


**I N F S Y S  
R E S E A R C H  
R E P O R T**



**INSTITUT FÜR INFORMATIONSSYSTEME  
ARBEITSBEREICH WISSENSBASIERTE SYSTEME**

**EVALUATING EPISTEMIC NEGATION  
IN ANSWER SET PROGRAMMING**

**YI-DONG SHEN   THOMAS EITER**

**INFSYS RESEARCH REPORT 1843-15-04**

**AUGUST 2015**

Institut für Informationssysteme  
AB Wissensbasierte Systeme  
Technische Universität Wien  
Favoritenstraße 9-11  
A-1040 Wien, Austria  
Tel: +43-1-58801-18405  
Fax: +43-1-58801-18493  
sek@kr.tuwien.ac.at  
www.kr.tuwien.ac.at



**kbs**  
*Knowledge-Based  
Systems Group*

## EVALUATING EPISTEMIC NEGATION IN ANSWER SET PROGRAMMING

Yi-Dong Shen<sup>1</sup>      Thomas Eiter<sup>2</sup>

**Abstract.** Epistemic negation along with default negation plays a key role in knowledge representation and nonmonotonic reasoning. However, the existing epistemic approaches such as those by Gelfond (1991), Truszczyński (2011) and Kahl [18] (2014) behave not satisfactorily in that they suffer from either epistemic circular justifications or the multiple world view problem and thus produce undesired results for some logic programs. In this paper we propose a new approach to handling epistemic negation which is free of epistemic circular justifications and the multiple world view problem, and thus offers a solution to evaluate epistemic specifications which were introduced by Gelfond (AAAI 1991) over two decades ago. We consider general logic programs consisting of rules of the form  $H \leftarrow B$ , where  $H$  and  $B$  are arbitrary first-order formulas possibly containing epistemic negation, and define an extended FLP answer set semantics for general logic programs, called *epistemic FLP semantics*, by introducing a novel program transformation and a new definition of world views. The proposed approach can readily be adapted to any other existing answer set semantics for extension with epistemic negation, such as those by Pearce (1996, 2006), Pearce et al. (2007), Truszczyński (2010), Bartholomew et al. (2011), Ferraris et al. (2011), and Shen et al. (2014). We also consider the computational complexity of epistemic FLP semantics and show that for a propositional program  $\Pi$  with epistemic negation, deciding whether  $\Pi$  has epistemic FLP answer sets is  $\Sigma_3^p$ -complete and deciding whether a propositional formula  $F$  is true in  $\Pi$  under epistemic FLP semantics is  $\Sigma_4^p$ -complete in general.

---

<sup>1</sup>State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences, Beijing 100190, China; email: ydshen@ios.ac.cn.

<sup>2</sup>Institut für Informationssysteme, Technische Universität Wien, Favoritenstraße 9-11, A-1040 Vienna, Austria; email: {eiter,fink,tkren,redl}@kr.tuwien.ac.at.

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Preliminaries</b>	<b>7</b>
2.1	A First-Order Logic Language . . . . .	7
2.2	Logic Programs with Epistemic Negation . . . . .	8
2.2.1	Syntax . . . . .	8
2.2.2	Grounding . . . . .	9
2.2.3	Satisfaction and Models . . . . .	9
<b>3</b>	<b>Program Transformation and Epistemic Reducts</b>	<b>10</b>
<b>4</b>	<b>FLP Answer Set Semantics with Epistemic Negation</b>	<b>11</b>
<b>5</b>	<b>Computational Complexity</b>	<b>17</b>
<b>6</b>	<b>Related Work</b>	<b>18</b>
<b>7</b>	<b>Summary and Future Work</b>	<b>20</b>
<b>A</b>	<b>Proofs</b>	<b>23</b>
<b>B</b>	<b>Programs Matching Normal Epistemic Specifications</b>	<b>28</b>

## 1 Introduction

Answer set programming (ASP) is a major logic programming paradigm rooted in knowledge representation and reasoning (KR) for modeling and solving knowledge-intensive search and optimization problems such as product configuration and planning [4]. In ASP, the semantics of a logic program is given by a set of intended models, called stable models or answer sets [17, 14]. Such answer sets can be defined in different ways; Lifschitz [21] listed 13 of them in the literature. These semantics agree for normal logic programs, but show discrepancies for more general logic programs such as logic programs with aggregates [29, 9], with external sources such as description logic programs (dl-programs) [8], and with propositional or first-order formulas [24, 26, 31, 3, 11]. Most recently Shen et al. [28] introduced a new one called the well-justified FLP answer set semantics, which is fundamentally distinct from other existing answer set semantics in that it is free of circular justifications, i.e., every answer set of a general logic program has a level mapping. This semantics has been implemented over the well-known ASP reasoner DLVHEX.<sup>1</sup>

Negation is a key mechanism in ASP for reasoning with incomplete knowledge. There are two major types of negation, *default negation* and *epistemic negation*. A third, called *strong negation*, also appears in the literature; when default negation is available, strong negation is easily compiled away using new predicate symbols [14] and thus it can be omitted. By abuse of notation, in this paper we use  $\neg$ , **not** and  $\sim$  to denote the three negation operators, respectively.<sup>2</sup> For a formula  $F$ , the default negation  $\neg F$  of formula  $F$  expresses that there is no *justification* for adopting  $F$  in an answer set and thus  $F$  can be assumed false by default in the answer set; in contrast, the epistemic negation **not**  $F$  of  $F$  expresses that there is no evidence *proving* that  $F$  is true, i.e.,  $F$  is false in some answer set. *Justification* in ASP is a concept defined over individual answer sets, while *provability* is a meta-level concept defined over a collection of answer sets. This means the two types of negation are orthogonal operations, where default negation works *locally* on each individual answer set, and epistemic negation works *globally* at a meta level over all answer sets.

With both default and epistemic negation, ASP is enabled to reason with different incomplete knowledge. For example, we can use the rule

$$\text{innocent}(X) \leftarrow \text{not guilty}(X)$$

to concisely express the *presumption of innocence*, which states that one is presumed innocent if there is no evidence proving s/he is guilty. We can also use rules of the form

$$\neg p(X) \leftarrow \text{not } p(X)$$

to explicitly state Reiter's *closed-world assumption* (CWA) [27], i.e., if there is no evidence proving  $p(X)$  is true we jump to the conclusion that  $p(X)$  is false.

However, observe that most of the existing answer set semantics, such as [14, 24, 26, 31, 3, 9, 11, 28], only support default negation and they do not allow for epistemic negation.

**Epistemic negation and specifications.** In fact, the need for epistemic negation was long recognized in ASP by Gelfond [15] and recently revisited in [16, 32, 18, 6]. In particular, [15] showed that formalization of CWA using default and strong negations with rules of the form

<sup>1</sup>[www.kr.tuwien.ac.at/research/systems/dlvhex](http://www.kr.tuwien.ac.at/research/systems/dlvhex)

<sup>2</sup>In many texts, *not* and  $\neg$  are used to denote the default and strong negation operators, respectively.

$$\sim p(X) \leftarrow \neg p(X)$$

as presented in [14], is problematic. He then proposed to address the problem using two epistemic modal operators  $\mathbf{K}$  and  $\mathbf{M}$ . Informally, for a formula  $F$ ,  $\mathbf{K}F$  expresses that  $F$  is true in every answer set, and  $\mathbf{M}F$  expresses that  $F$  is true in some answer set. Note that  $\mathbf{M}F$  can be viewed as shorthand for  $\neg\mathbf{K}\neg F$ .<sup>3</sup>

In the sequel, by an *object literal* we refer to an atom  $A$  or its strong negation  $\sim A$ ; a *default negated literal* is of the form  $\neg L$ , and a *modal literal* is of the form  $\mathbf{K}L$ ,  $\neg\mathbf{K}L$ ,  $\mathbf{M}L$  or  $\neg\mathbf{M}L$ , where  $L$  is an object literal.

Gelfond [15] considered disjunctive logic programs with modal literals, called *epistemic specifications*, which consist of rules of the form

$$L_1 \vee \cdots \vee L_m \leftarrow G_1 \wedge \cdots \wedge G_n \quad (1)$$

where each  $L$  is an object literal and each  $G$  is an object literal, a default negated literal, or a modal literal. Given a collection  $\mathcal{A}$  of interpretations as an *assumption*, a logic program  $\Pi$  is transformed into a *modal reduct*  $\Pi^{\mathcal{A}}$  w.r.t. the assumption  $\mathcal{A}$  by first removing all rules with a modal literal  $G$  that is not true in  $\mathcal{A}$ , then removing the remaining modal literals. The assumption  $\mathcal{A}$  is defined to be a *world view* of  $\Pi$  if it coincides with the collection of answer sets of  $\Pi^{\mathcal{A}}$  under [14].

