

Well-Founded semantics for Semi-Normal Extended Logic Programs

Martin Caminada*
Utrecht University

Abstract

In this paper we present a new approach for applying well-founded semantics to extended logic programs. The main idea is not to fundamentally change the definition of well-founded semantics (as others have attempted) but rather to define a few restrictions on the content of the extended logic program, that make it possible to apply “traditional” well-founded semantics in a very straightforward way.

Introduction

Well-founded semantics (van Gelder, Ross, & Schlipf 1991) has originally been stated as an alternative for stable model semantics in normal logic programs. Due to its skeptical nature, it has sometimes been regarded as an easily computable lower bound for the more credulous stable model semantics. At the same time, well-founded semantics avoids some of the problems of stable model semantics, in which relatively small pieces of information (like a rule $a \leftarrow \text{not } a$) can cause the total absence of stable models.

With the emergence of extended logic programming (Gelfond & Lifschitz 1991), several researchers have attempted to apply well-founded semantics to extended logic programs (Sakama 1992; Brewka 1996). The introduction of strong negation, however, introduces additional problems not present in normal (non-extended) logic programming. In this paper, we approach the issue of how to apply well-founded semantics for extended logic programs not by giving another complex and advanced specification of what well-founded semantics for extended logic programs should look like, but instead we state a few restrictions on the *content* of the extended logic programs. We then show that under these restrictions, a relatively simple and straightforward definition of well-founded semantics yields a decent and unproblematic well-founded model.

Basic Definitions

A *program* considered in this paper is an *extended logic program* (ELP) (Gelfond & Lifschitz 1991) containing rules with weak as well as strong negation.

definition 1. An extended logic program P is a finite set of clauses of the form:

$$c \leftarrow a_1, \dots, a_n, \text{not } b_1, \dots, \text{not } b_m (n \geq 0, m \geq 0)$$

where each c , a_i and b_j is a positive/negative literal and not stands for negation as failure. In the above rule, $b_j (1 \leq j \leq m)$ is called a weakly negated literal. The literal c is called the head of the rule, and the conjunction $a_1, \dots, a_n, \text{not } b_1, \dots, \text{not } b_m$ is called the body of the rule. A rule is called *strict* iff it contains no weakly negated literals (that is, if $m = 0$), otherwise, the rule is *defeasible*.

Notice that the head of a rule is never empty, although the body can be empty. If \perp is a literal, then we identify $\neg\perp$ with \perp . If P is an extended logic program, then $\text{strict}(P)$ stands for the set of strict rules in P , and $\text{defeasible}(P)$ stands for the set of defeasible rules in P .

The *closure* of a set of strict rules consists of all literals that can be derived with it, as is stated in the following definition.

definition 2. Let S be a set of strict rules. We define $Cl(S)$ as the smallest set of literals such that if S contains a rule $c \leftarrow a_1, \dots, a_n$ and $a_1, \dots, a_n \in Cl(S)$ then $c \in Cl(S)$.

If S is a set of strict rules and L a set of literals, then we write $Cl(S \cup L)$ as an abbreviation of $Cl(S \cup \{\perp \leftarrow \mid \perp \in L\})$.

definition 3. We say that a set of literals L is *consistent* iff L does not contain a literal \perp and its negation $\neg\perp$. We say that a set of strict rules S is *consistent* iff $Cl(S)$ is *consistent*.

The idea of P^L (the Gelfond-Lifschitz reduct of a logic program P under a set of literals L) is to remove each rule from P that is “defeated” by L (that is, to remove each rule containing a weakly negated literal in L) and then from the remaining rules to remove all remaining occurrences of weak negation.

definition 4. Let P be an extended logic program and let L be a set of literals. We define P^L as $\{c \leftarrow a_1, \dots, a_n \mid c \leftarrow a_1, \dots, a_n, \text{not } b_1, \dots, \text{not } b_m \in P(n, m \geq 0) \text{ and } \neg\exists b_j (1 \leq j \leq m) : b_j \in L\}$.

