom $\neg A$ by A and all occurrences of the atom A by $\neg A$, for all atoms A in S . Then T° is a static expansion of a belief theory T in AEB if and only if $\Psi(T^\circ)$ is a static expansion of the theory $\Psi(T)$. \square

2.1 Strong Negation and Logic Programming

Since logic programs under major semantics, including *stable semantics* [14], *well-founded semantics* [34], *partial stable* or *stationary semantics* [27] and *static semantics* [30], can be translated into belief theories in AEB via the embedding $T_{\mathcal{B}\neg}(P)$ defined in [30, 31]⁸, the introduction of strong negation into belief theories immediately introduces strong negation into logic programs with these semantics.

In particular, it follows from [15, 27] that stable semantics augmented with strong negation is equivalent to stable semantics with the so called “classical negation”, originally introduced by Gelfond and Lifschitz [15].

Theorem 4 Strong Negation Extends “Classical” Negation. *There is a one-to-one correspondence between stable models \mathcal{M} of a logic program P with “classical negation” and consistent static autoepistemic expansions T° of its translation $T_{\mathcal{B}\neg}(P)$ into belief theory that satisfy the condition $\mathcal{B}A \vee \mathcal{B}\neg A$, for all objective atoms A . We assume here that “classical negation” of an atom A is translated into its strong negation $\neg A$. \square*

⁸ According to $T_{\mathcal{B}\neg}(P)$, a logic programming rule:

$$A_1 \vee \dots \vee A_k \leftarrow B_1 \wedge \dots \wedge B_m \wedge \text{not } C_1 \wedge \dots \wedge \text{not } C_n.$$

is translated into: $B_1 \wedge \dots \wedge B_m \wedge \mathcal{B}\neg C_1 \wedge \dots \wedge \mathcal{B}\neg C_n \supset A_1 \vee \dots \vee A_k$.

From Proposition 2 and the fact that default negation, $not F$ is translated into $\mathcal{B}\neg F$, one immediately derives an important relationship between strong negation and default negation in logic programs:

Proposition 5 Strong Negation vs. Default Negation. *Suppose that P is a logic program. The inference rule $\frac{\neg F}{not F}$ is satisfied by all semantics of P obtained by the embedding of P into AEB via the translation $T_{\mathcal{B}\neg}(P)$ and for all objective formulae F such that P includes the strong negation constraint S_A , for every atom A that occurs in F . \square*

It is also worth noting that, for atomic objective formulae A , the default axioms discussed above have a particularly simple translation into logic programming rules: $\neg A \leftarrow not A$ and $A \leftarrow not \neg A$.

3 Explicit Negation

In this section we introduce an alternative notion of symmetric negation in the Autoepistemic Logic of Beliefs, which we call *explicit negation*. The resulting extension of AEB , called the *Autoepistemic Logic of Beliefs with Explicit Negation*, $AEBX$, demonstrates the flexibility of the AEB framework in expressing different forms of negation.

After providing motivation and formal definitions, we show that static expansions in $AEBX$ are sufficiently expressive to characterize the *well-founded semantics with explicit negation*, $WFSX$, a semantics of logic programs introduced earlier in [23]. We then contrast explicit negation with strong negation. Section 4 illustrates an application of explicit negation showing how belief revision in $AEBX$ can directly capture the contradiction removal techniques used for logic programs [4].

3.1 The Issue of Relevance

In logic programming, clauses are seen as inference rules rather than material implications. However, for definite and normal programs this distinction is immaterial. Due to the absence of negative facts in such programs, the addition of a contrapositive $\neg Head \rightarrow \neg Body$ of a program rule $Head \leftarrow Body$ has no bearing on the rest of the program.

On the other hand, in normal logic programs extended with a symmetric negation (hereafter, simply called *extended logic programs* or $ELPs$) some care is needed if one wants to preserve the procedural reading of logic program rules. According to this reading, as in a procedure, the truth value of the head of a rule is solely determined by the truth value of its body. Thus, the procedural reading indicates that computing the truth value of a literal can be done by relying solely on the procedural call graph implicitly defined by the rules for literals. In other words, literals that are not (transitively) called by the rules for another literal should not influence its truth value. This well-known property, called *relevance*

[10], is essential to guarantee the availability of (strictly) top-down evaluation procedures for a semantics. It is worth noting that the well-founded semantics for normal logic programs [12] obeys relevance (cf. [10]). On the other hand, neither the Stable Models Semantics, nor *AEB* when applied to theories resulting from logic programs with the strong negation constraints, satisfy this principle.

One paramount motivation for introducing *AEBX* is the desire to capture those semantics of logic program that satisfy relevance and thus permit efficient, top-down implementations. As shown in Example 2, queries for non-relevant semantics cannot, in general, be evaluated in a top-down fashion using standard logic programming implementation procedures, i.e., by simply following the program’s call graph. The other was the desire to ensure that *AEBX* is more expressive than *AEB* by providing a meaningful semantics to those programs that appear to have a well-defined intended meaning and yet do not have a consistent semantics when strong negation is used (see Example 3).