**Epistemic circular justification problem.** More recently, Gelfond [16] addressed the problem that applying the above approach to handle modal literals may produce unintuitive results. For example, consider a logic program  $\Pi = \{p \leftarrow \mathbf{K}p\}$ . The rule expresses that for any collection  $\mathcal{A}$  of answer sets of  $\Pi$  and any  $I \in \mathcal{A}$ , if  $p$  is true in all answer sets in  $\mathcal{A}$ , then  $p$  is true in  $I$ . This amounts to saying that if  $p$  is true in all answer sets, then  $p$  is always true (in particular in all answer sets). Obviously, this rule is not informative and does not contribute to constructively building any answer set; thus it can be eliminated from  $\Pi$ , leading to  $\Pi = \emptyset$ . As a result,  $\Pi$  is expected to have a unique answer set  $\emptyset$ . However,  $\{p\}$  would be an answer set of  $\Pi$  when applying the approach of [15]. To illustrate, consider an assumption  $\mathcal{A} = \{\{p\}\}$ , i.e.,  $p$  is assumed to be true in all interpretations in  $\mathcal{A}$ . Then,  $\mathbf{K}p$  is true in  $\mathcal{A}$  and we obtain the modal reduct  $\Pi^{\mathcal{A}} = \{p\}$ . This reduct has a unique answer set  $\{p\}$ , which coincides with the assumption  $\mathcal{A}$ . Thus  $\mathcal{A}$  is a world view of  $\Pi$  under [15]. Observe that this world view has an epistemic circular justification:

$$\exists_{I \in \mathcal{A}} p \in I \leftarrow \mathbf{K}p \leftarrow \forall_{I \in \mathcal{A}} p \in I$$

i.e.,  $p$  being true in an interpretation  $I = \{p\}$  of the world view  $\mathcal{A}$  is due to  $\mathbf{K}p$  being treated true in the program transformation for the modal reduct  $\Pi^{\mathcal{A}}$  (via the rule  $p \leftarrow \mathbf{K}p$ ), which in turn is due to  $p$  being true in all interpretations of the assumption  $\mathcal{A}$ .

In general, a world view  $\mathcal{A}$  is said to have an *epistemic circular justification* if some object literal  $L$  being true in some interpretation  $I \in \mathcal{A}$  is due to  $\mathbf{K}L$  (or its equivalent modal literals expressing that  $L$  is true in every interpretation  $J \in \mathcal{A}$ ) being treated true in the program transformation for the modal reduct of  $\Pi$  w.r.t.  $\mathcal{A}$ .

<sup>3</sup>Note that  $\neg\mathbf{K}F$  and  $\sim\mathbf{K}F$  are semantically equivalent. In [15],  $\mathbf{M}F$  is shorthand for  $\sim\mathbf{K}\sim F$ , while in [16], it is shorthand for  $\sim\mathbf{K}\neg F$ , which is semantically equivalent to  $\neg\mathbf{K}\neg F$ .

$G$	$G$ is true in $\mathcal{A}$	Otherwise
$\mathbf{K}L$	replace it with $L$	delete the rule
$\neg\mathbf{K}L$	replace it with $\top$	replace it with $\neg L$
$\mathbf{M}L$	replace it with $\top$	replace it with $\neg\neg L$
$\neg\mathbf{M}L$	replace it with $\neg L$	delete the rule

Table 1: The program transformation of [18], where  $G$  is a modal literal in the body of a rule. Note that  $\neg\neg L$  is a nested expression defined in [19], which is not equivalent to  $L$  as in classical logic.

To remedy the epistemic circular justification problem, Gelfond [16] revised the program transformation such that a modal reduct  $\Pi^{\mathcal{A}}$  is obtained from  $\Pi$  by first removing all rules of form (1) with a modal literal  $G$  that is not true in  $\mathcal{A}$ , then removing all modal literals  $\neg\mathbf{K}L$  and  $\mathbf{M}L$ , and finally replacing all modal literals  $\mathbf{K}L$  by  $L$  and  $\neg\mathbf{M}L$  by  $\neg L$ .

It is easy to check that the logic program  $\Pi = \{p \leftarrow \mathbf{K}p\}$  has a unique world view  $\{\emptyset\}$  when applying the revised program transformation. Unfortunately, the epistemic circular justification problem persists in other logic programs, such as  $\Pi = \{q \leftarrow \neg\mathbf{K}p, p \leftarrow \neg q\}$ . Consider an assumption  $\mathcal{A} = \{\{p\}\}$ . Since  $\mathbf{K}p$  is true in  $\mathcal{A}$ , the modal literal  $\neg\mathbf{K}p$  is not true in  $\mathcal{A}$  and thus the first rule is removed, yielding the modal reduct  $\Pi^{\mathcal{A}} = \{p \leftarrow \neg q\}$ . This reduct has a unique answer set  $\{p\}$ , which coincides with  $\mathcal{A}$ , hence  $\mathcal{A}$  is a world view of  $\Pi$ . Note that this world view has also an epistemic circular justification:

$$\exists_{I \in \mathcal{A}} p \in I \Leftarrow \neg q \Leftarrow \neg\mathbf{K}p \Leftarrow \forall_{I \in \mathcal{A}} p \in I$$

i.e.,  $p$  being true in an interpretation  $I = \{p\}$  of the world view  $\mathcal{A}$  is (via the rule  $p \leftarrow \neg q$ ) due to  $q$  being false in  $I$ , which in turn (via the rule  $q \leftarrow \neg\mathbf{K}p$ ) is due to  $\mathbf{K}p$  being treated true and thus  $\neg\mathbf{K}p$  treated false in the program transformation, which is due to  $p$  being true in all interpretations of the assumption  $\mathcal{A}$ .

**Multiple world view problem.** In addition to epistemic circular justifications, the approach of [16, 15] also suffers from the *multiple world view problem* caused by so called *M-cycles*. Consider the logic program  $\Pi = \{p \leftarrow \mathbf{M}p\}$ , which expresses that for any world view  $\mathcal{A}$  and any  $I \in \mathcal{A}$ , if  $p$  is true in some answer set in  $\mathcal{A}$ , then  $p$  is true in  $I$ . This amounts to saying that if  $p$  is true in some answer set, then  $p$  is always true, in particular in every answer set. Note that this statement for  $p$  is not circularly justified, and thus  $\{\{p\}\}$  is expected to be the only world view of  $\Pi$  (see Example 1 for more explanation). However, under the approach of [16, 15] this program has two world views,  $\{\{p\}\}$  and  $\{\emptyset\}$ . Due to this the approach is said to have the multiple world view problem [18].

**Recent advance.** Most recently, [18] further studied both epistemic circular justifications and the multiple world view problem with the approach of [16, 15] and proposed a new program transformation by appealing to *nested expressions* of [19]. Let  $\Pi$  be a logic program with rules of form (1) and  $\mathcal{A}$  a collection of interpretations as an assumption. The *Kahl modal reduct*  $\Pi^{\mathcal{A}}$  of  $\Pi$  w.r.t.  $\mathcal{A}$  is a *nested logic program* (i.e., a logic program with nested expressions), which is obtained from  $\Pi$  by replacing all modal literals or deleting rules according to Table 1. The assumption  $\mathcal{A}$  is defined

to be a *Kahl world view* of  $\Pi$  if it coincides with the collection of answer sets of the nested logic program  $\Pi^A$  under the semantics of [19].

Kahl [18] listed sixty-two interesting logic programs with modal literals and illustrated the approach with these programs. As a typical example, given an assumption  $\mathcal{A} = \{\emptyset\}$ , the Kahl modal reduct of the logic program  $\Pi = \{p \leftarrow \mathbf{M}p\}$  w.r.t.  $\mathcal{A}$  is  $\Pi^A = \{p \leftarrow \neg\neg p\}$ , which has two answer sets  $\emptyset$  and  $\{p\}$  under the semantics of [19]. Thus  $\{\emptyset\}$  is not a Kahl world view.

However, our careful study reveals that the approach of [18] may also produce undesired answer sets for some logic programs, even including a few of his sixty-two example logic programs. To sum up, we observe the following three critical shortcomings of the approach:

(1) The definition of the Kahl program transformation/modal reduct looks a bit ad hoc, and the variety of replacements for modal literals (see Table 1) lacks a deeper discussion or justification.<sup>4</sup>

(2) It is undesired to transform a logic program into a reduct containing nested expressions. As shown in [28], the existing semantics for nested expressions, such as [19, 12, 11], suffer from circular justifications. For example, for the logic program  $\Pi = \{p \leftarrow \neg\neg p\}$ ,  $I = \{p\}$  is an answer set under these semantics. Observe that this answer set has a circular justification via the self-supporting loop  $p \leftarrow \neg\neg p \leftarrow p$ , i.e.,  $p$  being true in  $I$  is due to  $I$  satisfying  $\neg\neg p$  (via the rule  $p \leftarrow \neg\neg p$ ), which in turn is due to  $p$  being true in  $I$ .

For a logic program with rules of form (1), it is desirable to transform it to a regular disjunctive logic program so that the standard answer set semantics of [17, 14] can be applied.

(3) The approaches of [15, 16] are said to suffer from the multiple world view problem because they yield for the logic program  $\Pi = \{p \leftarrow \mathbf{M}p\}$  two world views, viz.  $\mathcal{A}_1 = \{\{p\}\}$  and  $\mathcal{A}_2 = \{\emptyset\}$ ; however, to the best of our knowledge there has been no deeper discussion or justification in the literature for why  $\mathcal{A}_1$  is the right world view of this program and  $\mathcal{A}_2$  is not. In fact, there has been no formal definition of the multiple world view problem in the literature; cf. [16, 32, 18, 6]. Moreover, we observe that the approach of [18] also suffers from the multiple world view problem and produces undesired answer sets for some logic programs, as illustrated in the following example.

**Example 1** Consider again the logic program  $\Pi = \{p \leftarrow \mathbf{M}p\}$ . Intuitively, for this program it is expected to find all world views  $\mathcal{A}$  such that  $\mathbf{M}p$  is true in  $\mathcal{A}$  and every  $I \in \mathcal{A}$  is an answer set of  $\Pi$  after  $\mathbf{M}p$  is replaced by  $\top$ . We have indeed one such world view  $\mathcal{A} = \{\{p\}\}$ . Note that  $\{\emptyset\}$  should arguably not be a world view for this program, as it does not satisfy the above intuitive expectation. Applying the approach of [18] to this program will yield a unique world view  $\mathcal{A} = \{\{p\}\}$ , as expected.