Well-founded semantics (van Gelder, Ross, & Schlipf 1991) is a concept originally proposed for non-extended logic programs. As its original description is quite complex, we will use the following definition instead (inspired by (Brewka 1996)).

*This work has been supported by the EU ASPIC project.

definition 5. Let P be an extended logic program and L be a set of literals. We define $\gamma(L)$ (the standard stable operator) as $Cl(P^L)$. We define $\Gamma(L)$ as $\gamma(\gamma(L))$. The well-founded model of P is the smallest fixpoint of Γ .

The Problem

Well-founded semantics (WFS) has been applied successfully in non-extended logic programs (Dix 1995a; 1995b). Applying WFS for extended logic programs, however, introduces the problem that the well-founded model is not guaranteed to be consistent. Consider the following example, taken from (Caminada & Amgoud 2005).

example 1.

“John wears something that looks like a wedding ring.”

“John parties with his friends until late.”

“Someone wearing a wedding ring is usually married.”

“A party-animal is usually a bachelor.”

“A married person, by its definition, has a spouse.”

“A bachelor, by definition, does not have a spouse.”

These sentences are represented by the program P :

| | |
|--------------------------------------|--------------------------------------|
| $r \leftarrow$ | $p \leftarrow$ |
| $m \leftarrow r, \text{not } \neg m$ | $b \leftarrow p, \text{not } \neg b$ |
| $hs \leftarrow m$ | $\neg hs \leftarrow b.$ |

For example 1, applying the unaltered version of WFS yields a well-founded model of $\{r, p, m, b, hs, \neg hs\}$, which is inconsistent.

To cope with this problem, many approaches have been stated. Brewka, for instance, proposes to define the function $\Gamma(L)$ not as $\gamma(\gamma(L))$ but as $\gamma(Cn(\gamma(L)))$, where $Cn(L)$ is L if L is consistent, or Lit if L is not consistent (Brewka 1996). Another approach would be to apply paraconsistent reasoning, as for instance has been done in (Sakama 1992).

An alternative approach would be not to redefine the *semantics* of an ELP, but instead to state some additional conditions on the *content* of the extended logic program. The above example, for instance, would yield a perfectly acceptable outcome if the rules $\neg m \leftarrow \neg hs$ and $\neg b \leftarrow hs$ were added (which are essentially the contraposed versions of $hs \leftarrow m$ and $\neg hs \leftarrow b$). In that case, the well-founded model would be $\{r, p\}$. This approach would be quite similar to the work that Caminada and Amgoud have done in the field of formal argumentation, where similar difficulties occur (Caminada & Amgoud 2005).

Logic Programming as Argumentation

In this section, we will state some theory that allows us to link logic programming to formal argumentation. Using this theory, we will be able to apply the solution of (Caminada & Amgoud 2005) in the context of extended logic programming.

The first thing to do is to define the set of arguments and the defeat relation, given an (extended) logic program P . We choose a form of arguments that is different from (Dung 1995) and better suited to our purpose.

definition 6. Let P be an extended logic program.

- An argument A based on P is a finite tree of rules from P such that each node (of the form $c \leftarrow a_1, \dots, a_n, \text{not } b_1, \dots, \text{not } b_m$ with $n \geq 0$ and $m \geq 0$) has exactly n children, each having a different head $a_i \in \{a_1, \dots, a_n\}$. The conclusion of A ($\text{Conc}(A)$) is the head of its root.
- We say that an argument A_1 defeats an argument A_2 iff A_1 has conclusion c and A_2 has a rule containing $\text{not } c$.

We define Arguments_P as the set of arguments that can be constructed using P , and Defeat_P as the defeat relation under P . Let $\text{Args} \subseteq \text{Arguments}_P$. We define $\text{Concs}(\text{Args})$ as $\{\text{Conc}(A) \mid A \in \text{Args}\}$.