Example 2. Take the following theory T and corresponding logic program P :

$$\begin{array}{ll} \mathcal{B}\neg god_exists \supset god_exists & god_exists \leftarrow not \neg god_exists \\ \mathcal{B}\neg god_exists \supset \neg god_exists & \neg god_exists \leftarrow not god_exists \\ \mathcal{B}\neg god_exists \supset \neg go_to_church & \neg go_to_church \leftarrow not god_exists \\ go_to_church & go_to_church \end{array}$$

where the first two rules represent two conflicting default axioms, which read “I conclude God exists if I disbelieve Its non-existence” and “I conclude God doesn’t exist if I disbelieve Its existence”. The third rule states that “If I disbelieve the existence of God then I do not go to church”, and the fourth that “I go to church”. If \neg is understood as strong negation, i.e. if strong negation constraints are added to T , then the theory has one expansion, where god_exists because I go to church. Indeed, $go_to_church \supset \neg go_to_church$ (strong negation constraint), which by contraposition of the third clause entails $\neg \mathcal{B}\neg god_exists$. Therefore, since $\neg god_exists$ only appears in the second clause, $\neg god_exists$ holds in all minimal models, and by necessitation rule (N), $\mathcal{B}\neg god_exists$. Hence, by the first clause of T , god_exists holds in all expansions of the theory. Note that this conclusion is not *relevant*. Looking at the program P makes it clear that only the first two rules are (transitively) called by god_exists . However, the reader may check that the theory consisting solely of the first two clauses, does not have god_exists in all its expansions. With explicit negation (to be defined below) we have different conclusions, and the corresponding least expansion includes $\{go_to_church, \mathcal{B}\neg go_to_church\}$ but not god_exists nor $\mathcal{B}\neg god_exists$ (cf. Example 4). \square

Example 3. Take now the following theory and corresponding logic program:

$$\begin{array}{ll} \mathcal{B}\neg shave(x, x) \supset shave(John, x) & shave(john, X) \leftarrow not shave(X, X) \\ shave(y, x) \supset go_dine_out(x) & go_dine_out(X) \leftarrow shave(Y, X) \\ \neg shave(Peter, Peter) & \neg shave(peter, peter) \\ \neg go_dine_out(John) & \neg go_dine_out(john) \end{array}$$

The first rule states that “John shaves everyone not believed to shave themselves”. The second says that “If x has been shaved (by anyone) then x will go out to dine”. The third states that “Peter does not shave himself”, and the

fourth that “John has not gone out to dine”. We would like to know whether we believe John has shaved himself given that he has not gone out to dine. Note that believing he has not shaved himself leads to a contradiction, and that the conclusion that he has shaved himself is not true in all minimal models.

If the strong negation constraints are added to this theory then AEB assigns no consistent meaning to it (there is no consistent expansion). On the contrary, if explicit negation, to be defined below, is used instead the theory has one expansion that includes: $\{-go_dine_out(john), \mathcal{B}\neg go_dine_out(john)\}$ but not $shave(john, john)$. The absence of consistent expansions by using strong negation is brought about by the use of the contrapositive $\neg go_dine_out(X) \supset \neg shave(X, X)$ to conclude $\neg shave(john, john), \mathcal{B}\neg shave(john, john)$, and $shave(john, john)$ (contradiction). \square

As we have seen, the reason strong negation does not directly capture any semantics for extended logic programs complying with relevance, such as $WFSX$, is because of its very definition. Strong negation constraints, $\neg A \supset \neg A$, state that strongly negated facts or conclusions entail classically negated ones, thereby permitting the use of the contrapositives of the material implications resulting from the translation of the logic program rules.

What the above two paradigmatic examples have in common is the back propagation of truth values by strong negation, against the logic program rule arrow, into a loop of otherwise undefined literals (i.e., such that neither L nor $not\ L$ hold). In the Example 2 we have an even loop over default negation, and in the Example 3 an odd one. In the first example, the back propagation decides the loop one way, and in the second it comes up against the impossibility of resolving the loop by imposing a truth value. However, as we shall see below, explicit negation does not affect either loop.

To conclude, because of its use of classical negation contrapositives, strong negation leads both to logic programs without a semantics and to logic program theories with unwarranted (non-relevant) conclusions, i.e. conclusions not solely based on the procedural call graph of the logic program. To be able to capture relevant logic program semantics a weaker notion of symmetric negation is needed. Theories with explicit negation which are translatable into extended logic programs can be efficiently queried by the top-down procedural implementation technology of logic programs [2].

3.2 Introducing Explicit Negation

Explicit negation is added to AEB by means of an explicit negation operator $\overline{}$, thus defining the *Autoepistemic Logic of Beliefs with Explicit Negation*, $AEBX$. Specifically, this is accomplished by:

- Augmenting the original objective language \mathcal{L} with new *objective* propositional symbols \overline{A} , called *dual or explicit negation atoms*, resulting in a new objective language $\widehat{\mathcal{L}}$ and the new language of beliefs $\widehat{\mathcal{L}}_{\mathcal{A}\mathcal{E}\mathcal{B}}$. The extension of explicit negation to arbitrary positive objective formulae can be done in the same way as for strong negation.