Now consider another very similar logic program, which is borrowed from Example 29 in Appendix D of [18]:

$$\begin{array}{ll} \Pi : & p \leftarrow \mathbf{M}q \wedge \neg q & r_1 \\ & q \leftarrow \mathbf{M}p \wedge \neg p & r_2 \end{array}$$

<sup>4</sup>In fact, no existing approaches to epistemic specifications, such as [15, 16, 32, 18], have ever provided a deeper discussion or justification for the replacements of modal literals in their program transformations.

Following the same intuition as  $\{p \leftarrow \mathbf{M}p\}$ , for this program it is expected to find all world views  $\mathcal{A}$  such that both  $\mathbf{M}q$  and  $\mathbf{M}p$  are true in  $\mathcal{A}$  and every  $I \in \mathcal{A}$  is an answer set of  $\Pi$  after  $\mathbf{M}q$  and  $\mathbf{M}p$  are replaced by  $\top$ . It is easy to check that  $\mathcal{A} = \{\{p\}, \{q\}\}$  is such a world view that satisfies the intuitive expectation, but  $\{\emptyset\}$  does not satisfy the expectation and thus should not be a world view. However, applying the approach of [18] will produce both as the world views of this program, contradicting the expectation. This example suggests that the approach of [18] also suffers from the multiple world view problem.

**Our contributions.** The goal of this paper is to address the above problems and provide a better solution to epistemic negation as well as epistemic specifications of [15]. Our main contributions are summarized as follows:

(1) We use modal operator **not** to directly express epistemic negation and define general logic programs consisting of rules of the form  $H \leftarrow B$ , where  $H$  and  $B$  are arbitrary first-order formulas possibly containing epistemic negation. Modal formulas  $\mathbf{K}F$  and  $\mathbf{M}F$  are viewed as shorthand for  $\neg \mathbf{not} F$  and  $\mathbf{not} \neg F$ , respectively, and thus epistemic specifications of [15] are a special class of general logic programs.

(2) We propose to apply epistemic negation to minimize the knowledge in world views of a general logic program  $\Pi$ , i.e., we apply epistemic negation to formulas w.r.t. a world view and assume **not**  $F$  in  $\Pi$  to be true in the world view whenever possible; this is analogous to applying default negation to minimize the knowledge in answer sets, i.e., one applies default negation to formulas w.r.t. an answer set and assumes  $\neg F$  in  $\Pi$  to be true in the answer set whenever possible (CWA or minimal models). To this end, we introduce a novel and very simple program transformation based on epistemic negation and present a new definition of world views. Given a subset  $\Phi$  of the epistemic negations **not**  $F$  in  $\Pi$ , called a *guess*, we transform  $\Pi$  into an *epistemic reduct* based on  $\Phi$ , denoted  $\Pi^\Phi$ , by replacing every **not**  $F$  with  $\top$  if it is in  $\Phi$ , and with  $\neg F$ , otherwise. Let  $\mathcal{A}$  be the set of all answer sets of  $\Pi^\Phi$ .  $\mathcal{A}$  is called a *candidate world view* w.r.t.  $\Phi$  if it agrees with  $\Phi$  in the sense that every **not**  $F$  in  $\Pi$  is true in  $\mathcal{A}$  if it is in  $\Phi$ , and false, otherwise. A candidate world view  $\mathcal{A}$  w.r.t. a guess  $\Phi$  is defined to be a *world view* under our approach if  $\Phi$  is maximal (i.e., there is no candidate world view w.r.t. any guess  $\Phi' \supset \Phi$ ). Note that it is by applying a *maximal guess*  $\Phi$  that we minimize the knowledge in the world view w.r.t.  $\Phi$ . Obviously, such definitions of program transformations and world views fundamentally differ from those of [15, 16, 18].

We will give a formal definition for the multiple world view problem based on epistemic negation and show that world views under our approach are free of epistemic circular justifications and the multiple world view problem.

(3) The proposed approach can be used to extend any existing answer set semantics with epistemic negation, such as [24, 26, 31, 3, 9, 11, 28]. For illustration, we extend here as a show case the well-known *FLP semantics* of [9], yielding a new semantics called *epistemic FLP semantics*.

(4) As satisfiability of an arbitrary general logic program is undecidable, we address the computational complexity of epistemic FLP semantics for propositional programs. In particular, we show that deciding whether a propositional program  $\Pi$  has epistemic FLP answer sets is  $\Sigma_3^P$ -complete.



Furthermore, we show that query evaluation, i.e., deciding whether a propositional formula is true in every epistemic FLP answer set of some world view of  $\Pi$ , is  $\Sigma_4^p$ -complete in general; important fragments, e.g. programs that match normal epistemic specifications as in [15], have lower complexity.

The paper is organized as follows. In Section 2, we introduce a first-order logic language and define logic programs with first-order formulas and epistemic negation. In Section 3, we present a novel program transformation and define epistemic reducts. In Section 4 we extend FLP answer set semantics with epistemic negation, and in Section 5 we study in depth the computational complexity of the extended semantics. In Section 6, we discuss related work, and in Section 7 we conclude and mention some ongoing work.

In order not to distract from the flow of reading, proofs of theorems are in the appendix.

## 2 Preliminaries

In this section, we introduce a first-order logic language and define logic programs with first-order formulas and epistemic negation.

### 2.1 A First-Order Logic Language

We follow the notation in [28] and define a first-order logic language  $\mathcal{L}_\Sigma$  with equality over a *signature*  $\Sigma = (\mathcal{P}, \mathcal{F})$ , where  $\mathcal{P}$  and  $\mathcal{F}$  are countable sets of *predicate* and *function* symbols, respectively;  $\mathcal{C} \subseteq \mathcal{F}$  denotes the set of 0-ary function symbols, which are called *constants*. *Variables*, *terms*, *atoms* and *literals* are defined as usual. We denote variables with strings starting with  $X$ ,  $Y$  or  $Z$ .

*First-order formulas* (briefly *formulas*) are constructed as usual from atoms using connectives  $\neg$ ,  $\wedge$ ,  $\vee$ ,  $\supset$ ,  $\top$ ,  $\perp$ ,  $\exists$  and  $\forall$ , where  $\top$  and  $\perp$  are two 0-place logical connectives expressing *true* and *false*, respectively. Formulas are *closed* if they contain no free variables, i.e., each variable occurrence is in the scope of some quantifier. A *first-order theory* (or *theory*) is a set of closed formulas. Terms, atoms and formulas are *ground* if they have no variables, and *propositional* if they contain no variables, no function symbols except constants, and no equalities. By  $\mathcal{N}_\Sigma$  we denote the set of all ground terms of  $\Sigma$ , and by  $\mathcal{H}_\Sigma$  the set of all ground atoms.

In this paper we consider *SNA interpretations*, i.e., interpretations which employ the well-known *standard names assumption* (SNA) [5, 23]. An SNA interpretation (or interpretation for short)  $I$  of  $\mathcal{L}_\Sigma$  is a subset of  $\mathcal{H}_\Sigma$  such that for any ground atom  $A$ ,  $I$  *satisfies*  $A$  if  $A \in I$ , and  $I$  *satisfies*  $\neg A$  if  $A \notin I$ . The notion of *satisfaction/models* of a formula/theory in  $I$  is defined as usual. A theory  $T$  is *consistent* or *satisfiable* if  $T$  has a model.  $T$  *entails* a closed formula  $F$ , denoted  $T \models F$ , if all models of  $T$  are models of  $F$ .  $F$  is *true* (resp. *false*) in an interpretation  $I$  if  $I$  *satisfies* (resp. *does not satisfy*)  $F$ .

## 2.2 Logic Programs with Epistemic Negation

Based on the first-order logic language  $\mathcal{L}_\Sigma$  defined above, we introduce the syntax of a logic program with epistemic negation and define its grounding, satisfaction and models.

### 2.2.1 Syntax

As in [28] and early AI literature, default negation will be expressed by giving the connective  $\neg$  a special meaning; i.e., we use  $\neg$  to denote the default negation operator. We further extend the language  $\mathcal{L}_\Sigma$  to include the epistemic negation operator **not** and the rule operator  $\leftarrow$ .

*Epistemic formulas* are constructed from atoms using the connectives  $\neg, \wedge, \vee, \supset, \top, \perp, \exists, \forall$  together with the operator **not** in the same way as first-order formulas. An *epistemic negation* is an epistemic formula of the form **not**  $F$ , where  $F$  is an epistemic formula; it is *non-nested* if  $F$  contains no epistemic negation, and *nested*, otherwise. An epistemic formula is *closed* if it contains no free variables.

Let  $E$  be an epistemic formula and  $E'$  be  $E$  with every free variable replaced by a constant (this process is called *grounding*).  $E$  is *instance-closed* (w.r.t. epistemic negation) if for every epistemic negation **not**  $F$  in  $E'$ ,  $F$  itself is a closed epistemic formula. For instance,  $person(X) \wedge \mathbf{not} \mathit{guilty}(X)$  is an instance-closed epistemic formula, as when the free variable  $X$  is replaced by a constant, the epistemic negation **not**  $\mathit{guilty}(X)$  will become closed; however,  $\forall X (person(X) \wedge \mathbf{not} \mathit{guilty}(X))$  is not instance-closed because  $X$  is not a free variable and the epistemic negation **not**  $\mathit{guilty}(X)$  will not become closed after grounding.

We use instance-closed epistemic formulas to construct rules and logic programs.

**Definition 1 (General logic program)** A *general logic program* (logic program for short) is a finite set of *rules* of the form  $H \leftarrow B$ , where  $H$  and  $B$  are instance-closed epistemic formulas without nested epistemic negations.

Note that like the approaches of [16, 18], where modal operators **K** and **M** are not nested in logic programs, we do not consider logic programs with nested epistemic negations (**not** is also a modal operator); a nested epistemic negation like **not** (**not**  $F$ ) intuitively expresses "That  $F$  cannot be proved to be true cannot be proved to be true" and such expressions seem to have rare applications in practical scenarios.