We say that argument A is a *subargument* of argument B iff A is a subtree of B . We say that argument A is a *direct subargument* of argument B iff A is a subtree of B and there does not exist an argument C such that $C \neq A$, $C \neq B$, C is a subtree of B , and A is a subtree of C .

definition 7. We say that:

- a set of arguments Args is conflict-free iff Args does not contain two arguments A and B such that A defeats B
- a set of arguments Args defends an argument A iff for each argument B that defeats A , Args contains an argument C that defeats B .

definition 8. Let Args be a set of arguments. We define $f(\text{Args})$ as $\{A \mid \text{Args} \text{ does not contain an argument that defeats } A\}$ and $F(\text{Args})$ as $f(f(\text{Args}))$.

$F(\text{Args})$ can be seen as the set of arguments that are defended by Args (Dung 1995).

lemma 1. Let P be an extended logic program and let E be the smallest fixpoint of F under P . E is conflict-free.

Proof. As E is the smallest fixpoint of F under P , it holds that (Dung 1995) $E = \cup_{i=0}^{\infty} F^i(\emptyset)$. Suppose that E is not conflict-free. As F is a monotonic function and $F^0(\emptyset) = \emptyset$, there must be some smallest i ($i \geq 0$) such that $F^i(\emptyset)$ is conflict-free but $F^{i+1}(\emptyset)$ is not conflict-free. From definition 7 it then follows that $F^{i+1}(\emptyset)$ contains two arguments A and B such that A defeats B . The fact that A defeats B and $B \in F^{i+1}(\emptyset)$ means that there is an argument $C \in F^i(\emptyset)$ that defeats A . The fact that C defeats A and $A \in F^{i+1}(\emptyset)$ means that there is an argument $D \in F^i(\emptyset)$ that defeats C . But then $F^i(\emptyset)$ would not be conflict-free. Contradiction. \square

The following property follows from definition 6 and 2.

property 1. Let S be a set of strict rules and $\mathbf{1}$ be a literal. It holds that $\mathbf{1} \in Cl(S)$ iff there exists an argument A , based on S , such that $\text{Conc}(A) = \mathbf{1}$.

The following property follows from definition 4 and 6.

property 2. Let P be an extended logic program and L be a set of literals. There exists an argument A , based on P^L , with $\text{Conc}(A) = \mathbf{1}$ iff there exists an argument B , based on P , with $\text{Conc}(B) = \mathbf{1}$, such that B does not contain a weakly negated literal $k \in L$.

The function γ is actually quite similar to the function f , as is stated in the following theorem.

theorem 1. Let L be a set of literals and $Args$ be a set of arguments. If $L = \text{Concs}(Args)$ then $\gamma(L) = \text{Concs}(f(Args))$.

Proof. We need to prove two things:

1. $\gamma(L) \subseteq \text{Concs}(f(Args))$
Let $\mathbf{1} \in \gamma(L)$. This, by definition 5, means that $\mathbf{1} \in Cl(P^L)$. From property 1 it follows that there exists an argument (A), based on P^L , with $\text{Conc}(A) = \mathbf{1}$. Then, according to property 2, there exists an argument (B), based on P , with $\text{Conc}(B) = \mathbf{1}$, such that B does not contain a weakly negated literal $k \in L$. As $L = \text{Concs}(Args)$, the argument B is not defeated by $Args$. Therefore, $B \in f(Args)$. As B has conclusion $\mathbf{1}$ it holds that $\mathbf{1} \in \text{Concs}(f(Args))$
2. $\text{Concs}(f(Args)) \subseteq \gamma(L)$
Let $\mathbf{1} \in \text{Concs}(f(Args)) \subseteq \gamma(L)$. This means that $f(Args)$ contains some argument (say B) with conclusion $\mathbf{1}$. That is, there exists an argument (B) with conclusion $\mathbf{1}$ that is not defeated by $Args$. From property 2 it then follows that there exists an argument A , based on P^L (as $L = \text{Concs}(Args)$), with $\text{Conc}(A) = \mathbf{1}$. This, by property 1, means that $\mathbf{1} \in Cl(P^L)$, which by definition 5 means that $\mathbf{1} \in \gamma(L)$. □

The following theorem states that the well-founded model of a program P coincides with the conclusions of the grounded extension (Dung 1995) of the argument-interpretation of P .

theorem 2. Let P be an extended logic program. The grounded extension GE of $\langle Arguments_P, Defeat_P \rangle$ coincides with the smallest fixpoint (WFM) of Γ . That is: $\text{concs}(GE) = WFM$.