- Extending the logical closure operator Cn_{AEB} with the following *Coherence* inference rule, for every objective propositional symbol A ⁹

$$\frac{\overline{A}}{\mathcal{B}\neg A}$$

A Coherence inference rule says that if one derives the dual, one has to believe its negation, i.e. “ \overline{A} serves as evidence against A ”. Since the Coherence inference rules have no effect on belief theories that do not include explicit negation atoms, in the sequel we will assume them as part of the operator Cn_{AEB} without further mention, whenever explicit negation is used.

Example 4. The details of this example show the essence of how explicit negation treats both previous examples, and the way it differs from strong negation.

The theory $T^\circ = Cn_{AEB}(T \cup \{\mathcal{B}gtc, \mathcal{B}\neg gtc\})$ is, with the obvious abbreviations, an expansion of the *AEBX* theory in Example 2 (where strong negation is replaced by explicit negation).

This example illustrates why explicit negation does not affect the theory’s even loop and, for the same reason, why it does not affect the odd loop of Example 3. Indeed, in the general case, the explicit negation of the head of a program rule may be true even though its body is undefined (i.e., such that neither $\mathcal{B}Body$ nor $\mathcal{B}\neg Body$ hold in an expansion). In other words, explicit negation allows the overriding to false of a rule’s head when its body is undefined. Because of this feature, there is no backward propagation of falsity of the head to the rule’s body. On the other hand, when the rule’s body is true, then its head must necessarily be true, which, however, represents a forward, rather than backwards, propagation of truth values. \square

The definition of explicit negation, contrary to that of strong negation, does not prevent the existence in a model of both A and \overline{A} , for some atom A . However, this kind of “paraconsistency” in models does not spill over to *AEBX* expansions:

Proposition 6. *Let T° be a consistent expansion of a belief theory T in *AEBX*. Then, for no atom A :* $T^\circ \models A \wedge \overline{A}$

The following result shows that the relevant logic program semantics *WFSX* (defined in [23]) is embeddable in *AEBX*. This embeddability result requires, besides the translation defined in [30, 31], a preliminary *WFSX*-semantics preserving transformation of the logic program. This transformation consists in the complete elimination of objective literal goals from rule bodies by means of unfolding, that is by successively replacing them by their various alternative rule definitions. The obtained programs are said to be in “semantic kernel form”.

Theorem 7 Explicit Negation and *WFSX*. *Let P be an extended logic program in the semantic kernel form. There is a one-to-one correspondence between the partial stable models \mathcal{M} of P , as defined in [23], and the consistent static autoepistemic expansions T° of its translation $T_{\mathcal{B}\neg}(P)$ into an *AEBX* belief theory, where “explicit negation” of an atom A is translated into \overline{A} . \square*

⁹ If $A = \overline{F}$ then we have $\frac{F}{\mathcal{B}\neg F}$.

Up to now, we have mainly considered *AEBX* theories resulting from the translation of (non-disjunctive) logic programs. This is so because we have motivated explicit negation by contrast with strong negation on such programs. Nevertheless, there is motivation to extend *AEB* with explicit negation to general theories, and not just logic programs.

The theory of Example 1, with strong negation replaced by explicit negation everywhere, illustrates such general use. In fact, the results will be the same, except when later one adds *good_researcher(Paul)*, expressing that there is evidence against Paul being a good researcher¹⁰. As shown in Example 1, by using strong negation, one additionally concludes that Paul will receive a negative evaluation, and that Keith is a good researcher and so receives a positive evaluation. When explicit negation is used instead, we still conclude that Paul will receive a negative evaluation, but we no longer surmise that Keith is a good researcher. Instead, a milder conclusion follows, that Keith is *believed to be* a good researcher. This conclusion is not enough to give Keith a good evaluation. See Example 5, for a detailed explanation of how these results are obtained.

3.3 Explicit and Strong Negation Compared

Explicit and strong negation share some similarities. The (derived) inference rule from strong negation to beliefs shown in Proposition 2, also holds for explicit negation. Explicit negation is in fact characterized by taking this inference rule as primitive, since it is no longer derivable in the absence of the strong negation axiom.

Another important similarity between both regards the symmetry between atoms and their negations. The result of Proposition 3 still holds if strong negation is replaced by explicit negation. In other words, explicit negation is also symmetric.

Because of the similarities, strong and explicit negation are even equivalent for the class of theories whose expansions capture the semantics of logic programs under Stable Models. Indeed, the result in Theorem 4 remains true regardless of whether “classical” negation is translated into strong or into explicit negation.