For convenience, for a rule  $r : H \leftarrow B$  we refer to  $B$  and  $H$  as the body and head of  $r$ , denoted  $body(r)$  and  $head(r)$ , respectively. When  $head(r)$  is empty, we rewrite the rule as  $\perp \leftarrow body(r)$ ; when  $body(r)$  is empty, we omit the rule operator  $\leftarrow$ .

**Definition 2 (Normal epistemic program)** A *normal epistemic program* is a logic program consisting of rules of the form

$$A_0 \leftarrow A_1 \wedge \dots \wedge A_m \wedge \mathbf{not} A_{m+1} \wedge \dots \wedge \mathbf{not} A_n \quad (2)$$

where  $n \geq m \geq 0$  and each  $A_i$  is an atom without equality and function symbols except constants.

Furthermore, as usual a *normal logic program* consists of rules of the above form (2) except that epistemic negation **not** is replaced by default negation  $\neg$ ; a *positive logic program* is a  $\neg$ -free normal logic program.

**Definition 3 (Propositional program)** A *propositional program*  $\Pi$  is a logic program which contains no variables, no function symbols except constants, and no equalities. The *Herbrand base* of  $\Pi$  is defined as usual. Any subset of the Herbrand base is a *Herbrand interpretation* of  $\Pi$ .

### 2.2.2 Grounding

In a logic program  $\Pi$ , some rules may contain free variables. In ASP, these free variables will be instantiated by constants from a finite set – usually the set  $\mathcal{C}_\Pi$  of constants occurring in  $\Pi$ . With no loss in generality, we assume that  $\mathcal{C}_\Pi$  consists of all constants in  $\Pi$  (in case that some constant  $a$  of the domain does not appear in  $\Pi$ , we may have it by adding to  $\Pi$  a dummy rule  $p(a) \leftarrow p(a)$ ). Then for any logic program  $\Pi$ ,  $\mathcal{C}_\Pi$  is unique.

A *closed instance* of a rule  $r$  is  $r$  with all free variables replaced by constants in  $\mathcal{C}_\Pi$ . The *grounding* of  $\Pi$ , denoted  $ground(\Pi)$ , is the set of all closed instances of all rules in  $\Pi$ .

Note that each rule  $H \leftarrow B$  with the set  $S$  of free variables may also be viewed as a *globally* universally quantified rule  $\forall S(H \leftarrow B)$ , where the domain of each variable in  $S$  is  $\mathcal{C}_\Pi$  while the domain of the other (*locally* quantified) variables is  $\mathcal{N}_\Sigma$ . Only globally universally quantified variables will be instantiated over their domain  $\mathcal{C}_\Pi$  for the grounding  $ground(\Pi)$ .

To sum up, a logic program  $\Pi$  is viewed as shorthand for  $ground(\Pi)$ , where each free variable in  $\Pi$  is viewed as shorthand for constants in  $\mathcal{C}_\Pi$ .

As rule bodies and heads in  $\Pi$  are instance-closed epistemic formulas, for every epistemic negation **not**  $F$  in  $ground(\Pi)$ ,  $F$  itself is a closed epistemic formula. Therefore, in the sequel unless otherwise stated, for any epistemic negation **not**  $F$  we assume  $F$  is a closed epistemic formula.

### 2.2.3 Satisfaction and Models

Next, we extend the satisfaction relation of  $\mathcal{L}_\Sigma$  to logic programs. As epistemic negation works at a meta level over a collection of interpretations, the definition of satisfaction/models of epistemic formulas should be based on a collection of interpretations.

**Definition 4** Let  $\mathcal{A}$  be a collection of interpretations and  $I \in \mathcal{A}$ .

- (1) Let  $F$  be a closed formula. **not**  $F$  is satisfied by/true in  $\mathcal{A}$  if  $F$  is false in some  $J \in \mathcal{A}$ , and false, otherwise.  $I$  satisfies **not**  $F$  w.r.t.  $\mathcal{A}$  if **not**  $F$  is true in  $\mathcal{A}$ .
- (2)  $I$  satisfies a closed epistemic formula  $E$  w.r.t.  $\mathcal{A}$  if  $I$  satisfies  $E$  as in first-order logic except that the satisfaction of epistemic negations in  $E$  is determined by (1).
- (3)  $I$  satisfies a closed instance  $r$  of a rule w.r.t.  $\mathcal{A}$  if  $I$  satisfies  $head(r)$  w.r.t.  $\mathcal{A}$  once  $I$  satisfies  $body(r)$  w.r.t.  $\mathcal{A}$ .

- (4)  $\mathcal{A}$  is a *collection of models* of a logic program  $\Pi$  if every  $J \in \mathcal{A}$  satisfies all rules in  $ground(\Pi)$  w.r.t.  $\mathcal{A}$ . In this case, every  $J \in \mathcal{A}$  is a model of  $\Pi$  (w.r.t.  $\mathcal{A}$ ).  $\Pi$  is *consistent* if it has a model.
- (5) A closed epistemic formula  $E$  is *true* in  $\Pi$  if  $\Pi$  has some collection  $\mathcal{A}$  of models such that every  $J \in \mathcal{A}$  satisfies  $E$  w.r.t.  $\mathcal{A}$ .

Observe the following properties of satisfaction of a logic program  $\Pi$ .

First, when  $\Pi$  contains no epistemic negation, satisfaction in Definition 4 reduces to that in first-order logic. In this case, we omit  $\mathcal{A}$ . Then  $I$  satisfies a closed rule instance  $r$  if  $I$  satisfies  $head(r)$  or  $I$  does not satisfy  $body(r)$ , and  $I$  is a model of  $\Pi$  if it satisfies all rules in  $ground(\Pi)$ ; moreover, a model  $I$  is *minimal* if  $\Pi$  has no model  $J$  that is a proper subset of  $I$ .

Second, satisfaction of an epistemic negation  $\mathbf{not} F$  in  $I$  w.r.t.  $\mathcal{A}$  is determined only by  $\mathcal{A}$ , i.e.,  $I$  can be ignored.

Finally, when  $\mathcal{A}$  consists of a single interpretation  $I$ , the notions of satisfaction, models and consistency are the same as in first-order logic, i.e., they are determined only by  $I$  and  $\mathcal{A}$  can be ignored. In this special case, the epistemic negation operator  $\mathbf{not}$  coincides with the operator  $\neg$ , i.e.,  $I$  satisfies  $\mathbf{not} F$  w.r.t.  $\mathcal{A}$  iff  $I$  satisfies  $\neg F$  iff  $F$  is false in  $I$ , where  $F$  is a closed formula. Hence the following proposition is immediate.

**Proposition 1** *Let  $\Pi$  be a logic program and  $\Pi^\neg$  be  $\Pi$  with all epistemic negations  $\mathbf{not} F$  replaced by default negations  $\neg F$ . For any interpretation  $I$ ,  $\mathcal{A} = \{I\}$  is a collection of models of  $\Pi$  iff  $I$  is a model of  $\Pi^\neg$ .*

The following theorem is important, as it lays a theoretical basis for the introduction of our novel program transformation, which is described in the next section.

**Theorem 2** *Let  $\Pi$  be a logic program such that for every  $\mathbf{not} F$  in  $ground(\Pi)$ ,  $F$  is true in every model of  $\Pi$ . Let  $\Pi^\neg$  be  $\Pi$  with each epistemic negation  $\mathbf{not} F$  replaced by default negation  $\neg F$ . Then  $\Pi$  and  $\Pi^\neg$  have the same models.*

### 3 Program Transformation and Epistemic Reducts

In ASP, it is common to transform a logic program into a reduct which is free of negation or modal operators. For instance, for a normal logic program  $\Pi$ , the seminal *GL-reduct*  $\Pi^I$  w.r.t. a given interpretation  $I$  is obtained from  $ground(\Pi)$  by removing first all rules whose bodies contain a default negation  $\neg A$  with  $A \in I$ , and then all  $\neg A$  from the remaining rules [17]. Similarly, when  $\Pi$  is a logic program extended with modal operators  $\mathbf{K}$  and  $\mathbf{M}$ , [15, 16, 32, 18] defined a transformation w.r.t. a given set  $\mathcal{A}$  of interpretations by eliminating/replacing all modal literals in  $ground(\Pi)$  in terms of whether they are true or not in  $\mathcal{A}$ .

Note that these existing definitions of program transformations are based on an assumption which is either a given interpretation or a given set of interpretations.

In this paper we define a program transformation in an alternative way, which is based on a notion of *guess on provability* of epistemic negation, instead of on an assumption of a set of interpretations. Recall that an epistemic negation  $\mathbf{not} F$  expresses that there is no evidence proving that  $F$  is true, where  $F$  is *proved true* if it is true in every answer set.

**Definition 5 (Guess on provability of epistemic negation)** Let  $\Pi$  be a logic program, and furthermore let  $Ep(\Pi)$  denote the set of epistemic negations  $\mathbf{not} F$  in  $ground(\Pi)$ . A *guess of  $\Pi$  on provability of epistemic negation* is a subset  $\Phi$  of  $Ep(\Pi)$  such that for every  $\mathbf{not} F \in \Phi$ , it is guessed that  $F$  couldn't be proved true, and for every  $\mathbf{not} F \in Ep(\Pi) \setminus \Phi$ , it is guessed that  $F$  would be proved true.

Once a guess  $\Phi$  is given, we can transform program  $\Pi$  by replacing all epistemic negations in terms of  $\Phi$ . There would be different replacements for epistemic negations, which would lead to different program transformations. The simplest yet naive one is to replace  $\mathbf{not} F$  with  $\top$  if  $\mathbf{not} F \in \Phi$ , and with  $\perp$ , otherwise. It turns out that this transformation incurs both epistemic circular justifications and the multiple world view problem, analogously to the cases in [15].