Proof. From theorem 1 it follows that, if $L = \text{Concs}(Args)$, then $\gamma(\gamma(L)) = \text{Concs}(f(f(Args)))$, so $\Gamma(L) = \text{Concs}(F(L))$. Therefore, the smallest fixpoint of Γ is equal to the conclusions of the smallest fixpoint of F , which is the grounded extension. □

Semi-Normal Extended Logic Programs

In this section, we define some restrictions on an extended logic program. An extended logic program that satisfies these restrictions is called a *semi-normal* extended logic program (a term inspired by semi-normal default theories). We then show that a semi-normal extended logic program avoids problems like illustrated in example 1 by always having a consistent well-founded model.

definition 9. Let s_1 and s_2 be strict rules. We say that s_2 is a transposition of s_1 iff:

$$s_1 = c \leftarrow a_1, \dots, a_n \text{ and}$$

$$s_2 = \neg a_i \leftarrow a_1, \dots, a_{i-1}, \neg c, a_{i+1}, \dots, a_n \text{ for some } 1 \leq i \leq n.$$

The intuition behind transposition can be illustrated by translating a strict rule $c \leftarrow a_1, \dots, a_n$ to a material implication $c \subseteq a_1 \wedge \dots \wedge a_n$. This implication is logically equivalent to

$\neg a_i \subseteq a_1 \wedge \dots \wedge a_{i-1} \wedge \neg c \wedge a_{i+1} \wedge \dots \wedge a_n$, which is again translated to $\neg a_i \leftarrow a_1, \dots, a_{i-1}, \neg c, a_{i+1}, \dots, a_n$. Notice that, when $n = 1$, transposition coincides with classical contraposition.

definition 10. A defeasible rule is semi-normal iff it is of the form

$$c \leftarrow a_1, \dots, a_n, \text{not } b_1, \dots, \text{not } b_m, \text{not } \neg c.$$

definition 11. An extended logic program P is called semi-normal iff:

1. $strict(P)$ is consistent,
2. $strict(P)$ is closed under transposition, and
3. $defeasible(P)$ consists of semi-normal rules only

If A is an argument, then the *depth* of A is the number of nodes on the longest root-originated path in A . If A is an argument and r is a rule in A then the *depth of r in A* is the number of nodes on the shortest path from the root to a node labeled with r .

lemma 2. Let P be a semi-normal extended logic program, Ass (the assumptions) be a nonempty set of strict rules with empty antecedents $\{a_1 \leftarrow, \dots, a_n \leftarrow\}$ and A an argument with conclusion c based on $strict(P) \cup Ass$, such that A contains an assumption $a_i \leftarrow (1 \leq i \leq n)$ that does not occur in P . There exists an argument B , based on $strict(P) \cup Ass \cup \{\neg c \leftarrow\}$ such that B has a conclusion $\neg a_i$.

Proof. We prove this by induction on the depth of A .

basis Let's assume that the depth of A is 1. In that case, A consists of a single rule, which must then have an empty antecedent. Therefore, the root of A must be $c \leftarrow$. It then follows that $c = a_i$. Therefore, there exists an argument (A itself) based on $strict(P) \cup Ass \cup \{\neg c \leftarrow\}$ that has conclusion $\neg a_i$.