For logic programs, explicit negation can be seen as an approximation to strong negation, in the sense that all the relevant consequences of a belief theory with strong negation remain true in *AEBX* after strong negation is replaced everywhere by explicit negation:

Proposition 8. *Let T_s be a *AEB* theory with strong negation obtained from an extended logic program P via the translation described in Section ??, T_x be the *AEBX* theory obtained from T_s by replacing strong by explicit negation and deleting all the strong negation constraints in T_s , and let T_s^\diamond be an expansion of T_s . Then there exists an expansion T_x^\diamond of T_x such that, for every objective atom A , $\mathcal{B}A \in T_s^\diamond$ if and only if $\mathcal{B}A \in T_x^\diamond$, and $\mathcal{B}\neg A \in T_s^\diamond$ if and only if $\mathcal{B}\neg A \in T_x^\diamond$. \square*

¹⁰ In Example 1 we added instead *bad_researcher(Paul)*, standing for \neg *good_researcher(Paul)*.

Corollary 9 Explicit Negation Approximates Strong Negation. *Let T_s and T_x be as in Proposition 8. If some formula F of the form $\mathcal{B}L$, where L is either an objective atom A or its negation $\neg A$, holds in every expansion of T_x , then F also holds in every expansion of T_s . \square*

Consequently, query evaluation procedures for explicit negation in logic programs can be used as sound query evaluation procedures for strong negation in logic programs¹¹. However those procedures are not complete for strong negation since, as shown by Examples 2 and 3, the converse of Corollary 9 is not true in general. Another result contrasting explicit and strong negation is:

Theorem 10 Explicit Negation Extends Strong Negation. *There is a one to one correspondence between expansions of a theory T_s with strong negation and expansions of the AEBX theory T_x obtained by replacing in T_s strong by explicit negation, and by adding, for every atom A , the clause: $\overline{A} \supset \neg A$. \square*

Disjunctive Syllogism is a derived inference rule applicable to classical and strong negations, but which is not enjoyed by explicit negation. This rule typically allows one to conclude A on the strength of $A \vee B$ and $\neg B$ (or $\neg B$).

Example 5. Consider the theory:

$$\overline{\text{good_researcher}(\text{Paul})} \quad \text{good_researcher}(\text{Keith}) \vee \text{good_researcher}(\text{Paul})$$

stating that “there is evidence against Paul being a good researcher”, and that “at least one of Keith or Paul is a good researcher”.

Whereas with strong negation $\text{good_researcher}(\text{Keith})$ follows from the theory by virtue of the Disjunctive Syllogism, with explicit negation it is not so. Instead, a milder conclusion obtains: First, note that $\text{good_researcher}(\text{Keith}) \vee \text{good_researcher}(\text{Paul})$ is equivalent to the formula $\neg \text{good_researcher}(\text{Paul}) \supset \text{good_researcher}(\text{Keith})$. Now, applying Necessitation, $\mathcal{B}(\neg \text{good_researcher}(\text{Paul}) \supset \text{good_researcher}(\text{Keith}))$. By then applying axiom **(K)**, we conclude $\mathcal{B}\neg \text{good_researcher}(\text{Paul}) \supset \mathcal{B}\text{good_researcher}(\text{Keith})$. Since the premise of this implication follows from $\overline{\text{good_researcher}(\text{Paul})}$ by Coherence, we extract the milder conclusion $\mathcal{B}\text{good_researcher}(\text{Keith})$.

Intuitively, having evidence against Paul being a good researcher does not ensure that John is not a good researcher, and so does not warrant Keith being a good researcher. However, this information is enough for believing Paul is not a good researcher, and consequently to believe Keith is. Contrast this reading with the one of strong negation, where $\neg \text{good_researcher}(\text{Paul})$ means “John is a bad researcher”. In this case, and knowing that either Paul or Keith is a good researcher, it is expected that Keith is a good researcher. \square

¹¹ See [2] for a description of such evaluation procedures.

4 Application of Explicit Negation

Applications of the logic programming semantics *WFSX* have been studied in various domains, including hypothetical reasoning [24], model-based diagnosis [26], and declarative debugging of logic programs [25]. Many of these applications require belief revision methods, via contradiction removal techniques [4]. Here we show how to capture belief revision in *AEBX* and we illustrate how to extend the above mentioned applications from the class of logic programs to the significantly broader class of belief theories.

Unlike normal programs, extended logic programs with symmetric negation may be contradictory, i.e., may not have a (consistent) semantics. While for some programs this seems reasonable (e.g., for a program with the two facts a and \bar{a}), for some others this may not be justified:

Example 6. Consider the program $P = \{runs \leftarrow not\ broken; \overline{runs}\}$ stating that “if we assume that a car is not broken, then it runs”, and that “the car does not run”.

None of the logic programming semantics mentioned in this paper assigns a consistent meaning to this program. Indeed, since there are no rules in the program for *broken*, all of the semantics assume *not broken* true ($\neg broken$ is true in all minimal models of the corresponding *AEB* theory), and so both *runs* and \overline{runs} hold. If \overline{runs} in the logic program is understood as the strong negation of *runs* then, since $runs \wedge \neg runs \supset \perp$, no consistent expansion of the corresponding belief theory exists. The same happens if \overline{runs} is understood as the explicit negation of *runs*, cf. Proposition 6.