The key idea of our program transformation is as follows. We first assume that the guess on all  $\mathbf{not} F \in \Phi$  is correct and thus replace them with  $\top$ . Then, for every  $\mathbf{not} F \in Ep(\Pi) \setminus \Phi$ , instead of replacing it with  $\perp$ , we replace it with  $\neg F$ . The intuition and rationale for the latter replacement is as follows: if  $\Phi$  is a *correct guess*, once all epistemic negations  $\mathbf{not} F \in \Phi$  in  $ground(\Pi)$  are replaced with  $\top$ , which leads to a new program  $\Pi^\top$ , for every  $\mathbf{not} F$  in  $\Pi^\top$ , the formula  $F$  is supposed to be true in every answer set of  $\Pi^\top$ . Let  $\Pi^\Phi$  be  $\Pi^\top$  with each  $\mathbf{not} F$  replaced by  $\neg F$ ; then by Theorem 2, where *model* is analogously replaced by *answer set*, we expect that  $\Pi^\top$  and  $\Pi^\Phi$  have the same answer sets. This rational justification of the replacements for epistemic negations leads to the following novel program transformation.

**Definition 6 (Epistemic reducts)** Let  $\Pi$  be a logic program and let  $\Phi \subseteq Ep(\Pi)$  be a guess of  $\Pi$  on provability of epistemic negation. The *epistemic reduct*  $\Pi^\Phi$  of  $\Pi$  w.r.t.  $\Phi$  is obtained from  $ground(\Pi)$  by replacing every  $\mathbf{not} F \in \Phi$  with  $\top$ , and every  $\mathbf{not} F \in Ep(\Pi) \setminus \Phi$  with  $\neg F$ . We call  $\Pi$  *consistent* w.r.t.  $\Phi$  if  $\Pi^\Phi$  is consistent.

In the Introduction we mentioned that a world view  $\mathcal{A}$  is said to have an epistemic circular justification if some object literal  $L$  being true in some interpretation  $I \in \mathcal{A}$  is due to  $\mathbf{KL}$  (or its equivalent modal literals expressing that  $L$  is true in every interpretation  $J \in \mathcal{A}$ ) being treated true in the program transformation w.r.t.  $\mathcal{A}$ . In our language,  $\mathbf{KL}$  is shorthand for  $\neg \mathbf{not} L$ , and in our program transformation w.r.t. a guess  $\Phi$ ,  $\neg \mathbf{not} L$  will be either treated  $\neg \top$  (when  $\mathbf{not} L \in \Phi$ ), which evaluates to false, or treated  $\neg \neg L$  (when  $\mathbf{not} F \in Ep(\Pi) \setminus \Phi$ ), which evaluates to  $L$ . This means that our definition of the program transformation would never incur epistemic circular justifications and thus guarantees that world views based on the epistemic reducts will be free of epistemic circular justifications.

## 4 FLP Answer Set Semantics with Epistemic Negation

Now that all epistemic negations have been removed from a logic program  $\Pi$ , leading to an epistemic reduct  $\Pi^\Phi$  w.r.t. a guess  $\Phi$ , we can apply any existing answer set semantics to compute all answer sets  $\mathcal{A}$  of  $\Pi^\Phi$ . Such answer sets are our expected ones if they agree with the guess  $\Phi$ , i.e., every  $\mathbf{not} F \in \Phi$  is true and every  $\mathbf{not} F \in Ep(\Pi) \setminus \Phi$  is false in  $\mathcal{A}$ .

For simple illustration, in this section we extend the well-known FLP answer set semantics of [9] to logic programs with epistemic negation. The following definition is from [28], which lifts Faber et al.'s FLP semantics to general logic programs without epistemic negation.

**Definition 7** Let  $\Pi$  be a logic program without epistemic negation and  $I$  an interpretation. The *FLP-reduct* of  $\Pi$  w.r.t.  $I$  is  $f\Pi^I = \{r \in \text{ground}(\Pi) \mid I \text{ satisfies } \text{body}(r)\}$ , and  $I$  is an *FLP answer set* of  $\Pi$  if  $I$  is a (subset-)minimal model of  $f\Pi^I$ .

Our epistemic FLP answer set semantics (EFLP semantics for short) is defined as follows.

**Definition 8 (EFLP semantics)** Let  $\Pi$  be a logic program and  $\Phi$  a guess such that  $\Pi^\Phi$  is a consistent epistemic reduct. The collection  $\mathcal{A}$  of all FLP answer sets of  $\Pi^\Phi$  is a *candidate world view* of  $\Pi$  w.r.t.  $\Phi$  if (a)  $\mathcal{A} \neq \emptyset$ , (b) every **not**  $F \in \Phi$  is true in  $\mathcal{A}$ , and (c) every **not**  $F \in \text{Ep}(\Pi) \setminus \Phi$  is false in  $\mathcal{A}$ . A candidate world view  $\mathcal{A}$  w.r.t. a guess  $\Phi$  is a *world view* if  $\Phi$  is maximal (i.e., there is no candidate world view w.r.t. any guess  $\Phi' \supset \Phi$ ). Every FLP answer set in a world view is called an *EFLP answer set*.

Obviously, for logic programs without epistemic negation,  $\text{Ep}(\Pi) = \emptyset$  and EFLP semantics reduces to FLP semantics.

The rationale for the *maximality condition* of a guess for a world view is twofold. On the one hand, intuitively we expect to find world views  $\mathcal{A}$  of  $\Pi$  such that all epistemic negations **not**  $F$  in  $\text{Ep}(\Pi)$  are satisfied by  $\mathcal{A}$ . This intention is often violated because no such a world view would exist for some logic programs, such as  $\Pi = \{p \leftarrow \text{not } q, q \leftarrow \text{not } p\}$ . Therefore, in Definition 8 we relaxed the condition by requiring that a *maximal* subset  $\Phi \subseteq \text{Ep}(\Pi)$  of epistemic negations are satisfied by  $\mathcal{A}$ .

On the other hand, from a theoretical angle, since epistemic negation is at a meta level, towards a “type theory” or higher-order theory of epistemic negation, it makes sense to assume that **not**  $F$  in  $\Pi$  is preferably true. This is analogous to the situation for *type-0 objects* (answer sets), where we apply default negation to formulas w.r.t. an answer set and assume  $\neg F$  in  $\Pi$  to be true in the answer set whenever possible (CWA or minimal models); epistemic negation applies the same principle on *type-1 objects* (world views), i.e., we apply epistemic negation to formulas w.r.t. a world view and assume **not**  $F$  in  $\Pi$  to be true in the world view whenever possible (maximal guesses). As a result, *by applying default negation we minimize the knowledge in every answer set, and by applying epistemic negation we minimize the knowledge in every world view.*

It is with this condition of maximal epistemic guesses for world views that we are ready to introduce the following formulation of the multiple world view problem.

**Definition 9 (Multiple world view problem)** An epistemic answer set semantics is said to have the *multiple world view problem* if some logic program may have two world views such that one satisfies a proper superset of epistemic negations of the other.

Obviously, due to the condition of maximal guesses EFLP semantics is free of the multiple world view problem. This will be further illustrated in Example 6.

**Definition 10 (Query evaluation)** Let  $\Pi$  be a logic program and  $F$  a closed formula. We say  $F$  is true in  $\Pi$  under EFLP semantics if  $\Pi$  has a world view  $\mathcal{A}$  such that  $F$  is true in every EFLP answer set in  $\mathcal{A}$ .

The following result shows that EFLP answer sets of  $\Pi$  are models of  $\Pi$ .

**Theorem 3** *Let  $\Pi$  be a logic program and  $\mathcal{A}$  a world view of  $\Pi$  w.r.t. a guess  $\Phi$ . Then  $\mathcal{A}$  is a collection of models of  $\Pi$ .*

Let  $\Pi$  be a normal epistemic program and  $\Pi^\neg$  be the normal logic program obtained from  $\Pi$  by replacing epistemic negation **not** with default negation  $\neg$ . Let  $I$  be an interpretation and  $(\Pi^\neg)^I$  be the GL-reduct of  $\Pi^\neg$ ; then  $I$  is a *standard answer set* of  $\Pi^\neg$  if  $I$  is the least model of  $(\Pi^\neg)^I$  [17].

The following result shows that answer sets of  $\Pi$  under EFLP semantics coincide with those of  $\Pi^\neg$  under the standard answer set semantics.

**Theorem 4** *Let  $\Pi$  be a normal epistemic program and let  $\Pi^\neg$  be the normal logic program obtained from  $\Pi$  by replacing **not** with  $\neg$ . Then (1) every world view  $\mathcal{A}$  of  $\Pi$  fulfills  $|\mathcal{A}| = 1$  (i.e.,  $\mathcal{A}$  is a singleton), and (2)  $\mathcal{A} = \{I\}$  is a world view of  $\Pi$  iff  $I$  is an FLP answer set of  $\Pi^\neg$  iff  $I$  is a standard answer set of  $\Pi^\neg$ .*

Next we consider a few simple examples to illustrate the novelty and suitability of our approach.

**Example 2** *The following logic program uses epistemic negation to formalize the well-known presumption of innocence:*

$$\begin{array}{ll} \Pi_1 : & innocent(John) \vee guilty(John) & r_1 \\ & innocent(X) \leftarrow \mathbf{not} guilty(X) & r_2 \end{array}$$

The first rule  $r_1$  says that John is either innocent or guilty. The rule  $r_2$  asserts that one is presumed innocent if there is no evidence proving s/he is guilty; specifically it means that if  $guilty(John)$  cannot be proved to be true from  $r_1$ , then conclude  $innocent(John)$ . Evidently,  $guilty(John)$  cannot be proved true from  $r_1$ ; therefore all intuitive answer sets of  $\Pi_1$  should contain  $innocent(John)$ , i.e., John is innocent.