step Suppose the above lemma holds for all strict arguments of depth $\leq j$. We now prove that it also holds for all strict arguments of depth $j + 1$. Let A be an argument of depth $j + 1$, based on $strict(P) \cup Ass$, with conclusion c . Let $c \leftarrow \text{Conc}(A_1), \dots, \text{Conc}(A_m)$ be the root of A . Let A_i be a direct subargument of A that contains the assumption $a_i \leftarrow$. Because the set of strict rules in P is closed under transposition, there exists a rule $\neg \text{Conc}(A_i) \leftarrow \text{Conc}(A_1), \dots, \text{Conc}(A_{i-1}), \neg c, \text{Conc}(A_{i+1}), \dots, \text{Conc}(A_m)$. The fact that A_i has a depth $\leq j$ means that we can apply the induction hypothesis. That is, there exists an argument (say B'), based on $strict(P) \cup Ass \cup \{\neg \text{Conc}(A_i) \leftarrow\}$, with conclusion $\neg a_i$. Now, in B' , substitute $\neg \text{Conc}(A_i) \leftarrow$ by the subargument $\neg \text{Conc}(A_i) \leftarrow A_1, \dots, A_{i-1}, \neg c, A_{i+1}, \dots, A_m$. The resulting argument (call it B) is a strict argument, based on $strict(P) \cup Ass \cup \{\neg c \leftarrow\}$, with conclusion $\neg a_i$. □

theorem 3. Let $\langle Arguments_P, Defeat_P \rangle$ be an argumentation framework built from a semi-normal extended logic program P , and let E be the smallest fixpoint of F . It holds that $\text{Concs}(E)$ is consistent.




Figure 1: The working of theorem 3

Proof. Let E be the grounded extension of $\langle \text{Arguments}_P, \text{Defeat}_P \rangle$. Suppose the conclusions of E are not consistent. Then E contains an argument (say A) with conclusion c and an argument (say B) with conclusion $\neg c$. As $\text{strict}(P)$ is consistent, at least one of these two arguments must contain a defeasible rule. Let us, without loss of generality, assume that A contains at least one defeasible rule. Let d be a defeasible rule in A that has minimal depth. Notice that the depth of d must be at least 1, for if d were the top-rule of A , then B would defeat A and E would not be conflict-free (which conflicts with lemma 1). It now holds that every rule in A with a smaller depth than d is a strict rule (see also figure 1). Let A_i be a subargument of A that has d as its top-rule. We will now prove that there exists an argument (D') in E that defeats A_i . Let A_1, \dots, A_n be the subarguments of A that are at the same level as A_i in A . Lemma 2 tells us that with the conclusions of A_1, \dots, A_n, B it is possible to construct an argument with a conclusion that is the opposite of the conclusion of A_i . Call this argument D . Now, let D' be equal to D , but with the assumptions $\text{Conc}(A_1) \leftarrow, \dots, \text{Conc}(A_n) \leftarrow, \text{Conc}(B) \leftarrow$ substituted by the underlying arguments A_1, \dots, A_n, B . It holds that $D' \in E$ (this is because each defeater of D' is also a defeater of $A_1, \dots, A_n, B \in E$, and the fact that E is a fixpoint of F means it defends itself against this defeater, which means that $D' \in E$). D' , however, defeats A_i on d , so the fact that $D', A_i \in E$ means that E is not conflict-free, and (lemma 1) also no fixpoint of F . Contradiction. \square

theorem 4. *Let P be a semi-normal extended logic program. The smallest fixpoint WFM (the well-founded model) of Γ is consistent.*

Proof. This follows directly from theorem 2 and theorem 3. \square

Discussion

Many scholars in the field of defeasible reasoning distinguish two types of abstract rules: *strict rules* and *defeasible rules* (Pollock 1992; Nute 1994; Prakken & Sartor 1997; García & Simari 2004). A strict rule $a_1, \dots, a_n \rightarrow b$ basically means that if a_1, \dots, a_n hold, then it is *without any possible exception* also the case that b holds. A defeasible rule $a_1, \dots, a_n \Rightarrow b$ basically means that if a_1, \dots, a_n hold, then it is *usually* (or *normally*) the case that b holds.

One possible application of strict rules is to describe things that hold by definition (like ontologies). For instance, a cow is by definition a mammal and someone who is married by definition has a spouse. For this kind of rules, it appears that transposition is quite naturally applicable. If from

a_1, \dots, a_n it follows without any possible exception that b , then it also holds that from $a_1, \dots, a_{i-1}, \neg b, a_{i+1}, \dots, a_n$ it follows without any possible exception that $\neg a_i$.