But one can argue that, since the car does not run, it follows by “reductio ad absurdum” that our assumption that the car is not broken must be incorrect and thus has to be *revised*. \square

The Contradiction Removal with Explicit Negation (*CRSX*) technique of [4] removes contradiction from logic programs under *WFSX* (wherever possible), by revising any default assumption that would otherwise lead to contradiction. To do so, it adds to programs a rule of the form $L \leftarrow not\ L$, for each such assumption *not L*, with the effect that no model of the program can now contain *not L*. The alternative minimal contradiction removing sets of such “inhibition rules” can be added to the original program in order to obtain the alternative minimally revised programs.

Example 7. The only revision of the program in Example 6, is obtained by adding to it $broken \leftarrow not\ broken$. The revised program is no longer contradictory: its *WFSX* is $\{\overline{runs}, not\ runs\}$. Furthermore, this addition is minimal. \square

The addition of inhibition rules can be allowed for only some pre-designated set of literals – the *revisables*. These are the literals that can be independently revised. Other literals may have their truth value revised but only as a consequence of changes in the revisables. Restricting the revisables is crucial for controlling the causal level at which assumption withdrawal may be performed.

Indeed, in an application domain such as diagnosis, revising the functional normality assumption about a component may be conditional upon revising the functional normality assumption about some subcomponent. Conversely, it may suffice to hypothesize the abnormality of a subcomponent to put in question the normality of the component containing it. We will deal with both issues.

Example 8. Consider now the (contradictory) program:

$$\begin{array}{ll} runs \leftarrow not\ broken & broken \leftarrow flatTire \\ \overline{runs} & broken \leftarrow badBattery \end{array}$$

If addition of inhibition rules is allowed for any literal, then one such minimal addition removing the contradiction is $\{broken \leftarrow not\ broken\}$. This revision can be seen as a diagnosis of the car simply stating the car might be broken. However, in this case one would like the diagnosis to delve deeper into the car's functional structure, and obtain two minimal diagnoses: one suggesting a possible problem with a flat tire, and another suggesting a possible problem with a bad battery. This is achieved by declaring as revisables only *badBattery* and *flatTire*. \square

This declaration of revisables is akin to the declaration of abducible literals in abductive logic programming. There, a solution to an abductive query is obtained by minimally, and consistently, adding to the theory facts needed in order to prove the abductive query literal. Moreover, addition of facts is only allowed for literals declared as abducibles¹².

Example 9. Consider the program obtained from Example 8 by removing the fact \overline{runs} , i.e. the program $\{runs \leftarrow not\ broken; broken \leftarrow flatTire; broken \leftarrow badBattery\}$. If every literal is abducible, then one abductive solution for *not runs* is $\{broken\}$ (i.e., this is one possible explanation of the fact that the car is not running). To obtain only the deeper diagnosis of the car, simply declare as abducibles $\{flatTire, badBattery\}$: the abductive solutions to that query are then $\{flatTire\}$ and $\{badBattery\}$. \square

This technique can easily be imported into *AEBX*. Define the revisions of a theory T which has no expansions (a *contradictory* theory), as the non-contradictory theories obtained by minimally adding clauses of the form $\mathcal{B} \neg L \supset L$, which disallow (consistent) expansions in which $\mathcal{B} \neg L$ is true¹³.

Note that, as opposed to logic programs, in *AEBX* there is no need for recourse to a meta-linguistic declaration of revisables. The language itself is sufficiently expressive to handle the definition of revisables. To realize this let us get back to the car example, which in *AEBX* is rendered as the belief theory T :

$$\begin{array}{ll} \mathcal{B} \neg broken \supset runs & badBattery \supset broken \\ \overline{runs} & flatTire \supset broken \end{array}$$

¹² A formal comparison of contradiction removal and abduction in logic programs can be found in [9].

¹³ An expansion with $\mathcal{B} \neg L$ would also contain L and $\mathcal{B}L$, since expansions are closed under necessitation, and would thus be inconsistent.

where *broken* is non-revisable. In *CRSX* this declaration of non-revisability means that *broken* may have its truth value revised only indirectly, as a consequence of changes in the revisables, but never directly on account of the addition of an inhibition rule for it. In *AEBX* it means that beliefs about *broken* may change only indirectly, as a consequence of changes in beliefs about revisables, but never directly because of the addition of an inhibition clause for it. To achieve this, we must require that belief changes about *broken* be solely determined by the belief changes about the revisables, so that adding a single inhibition clause, for *broken* only, is insufficient because the resulting theory has no consistent expansions. By Necessitation and the axiom **(K)**, the closure of T contains:

$$\mathcal{B}flatTire \vee \mathcal{B}badBattery \supset \mathcal{B}broken.$$

Thus, belief in the truth of *broken* follows from belief in *flatTire* or belief in *badBattery*. For the falsity of belief in *broken* to be determined by the falsity of beliefs in *flatTire* and *badBattery* we need an additional statement, which ensures that if both *flatTire* and *badBattery* are disbelieved then *broken* must be disbelieved as well, i.e.:

$$\mathcal{B}\neg flatTire \wedge \mathcal{B}\neg badBattery \supset \mathcal{B}\neg broken \quad (1)$$

Example 10. The belief theory T , augmented with clause (1) has two revisions: $T \cup \{\mathcal{B}\neg badBattery \supset badBattery\}$ and $T \cup \{\mathcal{B}\neg flatTire \supset flatTire\}$ each corresponding to one of the desired deeper diagnoses.