The grounding of  $\Pi_1$  is

$$\begin{array}{ll} \mathit{ground}(\Pi_1) : & innocent(John) \vee guilty(John) & r_1 \\ & innocent(John) \leftarrow \mathbf{not} guilty(John) & r'_2 \end{array}$$

It contains only one epistemic negation, thus  $Ep(\Pi_1) = \{\mathbf{not} guilty(John)\}$ . So we have two guesses:  $\Phi_1 = \{\mathbf{not} guilty(John)\}$  and  $\Phi_2 = \emptyset$ .

We start with the largest guess  $\Phi_1$  and check if there is a world view w.r.t.  $\Phi_1$ . Note  $Ep(\Pi_1) \setminus \Phi_1 = \emptyset$ . The epistemic reduct w.r.t.  $\Phi_1$  is

$$\begin{array}{ll} \Pi_1^{\Phi_1} : & innocent(John) \vee guilty(John) & r_1 \\ & innocent(John) \leftarrow \top & r''_2 \end{array}$$

$\Pi_1^{\Phi_1}$  is consistent and has a unique FLP answer set  $I = \{\text{innocent}(\text{John})\}$ . Let  $\mathcal{A} = \{I\}$ . As  $\text{guilty}(\text{John})$  is false in  $I$ , **not**  $\text{guilty}(\text{John})$  is true in  $\mathcal{A}$ , thus  $\mathcal{A}$  is a candidate world view of  $\Pi_1$ . Since  $\Phi_1$  is the largest guess,  $\mathcal{A}$  is a world view of  $\Pi_1$  w.r.t.  $\Phi_1$ .

As  $\Phi_2 \subset \Phi_1$ , there would be no world view w.r.t.  $\Phi_2$ . Hence,  $\Pi_1$  has only one world view  $\mathcal{A} = \{\{\text{innocent}(\text{John})\}\}$ , meaning that under the presumption of innocence, John is innocent. This conforms to our expectation.

**Example 3** To demonstrate the necessity of epistemic reasoning with epistemic negation, [15] introduced the following well-known college scholarship awarding problem with the program:

$$\begin{array}{ll}
\Pi_2: \text{eligible}(X) \leftarrow \text{highGPA}(X) & r_1 \\
\text{eligible}(X) \leftarrow \text{minority}(X) \wedge \text{fairGPA}(X) & r_2 \\
\sim \text{eligible}(X) \leftarrow \text{lowGPA}(X) & r_3 \\
\text{interview}(X) \leftarrow \neg \text{eligible}(X) \wedge \neg \sim \text{eligible}(X) & r_4 \\
\text{fairGPA}(\text{Mike}) \vee \text{highGPA}(\text{Mike}) & r_5
\end{array}$$

Gelfond argued that Mike was intended to be interviewed, i.e.,  $\text{interview}(\text{Mike})$  should be included in every answer set of  $\Pi_2$ ; however, he observed that the fourth rule  $r_4$  was not powerful enough to formalize the intended statement that the students whose eligibility is not determined by the college rules should be interviewed by the scholarship committee. It was due to this observation that [15, 16] proposed his modal formalism in which  $r_4$  was replaced by the rule

$$\text{interview}(X) \leftarrow \neg \mathbf{K} \text{eligible}(X) \wedge \neg \mathbf{K} \sim \text{eligible}(X).$$

Next, we show that this problem can be suitably handled using epistemic negation under EFLP semantics. As shown in [14], for any atom  $p(X)$ , the strong negation  $\sim p(X)$  can be compiled away by introducing a fresh predicate  $p'(X)$  along with a rule  $\perp \leftarrow p(X) \wedge p'(X)$ . We formulate the above problem by rewriting  $\Pi_2$  by replacing  $\neg$  with **not** and  $\sim \text{eligible}(X)$  with  $\text{eligible}'(X)$ , and adding the rule  $\perp \leftarrow \text{eligible}(X) \wedge \text{eligible}'(X)$ ; this yields a new program  $\Pi_2'$  with the grounding

$$\begin{array}{ll}
\text{ground}(\Pi_2'): \text{eligible}(\text{Mike}) \leftarrow \text{highGPA}(\text{Mike}) & r'_1 \\
\text{eligible}(\text{Mike}) \leftarrow \text{minority}(\text{Mike}) \wedge \text{fairGPA}(\text{Mike}) & r'_2 \\
\text{eligible}'(\text{Mike}) \leftarrow \text{lowGPA}(\text{Mike}) & r'_3 \\
\text{interview}(\text{Mike}) \leftarrow \mathbf{not} \text{eligible}(\text{Mike}) \wedge \mathbf{not} \text{eligible}'(\text{Mike}) & r'_4 \\
\text{fairGPA}(\text{Mike}) \vee \text{highGPA}(\text{Mike}) & r_5 \\
\perp \leftarrow \text{eligible}(\text{Mike}) \wedge \text{eligible}'(\text{Mike}) & r_6
\end{array}$$

$Ep(\Pi_2') = \{\mathbf{not} \text{eligible}(\text{Mike}), \mathbf{not} \text{eligible}'(\text{Mike})\}$ , so there are four guesses. We start with the largest guess  $\Phi_1 = \{\mathbf{not} \text{eligible}(\text{Mike}), \mathbf{not} \text{eligible}'(\text{Mike})\}$  and check if  $\Pi_2'$  has a world view w.r.t.  $\Phi_1$ .

The epistemic reduct  $\Pi_2'^{\Phi_1}$  w.r.t.  $\Phi_1$  is  $\text{ground}(\Pi_2')$  except that  $r'_4$  is replaced by the rule

$$\text{interview}(\text{Mike}) \leftarrow \top \wedge \top \quad r''_4$$

$\Pi_2'^{\Phi_1}$  is consistent and has the following two FLP answer sets:



$$I_1 = \{fairGPA(Mike), interview(Mike)\}$$

$$I_2 = \{highGPA(Mike), eligible(Mike), interview(Mike)\}$$

Let  $\mathcal{A} = \{I_1, I_2\}$ . As the atoms  $eligible(Mike)$  and  $eligible'(Mike)$  are both false in  $I_1$ , both **not**  $eligible(Mike)$  and **not**  $eligible'(Mike)$  are true in  $\mathcal{A}$ . As  $Ep(\Pi'_2) \setminus \Phi_1 = \emptyset$  and  $\Phi_1$  is the largest guess,  $\mathcal{A}$  is a unique world view of  $\Pi'_2$ .

Note that  $interview(Mike)$  appears in every answer set in  $\mathcal{A}$ , meaning *Mike* should be interviewed, as we expected.

**Remark 1** We can also formulate the above problem without using strong negation by rewriting  $\Pi_2$  with  $\neg$  replaced by **not**, and  $\sim$  by  $\neg$ ; thus  $r_4$  is rewritten as

$$interview(X) \leftarrow \mathbf{not} \text{ eligible}(X) \wedge \mathbf{not} \neg \text{ eligible}(X).$$

This rule directly expresses that a student should be interviewed if we can neither prove s/he is eligible nor not eligible. Then, the grounding of the new program consists of  $r'_1$ ,  $r'_2$ ,  $r_5$ , and the following two rules:

$$\begin{array}{ll} \neg \text{ eligible}(Mike) \leftarrow lowGPA(Mike) & r_3^- \\ interview(Mike) \leftarrow \mathbf{not} \text{ eligible}(Mike) \wedge \mathbf{not} \neg \text{ eligible}(Mike) & r_4^- \end{array}$$

It is easy to check that this new program has a unique world view  $\mathcal{A} = \{I_1, I_2\}$  w.r.t. the largest guess  $\Phi_1 = \{\mathbf{not} \text{ eligible}(Mike), \mathbf{not} \neg \text{ eligible}(Mike)\}$ , where

$$I_1 = \{fairGPA(Mike), interview(Mike)\}$$

$$I_2 = \{highGPA(Mike), eligible(Mike), interview(Mike)\}$$

which shows that *Mike* should be interviewed, as expected.

**Example 4 (Closed world rules)** Under EFLP semantics, we can directly formulate the closed world assumption using closed world rules of the form  $\neg p \leftarrow \mathbf{not} p$ , which expresses that when failing to prove  $p$  to be true, we assert  $\neg p$ . Moreover, we can also state its opposite using rules of the form  $p \leftarrow \mathbf{not} \neg p$ . We can further combine them, leading to the following interesting logic program.

$$\begin{array}{ll} \Pi : \quad \neg p \leftarrow \mathbf{not} p & r_1 \\ \quad \quad p \leftarrow \mathbf{not} \neg p & r_2 \end{array}$$

It is easy to check that this program has two world views:  $\mathcal{A}_1 = \{\emptyset\}$  w.r.t. the guess  $\Phi_1 = \{\mathbf{not} p\}$  and  $\mathcal{A}_2 = \{\{p\}\}$  w.r.t.  $\Phi_2 = \{\mathbf{not} \neg p\}$ . This conforms to our intuition that either  $\neg p$  or  $p$  can be concluded from  $\Pi$ , depending on whether we choose to apply CWA on  $p$  (rule  $r_1$ ) or on  $\neg p$  (rule  $r_2$ ).

**Remark 2** Gelfond [14] used a rule  $\sim p \leftarrow \neg p$  (instead of  $\neg p \leftarrow \mathbf{not} p$ ) to formalize CWA on an atom  $p$ , which achieves the effect that if  $p$  is not in an answer set  $I$ , then  $\sim p$  is in  $I$ . As indicated by [15], this formalization of CWA is problematic. For example, by this formalization the logic program  $\Pi = \{p \vee q\}$  extended with the rule  $\sim p \leftarrow \neg p$  would have two answer sets, viz.