In essence, one could say that the problems of example 1 are caused by the fact that two conclusions (m and b) are conflicting (as m implies hs , and b implies $\neg hs$) but the standard entailment of ELP is too weak to discover this conflict. Transposition (for strict rules) can thus be seen as a way of strengthening the entailment, so that this kinds of hidden conflicts become explicit, and therefore manageable.

Some formalisms for defeasible reasoning, like (Pollock 1992; 1995), have strict rules that coincide with classical (propositional or first order) reasoning. That is, there exists a strict rule $a_1, \dots, a_n \rightarrow b$ iff $a_1, \dots, a_n \vdash b$. In such a formalism, example 1 could be represented by the defeasible rules $r \Rightarrow m$ and $p \Rightarrow b$ and by the propositions $r, p, m \supset hs$ and $b \supset \neg hs$. Using these propositions one can then construct the strict rules $m, (m \supset hs) \rightarrow hs$ and $b, (b \supset \neg hs) \rightarrow \neg hs$, as well as the strict rules $\neg hs, (m \supset hs) \rightarrow \neg m$ and $hs, (b \supset \neg hs) \rightarrow \neg b$. These rules can be used not only to construct arguments for m and b but also to construct the much needed counterarguments deriving $\neg m$ and $\neg b$. By basing strict rules on classical entailment, Pollock is able to specify a formalism that avoids many of the difficult issues that have been plaguing the field of extended logic programming.

It is not difficult to see that transposition is a valid principle in classical logic (from $a_1, \dots, a_n \vdash b$ it follows that $a_1, \dots, a_{i-1}, \neg b, a_{i+1}, \dots, a_n \vdash \neg a_i$). In general, the set of strict rules generated by classical entailment satisfies many interesting properties. With transposition we have isolated the specific property of classical logic that is actually needed to avoid problems like illustrated by example 1. We simply apply the part of classical logic that we actually need, without having to go through the complexities of having to implement a full-blown classical logic theorem prover to generate the set of strict rules, as is for instance done in (Pollock 1995). The main cost of our approach is in generating the transpositions of the strict rules. For each strict rule, n transpositions are generated, where n is the number of literals in the body of the rule.

As for the defeasible rules, Pollock distinguishes two ways in which these can be argued against: rebutting and undercutting (Pollock 1992; 1995). Rebutting essentially means deriving the opposite consequent (head) of the rule, whereas undercutting basically means that there is some additional information under which the antecedent (body) of the rule is no longer a reason for the consequent (head) of the rule. For instance, suppose that we have the defeasible rule that an object that looks red usually is red. A rebutter would be that the object is not red, because it is known to be blue. An undercutter would be that the object is illuminated by a red light. This is not a reason for it *not* being red, but merely means that the fact that it looks red can no longer be regarded as a valid reason for it actually being red. Thus, rebutting attacks the consequent (head) of a rule, whereas undercutting attacks merely the connection between the antecedent (body) and the consequent (head) of a rule. Pollock claims, based on his philosophical work regarding episte-

mology, all forms of defeat can be reduced to rebutting and undercutting (Pollock 1992). This observation is important, as both of these forms of defeat can be modeled using semi-normal defeasible rules in extended logic programs.

Many problems in logic programming are caused by specific logic programs containing anomalous information (a rule like $a \leftarrow \text{not } a$ could for instance cause the absence of stable models). If one wants to apply standard and relatively straightforward semantics then one needs to make sure that a logic program does not contain such anomalies. If one provides anomalous input (like stating that a married person always has a spouse, without stating that someone who does not have a spouse is not married, using a formalism (ELP) that is not powerful enough to make this inference itself) then one should not be surprised that the outcome (the well-founded model) is anomalous as well. For reasons described above, we think that that the concept of semi-normal extended logic programs can serve as a quite natural and reasonable restriction of which programs can be regarded to be free of anomalies.