Observe that $T' = T \cup \{\mathcal{B}\neg broken \supset broken\}$ is not a revision. Indeed, all minimal models of the theory have $\neg flatTire$ and $\neg badBattery$, because there are no clauses defined for neither *flatTire* nor *badBattery*. Thus every expansion of T' must contain $\{\mathcal{B}\neg flatTire, \mathcal{B}\neg badBattery\}$. Now, by clause (1), every expansion of T' has $\mathcal{B}\neg Broken$, and thus, by the inhibition clause for *broken* and necessitation, is inconsistent. \square

Notice the similarity between clause (1) and Clark's completion [7] of *broken*. The latter implies that if both *flatTire* and *badBattery* are false then *broken* is false, whilst (1) states the same about the corresponding beliefs. For this reason (1) is called the *belief completion clause* for *broken*. The formal definition of belief completion rules can be found in [1].

The specification of revisables is captured in the language of *AEBX* by adding for each literal which is *not* a revisable the corresponding belief completion clause. This language level declarative specification achieves the same effect as the declaration of revisables for logic programs, which in the latter case can only be made meta-linguistically. Though not detailed here, it can also be used to explain the intuitive notion that a literal is abducible unless the rules defining it are closed by a completion clause. Moreover it has greater generality, by allowing conditions to be added to belief completion clauses. For example, if instead of simply stating that *broken* is not a revisable, we want to say that *broken* is not revisable only if the atom *deeper* is a consequence of the theory. In such a case we would just add:

$$deeper \wedge \mathcal{B}\neg flatTire \wedge \mathcal{B}\neg badBattery \supset \mathcal{B}\neg broken.$$

This allows for flexible control over the level of revision: adding the fact *deeper*, or concluding it, results in only causally deeper faults being obtained. Otherwise superficial faults can be obtained. Diverse complex diagnosis preferences and strategies can be encoded in *AEBX* (see [8] to find out how this can be accomplished by changing the revisables in logic programs).

We conclude by remarking that either one of symmetric negation in *AEB* can be used to obtain a solution to belief revision not relying on a theory transformation: instead of the addition of minimal sets of inhibition rules, alternative revisions can be obtained directly by modifying the definition of expansions leading to the notion of *Careful Autoepistemic Expansions* (see [1]).

5 Concluding Remarks

We have shown that, to represent and reason about knowledge, one requires, besides the usual default negation, another form of negation, symmetric with respect to default negation, and that classical negation does not fulfill that rôle.

We then introduced, into the quite general framework of Autoepistemic Logic of Belief two types of symmetric negation, strong and explicit, and illustrated their application on a number of knowledge representation examples.

Although similar, in logic programs strong and explicit negation differ in the use made of contrapositives by strong negation, and in the amenability by explicit negation to strictly top-down querying procedures. They are alike in that they both capture the Answer-Sets semantics [15], differ from classical negation, and enjoy symmetry. Of the two, explicit negation is weaker and easier to implement, and thus it can be used as an approximation to strong negation.

References

1. J. Alferes, L. Pereira, and T. C. Przymusiński. Belief revision in logic programming and non-monotonic reasoning. In *Progress in AI, Proceedings of the International Artificial Intelligence Conference, EPIA '95*, pages 41–56, 1995.
2. J. J. Alferes, C. V. Damásio, and L. M. Pereira. SLX – A top-down derivation procedure for programs with explicit negation. In M. Bruynooghe, editor, *International Symposium on Logic programming*. MIT Press, 1994.
3. J. J. Alferes and L. M. Pereira. On logic program semantics with two kinds of negation. In K. Apt, editor, *Int. Joint Conf. and Symp. on LP*. MIT Press, 1992.
4. J. J. Alferes and L. M. Pereira. Contradiction: when avoidance equal removal. In R. Dyckhoff, editor, *4th Int. Ws. on Extensions of LP*, volume 798 of *LNAI*. Springer-Verlag, 1994.
5. J. J. Alferes, L. M. Pereira, and T. C. Przymusiński. “Classical” negation in non monotonic reasoning and logic programming. In H. Kautz and B. Selman, editors, *4th Int. Symposium on Artificial Intelligence and Mathematics*. Florida Atlantic University, 1996.
6. Stefan Brass, Juergen Dix, and Teodor C. Przymusiński. Characterizations and implementation of static semantics of disjunctive programs. (In preparation), University of California at Riverside, 1996.