$I_1 = \{\sim p, q\}$  and  $I_2 = \{p\}$ ; thus both  $p$  and  $q$  are *unknown* in  $\Pi$  under the semantics of [14].<sup>5</sup> In contrast, if  $\Pi$  is extended with the CWA rule  $\neg p \leftarrow \mathbf{not} p$ , we would obtain a unique world view  $\{\{q\}\}$  under EFLP semantics, so  $p$  is false and  $q$  is true in  $\Pi$ , as expected.<sup>6</sup>

**Example 5** Let  $\Pi = \{r \leftarrow \mathbf{not} p \wedge \neg r, p \leftarrow \neg q, q \leftarrow \neg p\}$ . Intuitively, the last two rules say that either  $\{p\}$  or  $\{q\}$  is an answer set of  $\Pi$ , and the first rule is a constraint that every answer set of  $\Pi$  should contain  $p$ . Therefore,  $\Pi$  intuitively should have a single answer set  $\{p\}$ . It is easy to check that under EFLP semantics this program has a single world view  $\mathcal{A} = \{\{p\}\}$  w.r.t. the guess  $\Phi = \emptyset$ . Note that  $\mathcal{A}$  is not a world view under the approach of [16], where  $\mathbf{not} p$  in  $\Pi$  is replaced by  $\neg \mathbf{K}p$ .

**Example 6** Consider again the logic program  $\Pi = \{p \leftarrow \mathbf{M}p\}$  in Example 1. As mentioned earlier,  $\mathbf{M}p$  is shorthand for  $\mathbf{not} \neg p$ , so this program can be rewritten as  $\Pi = \{p \leftarrow \mathbf{not} \neg p\}$ .  $\Phi = \{\mathbf{not} \neg p\}$  is the largest guess and it is easy to check that  $\{\{p\}\}$  is a unique world view of  $\Pi$  w.r.t.  $\Phi$  under EFLP semantics.

Similarly, the second logic program in Example 1 can be rewritten as

$$\begin{array}{ll} \Pi : & p \leftarrow \mathbf{not} \neg q \wedge \neg q & r_1 \\ & q \leftarrow \mathbf{not} \neg p \wedge \neg p & r_2 \end{array}$$

It has a single world view  $\mathcal{A} = \{\{p\}, \{q\}\}$  w.r.t. the guess  $\Phi = \{\mathbf{not} \neg q, \mathbf{not} \neg p\}$ .

For the second program, as shown in Example 1 applying the approaches of [15, 16, 18] will produce two world views, including the undesired one  $\{\emptyset\}$ .

Note that EFLP semantics minimizes the knowledge in world views by requiring that every world view should satisfy as many epistemic negations as possible; this provides a satisfying justification for the above two programs why  $\mathcal{A}_1 = \{\{p\}\}$  is the right world view and  $\mathcal{A}_2 = \{\emptyset\}$  is not.

**Example 7** Consider the logic program  $\Pi = \{p \leftarrow \mathbf{not} p \vee p\}$ .<sup>7</sup> Note that  $\mathbf{not} p \vee p$  is a tautology in that for any collection  $\mathcal{A}$  of interpretations, every  $I \in \mathcal{A}$  satisfies  $\mathbf{not} p \vee p$  w.r.t.  $\mathcal{A}$ . Then  $p$  can be inferred by applying the rule  $p \leftarrow \mathbf{not} p \vee p$ ; thus  $\Pi$  is supposed to have a unique world view  $\mathcal{A} = \{\{p\}\}$ .

This program has two guesses:  $\Phi_1 = \{\mathbf{not} p\}$  and  $\Phi_2 = \emptyset$ . It is easy to check that  $\Pi$  has no world view w.r.t.  $\Phi_1$ , but has a world view  $\mathcal{A} = \{\{p\}\}$  w.r.t.  $\Phi_2$ .

Note that the approaches of [15, 16, 18] are not applicable to this program.

<sup>5</sup>Under the semantics of [14],  $p$  is true in an answer set  $I$  if  $p \in I$ , false if  $\sim p \in I$ , and *unknown*, otherwise.  $p$  is true/false in a logic program  $\Pi$  if it is true/false in all answer sets of  $\Pi$ , and *unknown*, otherwise.

<sup>6</sup> Alternatively, one might want to extend  $\Pi$  with the rule  $\sim p \leftarrow \mathbf{not} p$ . Then  $\sim p$  can be compiled away similarly as in Example 3. The resulting program  $\Pi' = \{p \vee q, p' \leftarrow \mathbf{not} p, \perp \leftarrow p \wedge p'\}$  has a unique world view  $\{\{q, p'\}\}$  under EFLP semantics, meaning  $p$  is false and  $q$  is true in  $\Pi$ .

<sup>7</sup>Note that  $\vee$  is classical logic disjunction connective, instead of epistemic one. Epistemic disjunctions are usually expressed using the epistemic operator  $|$  in the literature. A classical disjunction  $\neg p \vee p$  is a tautology, but an epistemic disjunction  $\neg p | p$  is not a tautology since it does not follow the law of the excluded middle (see [13] for detailed explanations).

Observe that in our logic language  $\mathcal{L}_\Sigma$  the formula  $\neg p \vee p$  is also a tautology, and  $p \leftarrow \neg p \vee p$  is not equivalent to  $\{p \leftarrow \neg p, p \leftarrow p\}$ . Therefore, under EFLP semantics the rule  $p \leftarrow \mathbf{not} p \vee p$  is not equivalent to  $\{p \leftarrow \mathbf{not} p, p \leftarrow p\}$ .

## 5 Computational Complexity

For a logic program  $\Pi$ , each guess  $\Phi$  leads to at most one candidate world view w.r.t.  $\Phi$ . As there are at most  $2^{|Ep(\Pi)|}$  guesses (where  $|Ep(\Pi)|$  is the number of epistemic negations in  $Ep(\Pi)$ ),  $\Pi$  has at most  $2^{|Ep(\Pi)|}$  candidate world views. In view of the fact that if  $\Pi$  has a world view w.r.t.  $\Phi$  then it will have no world view w.r.t. any  $\Phi' \supset \Phi$  or  $\Phi' \subset \Phi$ , the following result is immediate.

**Proposition 5** *A logic program  $\Pi$  has at most  $\binom{n}{\lfloor n/2 \rfloor}$  many world views under EFLP semantics, where  $n = |Ep(\Pi)|$ .*

Indeed,  $\binom{n}{\lfloor n/2 \rfloor}$  is the maximal size of an antichain in the powerset lattice of a set of cardinality  $n$  [30]. In the two most extreme cases that  $\Pi$  has a world view w.r.t.  $\Phi$ , where  $\Phi = Ep(\Pi)$  or  $\Phi = \emptyset$ ,  $\Pi$  has only one world view.

As the satisfiability of arbitrary first-order theories is undecidable, we only address the computational complexity of EFLP answer sets for propositional programs under Herbrand interpretations. Recall that in this case, deciding the existence of FLP answer sets of a given logic program  $\Pi$  without epistemic negation is  $\Sigma_2^p$ -complete [28], where the  $\Sigma_2^p$ -hardness is inherited from disjunctive logic programs [9, 7].

We first consider the complexity of recognizing a suitable guess.

**Theorem 6** *Given a propositional program  $\Pi$  and a guess  $\Phi$  for it, deciding whether  $\Pi$  has a candidate world view w.r.t.  $\Phi$  is  $D_2^p$ -complete.*

The class  $D_2^p$  consists of all problems  $(P_1, P_2)$  whose instances are pairs  $(I_1, I_2)$  of instances  $I_1$  of a problem  $P_1$  in  $\Sigma_2^p$  and  $I_2$  of a problem  $P_2$  in  $\Pi_2^p$ , respectively.

Informally, we must check the conditions (a)-(c) of a candidate world view in Definition 8, where (a) establishes *a fortiori* also consistency of  $\Pi^\Phi$ ; as  $\Pi^\Phi$  is constructible in polynomial time, we can solve (a) and (b) in  $\Sigma_2^p$  and (c) in  $\Pi_2^p$ , as deciding whether  $\Pi^\Phi$  has some FLP answer set that satisfies (resp. does not satisfy) a formula  $F$  is in  $\Sigma_2^p$ ; thus, the problem is in  $D_2^p$ . The  $D_2^p$ -hardness is shown by a reduction from deciding whether, given a pair  $(\Pi_1, \Pi_2)$  of propositional programs, both  $\Pi_1$  has some FLP-answer set and  $\Pi_2$  has no FLP answer set; this problem is  $D_2^p$ -complete. Informally,  $\Pi_2$  is modified to a program  $\Pi_2'$  whose answer sets amount to those of  $\Pi_2$  plus an extra answer set  $\{A\}$ , where  $A$  is a fresh atom;  $\Pi$  contains  $\Pi_1$  and  $\Pi_2'$  and has a candidate world view w.r.t.  $\Phi = \emptyset$ , where  $EP(\Pi) = \{\mathbf{not} A\}$ , just if  $\Pi_1$  has some FLP answer set and  $A$  is true in all FLP answer sets of  $\Pi_2'$  (i.e.,  $\Pi_2$  has no FLP answer set).

From this result, we obtain that deciding program consistency under candidate world views is at the third level of the polynomial hierarchy. More precisely,

**Theorem 7** *Given a propositional program  $\Pi$ , deciding whether  $\Pi$  has (i) some candidate world view and (ii) some world view, i.e., deciding EFLP answer set existence, are both  $\Sigma_3^p$ -complete.*

Indeed, a guess  $\Phi \subseteq Ep(\Pi)$  such that  $\Pi$  has some candidate world view w.r.t.  $\Phi$  can be verified in polynomial time with an oracle for  $D_2^p$ , and hence also with one for  $\Sigma_2^p$ . This places the problem in  $\Sigma_3^p$ . The matching  $\Sigma_3^p$ -hardness can be shown by a reduction from evaluating quantified Boolean formulas of the form  $\exists Z \forall Y \exists X \phi$ . Informally, epistemic negations **not**  $X$ , **not**  $\bar{X}$  are used to guess an assignment to the  $X$  atoms, and some other epistemic negation serves to check that for each assignment to the  $Y$  atoms, some assignment to  $Z$  makes  $\phi$  true; this test for  $Y$  and  $Z$  is encoded to cautious reasoning from the FLP answer set of a disjunctive logic program, which is  $\Pi_2^p$ -complete [7].