Quality Postulates

One way to evaluate the different approaches for providing a suitable semantics for ELP is by providing quality postulates (Caminada & Amgoud 2005). The idea is to state a number of general properties that should be satisfied by any formalism for defeasible reasoning, including ELP. In (Caminada & Amgoud 2005; ASPIC-consortium 2005) the following quality postulates have been stated:

- direct consistency. Let P be an extended logic program such that $\text{strict}(P)$ is consistent, and let M be a model of P (under some specified semantics). It must hold that M is consistent.
- closedness. Let P be an extended logic program and let M be a model under P (under some specified semantics). It must hold that $\text{Cl}(\text{strict}(P) \cup M) = M$.
- indirect consistency. Let P be an extended logic program such that $\text{strict}(P)$ is consistent, and let M be a model of P (under some specified semantics). It must hold that $\text{Cl}(\text{strict}(P) \cup M)$ is consistent.

The quality postulate of direct consistency is quite straightforward and is satisfied by most formalisms that we know of. The quality postulate of closedness basically states that, as far as the strict rules are concerned, the model is “complete”. The quality postulate of indirect consistency does by itself not require that the model is closed under the strict rules, but instead requires the more modest property that if one would compute the closure of the model under the strict rules, the result would at least not contain any inconsistencies.

The above three quality postulates are not completely independent. Indirect consistency, for instance, implies direct consistency. Similarly, closedness and direct consistency imply indirect consistency.

To illustrate the value of the above three quality postulates, consider a person who knows a set of strict and defeasible rules, encodes these as a semi-normal extended logic

program and then examines a model generated by an ELP inference engine. If the ELP inference engine would (in example 1) provide a model containing m but not containing hs (thus violating closedness) then the user may conclude that the ELP inference engine apparently “forgot” something. Worse yet, if the ELP inference engine provides a model containing m and b (thus violating indirect consistency) then the user may reason like: “My inference engine says that m , and I know that from m it always follows that hs , therefore hs . My inference engine also says that b and I know that from b it always follows that $\neg hs$, therefore $\neg hs$.” It is our view that, from an agent perspective, a formalism that does not satisfy indirect consistency cannot be used to generate the beliefs of an agent, as we think that an agent should never run into inconsistencies once it starts to do additional reasoning on its own beliefs.

Although ELP-models should ideally be closed under the strict rules of P , they should not necessarily be closed under the defeasible rules of P . If a is given and there exists a rule “if a then normally b ”, then one cannot simply derive b since the situation may not be normal. The quality postulate of closedness is thus only relevant with respect to strict rules.

A fourth quality postulate that has, as far as we know, not been published earlier is that of crash-resistancy:

- crash-resistancy. There should not exist an extended logic program P , with $\text{strict}(P)$ consistent, such that for any extended logic program P' , with $\text{strict}(P')$ consistent, that does not share any atoms with P , it holds that P has the same models (under some specific semantics) as $P \cup P'$.

Crash-resistancy basically states that it should not be possible for an extended logic program to contain some pieces of information (P) that makes other totally unrelated pieces of information (P') totally irrelevant when added.

The above four quality postulates are violated by various approaches that aim to provide extended logic programs with a suitable semantics. Indirect consistency, for instance, is problematic in approaches that are based on paraconsistent reasoning. When the approach of, for instance, (Sakama 1992) is applied to example 1, it produces a well-founded model $\{\{r, p, m, b, hs, \neg hs\}, \{\neg r, \neg p, \neg m, \neg b\}\}$. Using Ginsberg’s 7-valued default bilattice, this means that only r, p, m and b (but not hs or $\neg hs$) are considered true, thus violating closedness and indirect consistency.

Brewka’s approach to well-founded semantics (Brewka 1996), on the other hand, violates direct consistency as well as crash-resistancy. In example 1, $\text{strict}(P)$ is consistent, but Brewka’s approach nevertheless yields the inconsistent set Lit , which violates direct consistency. As the outcome of Lit is obtained even when one adds syntactically totally unrelated rules to P , crash-resistancy is violated as well.