7. K.L. Clark. Negation as failure. In H. Gallaire and J. Minker, editors, *Logic and Data Bases*, pages 293–322. Plenum Press, New York, 1978.
8. C. V. Damásio, W. Nejdl, L. M. Pereira, and M. Schroeder. Model-based diagnosis preferences and strategies representation with logic meta-programming. In K. Apt and F. Turini, editors, *Meta-logics and Logic Programming*, pages 311–338. MIT Press, 1995.
9. C. V. Damásio and L. M. Pereira. Abduction over 3-valued extended logic programs. In V. Marek, A. Nerode, and M. Truszczynski, editors, *Proceedings of the Third International Conference on Logic Programming and Non-Monotonic Reasoning*. LPNMR’95, Springer-Verlag, 1995.
10. Jürgen Dix. A Classification-Theory of Semantics of Normal Logic Programs: II. Weak Properties. *Fundamenta Informaticae*, XXII(3):257–288, 1995.
11. P. M. Dung and P. Ruamviboonsuk. Well founded reasoning with classical negation. In A. Nerode, W. Marek, and V. S. Subrahmanian, editors, *LP & NMR*, pages 120–132. MIT Press, 1991.
12. A. Van Gelder, K. A. Ross, and J. S. Schlipf. The well-founded semantics for general logic programs. *Journal of the ACM*, 38(3):620–650, 1991.
13. M. Gelfond. Logic programming and reasoning with incomplete information. Technical report, University of Texas at El Paso, 1992.
14. M. Gelfond and V. Lifschitz. The stable model semantics for logic programming. In R. Kowalski and K. Bowen, editors, *Proceedings of the Fifth Logic Programming Symposium*, pages 1070–1080, Cambridge, Mass., 1988. Association for Logic Programming, MIT Press.
15. M. Gelfond and V. Lifschitz. Logic programs with classical negation. In *Proceedings of the Seventh International Logic Programming Conference, Jerusalem, Israel*, pages 579–597, Cambridge, Mass., 1990. Association for Logic Programming, MIT Press.
16. Michael Gelfond, Halina Przymusinski, and Teodor C. Przymusinski. On the relationship between circumscription and negation as failure. *Journal of Artificial Intelligence*, 38(1):75–94, February 1989.
17. J. McCarthy. Circumscription – a form of non-monotonic reasoning. *Journal of Artificial Intelligence*, 13:27–39, 1980.
18. J. Minker. On indefinite data bases and the closed world assumption. In *Proc. 6-th Conference on Automated Deduction*, pages 292–308. Springer Verlag, 1982.
19. R.C. Moore. Semantic considerations on non-monotonic logic. *Journal of Artificial Intelligence*, 25:75–94, 1985.
20. D. Nelson. Constructible falsity. *JSL*, 14:16–26, 1949.
21. D. Pearce. Reasoning with negative information II: Hard negation, strong negation and logic programs. In D. Pearce and H. Wansing, editors, *Nonclassical Logics and Information Processing*, pages 63–79. Springer-Verlag, 1990.
22. D. Pearce and G. Wagner. Reasoning with negative information I: Strong negation in logic programs. In L. Haaparanta, M. Kusch, and I. Niiniluoto, editors, *Language, Knowledge and Intentionality*, pages 430–453. Acta Philosophica Fennica 49, 1990.
23. L. M. Pereira and J. J. Alferes. Well founded semantics for logic programs with explicit negation. In B. Neumann, editor, *European Conf. on AI*, pages 102–106. John Wiley & Sons, 1992.
24. L. M. Pereira, J. N. Aparício, and J. J. Alferes. Non-monotonic reasoning with logic programming. *Journal of Logic Programming. Special issue on Nonmonotonic reasoning*, 17(2, 3 & 4):227–263, 1993.

25. L. M. Pereira, C. Damásio, and J. J. Alferes. Debugging by diagnosing assumptions. In P. A. Fritzson, editor, *Automatic Algorithmic Debugging, AADEBUG'93*, volume 749 of *Lecture Notes in Computer Science*, pages 58–74. Springer-Verlag, 1993.
26. L. M. Pereira, C. Damásio, and J. J. Alferes. Diagnosis and debugging as contradiction removal. In L. M. Pereira and A. Nerode, editors, *2nd Int. Ws. on LP & NMR*, pages 316–330. MIT Press, 1993.
27. T. C. Przymusinski. The well-founded semantics coincides with the three-valued stable semantics. *Fundamenta Informaticae*, 13(4):445–464, 1990.
28. T. C. Przymusinski. Stable semantics for disjunctive programs. *New Generation Computing Journal*, 9:401–424, 1991. (Extended abstract appeared in: Extended stable semantics for normal and disjunctive logic programs. *Proceedings of the 7-th International Logic Programming Conference, Jerusalem*, pages 459–477, 1990. MIT Press.).
29. T. C. Przymusinski. Autoepistemic logic of knowledge and beliefs. (In preparation), University of California at Riverside, 1995. (Extended abstract appeared in ‘A knowledge representation framework based on autoepistemic logic of minimal beliefs’ In *Proceedings of the Twelfth National Conference on Artificial Intelligence, AAAI-94, Seattle, Washington, August 1994*, pages 952–959, Los Altos, CA, 1994. American Association for Artificial Intelligence, Morgan Kaufmann.).
30. T. C. Przymusinski. Static semantics for normal and disjunctive logic programs. *Annals of Mathematics and Artificial Intelligence*, Special Issue on Disjunctive Programs(14):323–357, 1995.
31. Teodor Przymusinski. Semantics of normal and disjunctive logic programs: A unifying framework. In J. Dix, L. Pereira, and T. Przymusinski, editors, *Proceedings of the Workshop on Non-Monotonic Extensions of Logic Programming at the Eleventh International Logic Programming Conference, ICLP'95, Santa Margherita Ligure, Italy, June 1994*, pages 43–67. Springer Verlag, 1995.
32. R. Reiter. On closed-world data bases. In H. Gallaire and J. Minker, editors, *Logic and Data Bases*, pages 55–76. Plenum Press, New York, 1978.
33. R. Reiter. A logic for default theory. *Journal of Artificial Intelligence*, 13:81–132, 1980.
34. A. Van Gelder, K. A. Ross, and J. S. Schlipf. The well-founded semantics for general logic programs. *Journal of the ACM*, 1990. (to appear). Preliminary abstract appeared in Seventh ACM Symposium on Principles of Database Systems, March 1988, pp. 221–230.
35. G. Wagner. Reasoning with inconsistency in extended deductive databases. In L. M. Pereira and A. Nerode, editors, *2nd Int. Ws. on LP & NMR*, pages 300–315. MIT Press, 1993.

A Autoepistemic Logic of Beliefs

Here we briefly recall the definition of the *Autoepistemic Logic of Beliefs*, *AEB*, introduced by Przymusinski in [29]. Consider a fixed propositional language \mathcal{L} with standard connectives (\vee , \wedge , \supset , \neg) and the propositional letter \perp (denoting *false*). The language of *AEB* is a propositional modal language \mathcal{L}_{AEB} obtained by augmenting the language \mathcal{L} with a modal operator \mathcal{B} , called the *belief* operator. The atomic formulae of the form $\mathcal{B}F$, where F is an arbitrary formula of

\mathcal{L}_{AEB} , are called *belief atoms*. The formulae of the original language \mathcal{L} are called *objective*. Observe that arbitrarily deep level of *nested beliefs* is allowed in belief theories. Any theory T in the language \mathcal{L}_{AEB} is called a *belief theory*.

Definition 11 Belief Theory. By an *autoepistemic theory of beliefs*, or just a *belief theory*, we mean an arbitrary theory in the language \mathcal{L}_{AEB} , i.e., a (possibly infinite) set of arbitrary clauses of the form:

$$B_1 \wedge \dots \wedge B_k \wedge \mathcal{B}G_1 \wedge \dots \wedge \mathcal{B}G_l \wedge \neg \mathcal{B}F_1 \wedge \dots \wedge \neg \mathcal{B}F_n \supset A_1 \vee \dots \vee A_m$$

where $k, l, m, n \geq 0$, A_i s and B_i s are objective atoms and F_i s and G_i s are arbitrary formulae of \mathcal{L}_{AEB} . Such a clause says that if the B_i s are true, the G_i s are believed, and the F_i s are not believed then one of the A_i s is true. \square

We assume the following two simple axiom schemata and one inference rule describing the arguably obvious properties of belief atoms:

(D) Consistency Axiom: $\neg \mathcal{B}\perp$.

(K) Normality Axiom: For any formulae F, G : $\mathcal{B}(F \supset G) \supset (\mathcal{B}F \supset \mathcal{B}G)$.

(N) Necessitation Rule: For any formula F : $\frac{F}{\mathcal{B}F}$

Definition 12 Formulae Derivable from a Belief Theory. For any belief theory T , we denote by $Cn_{AEB}(T)$ the smallest set of formulae of the language \mathcal{L}_{AEB} which contains the theory T , all the (substitution instances of) the axioms (K) and (D) and is closed under both standard propositional consequence and the necessitation rule (N). We say that a formula F is *derivable* from theory T in the logic AEB if F belongs to $Cn_{AEB}(T)$. A belief theory T is *consistent* if the theory $Cn_{AEB}(T)$ is consistent. \square

Definition 13 Minimal Models. [29] By a *minimal model* of a belief theory T we mean a model M of T with the property that there is *no* smaller model N of T which coincides with M on belief atoms $\mathcal{B}F$. If a formula F is true in all minimal models of T then we write: $T \models_{\min} F$ and say that F is *minimally entailed* by T . \square

The intended meaning of belief atoms $\mathcal{B}F$ is based on the principle of *predicate minimization*:

$$\mathcal{B}F \equiv F \text{ is minimally entailed} \equiv F \text{ is true in all minimal models.}$$

Definition 14 Static Autoepistemic Expansion. [29] A belief theory T° is called a *static autoepistemic expansion* of a belief theory T if it satisfies the following fixed-point equation $T^\circ = Cn_{AEB}(T \cup \{\mathcal{B}F : T^\circ \models_{\min} F\})$, where F ranges over all formulae of \mathcal{L}_{AEB} . \square