The quality postulate of crash-resistancy is violated by the stable model semantics of answer set programming, where a simple rule like $a \leftarrow \text{not } a$ yields no stable models at all, regardless of what additional (unrelated) information is contained in the logic program. A common opinion in the ELP-research community is that programs that have no stable models are by definition anomalous and unnatural. We

hereby would like to argue against this view. Consider a situation in where persons are usually believed in what they say, unless information of the contrary is available (rebut) or the person is known to be unreliable (undercut). Now consider the following three persons, who give the following statements:

- Bert: “Ernie is unreliable.”
- Ernie: “Elmo is unreliable.”
- Elmo: “Bert is unreliable.”

This would correspond with the following extended logic program:

- `bert_says_u_ernie` ←
- `u_ernie` ← `bert_says_u_ernie`, `not` `¬u_ernie`, `not` `u_bert`
- `ernie_says_u_elmo` ←
- `u_elmo` ← `ernie_says_u_elmo`, `not` `¬u_elmo`, `not` `u_ernie`
- `elmo_says_u_bert` ←
- `u_bert` ← `elmo_says_u_bert`, `not` `¬u_bert`, `not` `u_elmo`

It is perfectly possible for a situation to occur in which three persons, sitting in a circle, claim their direct neighbour is unreliable. How this conflict should be dealt with is an issue open for discussion, but it should at least not cause the hearer to enter a state of total ignorance in which also all other entailment is completely blocked. It is our opinion, also for reasons described in (Dung 1995) that the problems of stable model semantics are very often caused by the nature of the semantics itself, and not by an “anomalous” extended logic program.

Summary and Conclusions

One of the advantages of the approach as sketched in the current paper is that it satisfies each of the quality postulates direct consistency, indirect consistency, closedness and crash-resistancy. Furthermore, it does so without the need of an advanced semantics that is complex and potentially difficult to understand. Although the approach only works for the somewhat restricted notion of semi-normal extended logic programs, we believe that these restrictions are in essence quite natural and can be given a decent philosophical justification.

References

- ASPIC-consortium. 2005. Deliverable D2.5: Draft formal semantics for ASPIC system.
- Brewka, G. 1996. Well-founded semantics for extended logic programs with dynamic preferences. *J. Artif. Intell. Res. (JAIR)* 4:19–36.
- Caminada, M., and Amgoud, L. 2005. An axiomatic account of formal argumentation. In *Proceedings of the AAAI-2005*, 608–613.
- Dix, J. 1995a. A classification theory of semantics of normal logic programs: I. strong properties. *Fundam. Inform.* 22(3):227–255.

- Dix, J. 1995b. A classification theory of semantics of normal logic programs: II. weak properties. *Fundam. Inform.* 22(3):257–288.
- Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n -person games. *Artificial Intelligence* 77:321–357.
- García, A., and Simari, G. 2004. Defeasible logic programming: an argumentative approach. *Theory and Practice of Logic Programming* 4(1):95–138.
- Gelfond, M., and Lifschitz, V. 1991. Classical negation in logic programs and disjunctive databases. *New Generation Computing* 9(3/4):365–385.
- Nute, D. 1994. Defeasible logic. In Gabbay, D.; Hogger, C. J.; and Robinson, J. A., eds., *Handbook of Logic in Artificial Intelligence and Logic Programming*. Oxford: Clarendon Press. 253–395.
- Pollock, J. L. 1992. How to reason defeasibly. *Artificial Intelligence* 57:1–42.
- Pollock, J. L. 1995. *Cognitive Carpentry. A Blueprint for How to Build a Person*. Cambridge, MA: MIT Press.
- Prakken, H., and Sartor, G. 1997. Argument-based extended logic programming with defeasible priorities. *Journal of Applied Non-Classical Logics* 7:25–75.
- Sakama, C. 1992. Extended well-founded semantics for paraconsistent logic programs. In *FGCS*, 592–599.
- van Gelder, A.; Ross, K. A.; and Schlipf, J. S. 1991. The well-founded semantics for general logic programs. *J. ACM* 38(3):620–650.