Note that some world view exists iff some candidate world view exists; this justifies the second part of Theorem 7, and consistency checking under candidate world view and world view semantics are equally hard. On the other hand, under EFLP semantics formula evaluation is harder.

**Theorem 8** *Given a propositional program  $\Pi$  and a propositional formula  $F$ , deciding whether  $F$  is true in  $\Pi$  (i) w.r.t. candidate world views is  $\Sigma_3^p$ -complete and (ii) w.r.t. world views, i.e., under EFLP semantics, is  $\Sigma_4^p$ -complete.*

While under candidate world views, it suffices to guess  $\Phi \subseteq Ep(\Pi)$  such that  $\Pi$  has a candidate world view  $\mathcal{A}$  w.r.t.  $\Phi$  and  $\Pi^\Phi$  cautiously entails  $F$  (which can be checked with an  $\Sigma_2^p$  oracle), in case of EFLP semantics, in addition maximality of  $\Phi$  must be tested, i.e., no  $\Phi' \supset \Phi$  has some candidate world view; this however turns out to be  $\Pi_3^p$ -complete, which lifts the problem to the fourth level of the polynomial hierarchy.

We note that the above results hold for propositional programs that amount to epistemic specifications, i.e., the rules match the syntax of form (1), where  $\mathbf{K}F$  amounts to  $\neg \mathbf{not} F$  and  $\mathbf{M}F$  to  $\mathbf{not} \neg F$ ; in case of normal rules, i.e.,  $m = 1$ , the complexity drops by one level of the polynomial hierarchy, and we obtain  $D_1^p$ -,  $\Sigma_2^p$  and  $\Sigma_3^p$ -completeness in place of  $D_2^p$ -,  $\Sigma_3^p$  and  $\Sigma_4^p$ -completeness, respectively (for details, see B). Indeed, as  $\neg\neg A = A$  under FLP semantics, for normal specifications the reduct  $\Pi^\Phi$  amounts to a normal logic program, for which deciding FLP answer set existence is NP-complete. Furthermore, we remark that similar complexity results can be derived for other answer set semantics, and in particular for well-justified FLP answer sets [28].

## 6 Related Work

The need for using epistemic negation in knowledge representation and reasoning was long recognized by [15] and recently further emphasized in [16, 32, 10, 33, 18, 6]. In particular, [15] first introduced modal operators  $\mathbf{K}$  and  $\mathbf{M}$  to ASP and interpreted them in three-valued interpretations (called *three-valued possible worlds*). Formulas with such modal operators are called *strongly introspective* and logic programs with rules of form (1) called *epistemic specifications*. Truszczyński [32] revisited this formalism and redefined its semantics in two-valued interpretations.

It turns out that the approaches of [15, 32] suffer from epistemic circular justifications in logic programs like  $\Pi = \{p \leftarrow \mathbf{K}p\}$ , where an undesired world view  $\{\{p\}\}$  is produced. To remedy this, [16] further updated his program transformation. However, the epistemic circular justification problem persists in other logic programs, such as  $\Pi = \{q \leftarrow \neg \mathbf{K}p, p \leftarrow \neg q\}$ , where an undesired

world view  $\{\{p\}\}$  is produced when applying the approaches of [16, 15, 32]. In addition, these approaches also suffer from the multiple world view problem in logic programs like  $\Pi = \{p \leftarrow \mathbf{M}p\}$ , where two world views,  $\{\{p\}\}$  and  $\{\emptyset\}$ , are produced.

To address both circular justifications and the multiple world view problem, [18] proposed a more involved program transformation by appealing to nested expressions of [19]. However this approach has some clear shortcomings as listed in the Introduction. Specifically, as shown in [28] logic programs with nested expressions suffer from circular justifications under the existing semantics such as [19, 12, 11]. Moreover, as illustrated in Example 1, this approach also suffers from the multiple world view problem.

In an alternative venue, [34] developed *epistemic equilibrium models* for epistemic specifications by introducing modal operators  $\mathbf{K}$  and  $\mathbf{M}$  to Pearce’s equilibrium logic [25, 24]. This approach suffers from epistemic circular justifications and the multiple world view problem; for example, both  $\{\{p\}\}$  and  $\{\emptyset\}$  are world views of  $\Pi = \{p \leftarrow \mathbf{K}p\}$  and  $\Pi = \{p \leftarrow \mathbf{M}p\}$ . Very recently, [6] presented a new definition of epistemic equilibrium models (EEMs) and further defined *autoepistemic equilibrium models (AEEMs)* as world views that are maximal EEMs under set inclusion and a special partial preorder over S5 models. A dozen of logic programs were listed all of which except for that program in Example 1 have the AEEMs that coincide with the world views of [18]; the AEEM of the program in Example 1 coincides with the world view under our EFLP semantics. On the one hand, it is unclear whether AEEMs are free of epistemic circular justifications and the multiple world view problem for general logic programs. On the other hand, as discussed in [28], Pearce’s equilibrium semantics coincides with the answer set semantics of [12] for propositional logic programs and with that of [11] in the first-order case, and these semantics suffer from circular justifications. Since AEEMs are epistemic extensions of Pearce’s equilibrium models, circular justifications inevitably convey to AEEMs, thus sometimes leading to undesired world views. For example, for the program  $\Pi = \{p \leftarrow \neg\neg p, p \leftarrow \neg p\}$ ,  $\{p\}$  is a unique minimal equilibrium model/answer set under [24] and thus  $\{\{p\}\}$  is an AEEM under [6]. This answer set/world view is undesired because it has a circular justification via the self-supporting loop  $p \leftarrow \neg\neg p \leftarrow p$ . Finally, this approach may miss some desired world views; for instance,  $\{\{p\}\}$  is expected to be a world view of the program  $\Pi = \{p \leftarrow \mathbf{not} p \vee p\}$  (see Example 7), but it is not an AEEM of  $\Pi$  where  $\mathbf{not} p$  is replaced by  $\neg\mathbf{K}p$ .

Finally, we mention that [20] defined a modal logic of *Minimal Knowledge and Negation as Failure* (MKNF), which has two modal operators, viz.  $\mathbf{K}$  as defined above and  $\mathbf{not}$  like our epistemic negation operator. MKNF logic has recently been exploited for the integration of description logics and rules in the Semantic Web [23]. As indicated in the end of [20], “MKNF does not cover the important concept of strong introspection introduced in [15].” Thus, applying it to handle epistemic negation would yield unintuitive results. For instance, the logic program in Example 2 will be identified with  $P_1$  consisting of the following formulas:

$$\begin{array}{ll} \mathbf{K} \textit{innocent}(\textit{John}) \vee \mathbf{K} \textit{guilty}(\textit{John}) & f_1 \\ \mathbf{not} \textit{guilty}(X) \supset \mathbf{K} \textit{innocent}(X) & f_2 \end{array}$$

Under MKNF,  $P_1$  has two possible collections of models (where each is viewed as an S5 structure):

$$\mathcal{A}_1 = \{\{\textit{innocent}(\textit{John})\}, \{\textit{innocent}(\textit{John}), \textit{guilty}(\textit{John})\}\}$$

$$A_2 = \{\{guilty(John)\}, \{guilty(John), innocent(John)\}\}$$

The first collection means *John* is innocent, while the second says *John* is guilty, which violates our intuition.

## 7 Summary and Future Work

We have presented a novel approach to evaluating epistemic negation; specifically, we introduced a novel program transformation based on epistemic negations and presented a new definition of epistemic answer set semantics for general logic programs which minimizes the knowledge in world views by requiring that every world view should satisfy as many epistemic negations as possible. Thus this approach overcomes both the epistemic circular justification problem and the multiple world view problem, and provides an appropriate solution to epistemic specifications introduced by [15].

We considered general logic programs consisting of rules of the form  $H \leftarrow B$ , where  $H$  and  $B$  are arbitrary first-order formulas possibly containing epistemic negation, and defined epistemic FLP answer set semantics for general logic programs. The proposed approach can readily be adapted to any other existing answer set semantics for extension with epistemic negation, such as those in [24, 26, 31, 3, 11, 28].

We showed that for a propositional program  $\Pi$  with epistemic negation, deciding whether  $\Pi$  has EFLP answer sets is  $\Sigma_3^p$ -complete and deciding whether a propositional formula  $F$  is true in  $\Pi$  under EFLP semantics is  $\Sigma_4^p$ -complete in general, but has lower complexity for important program classes.

As ongoing work, we are considering an implementation of EFLP semantics on some state-of-the-art ASP solvers, such as DLVHEX or CLASP.<sup>8</sup> Specifically we are interested in implementing the well-justified FLP semantics [28] extended with epistemic negation. The well-justified FLP semantics enhances FLP semantics with a level mapping mechanism such that every answer set of a general logic program has a level mapping and thus is free of circular justifications.

## References

- [1] M. Balduccini and T. Cao Son, editors. *Logic Programming, Knowledge Representation, and Nonmonotonic Reasoning - Essays Dedicated to Michael Gelfond on the Occasion of His 65th Birthday*, volume 6565 of *Lecture Notes in Computer Science*. Springer, 2011.
- [2] C. Baral, G. Greco, N. Leone, and G. Terracina, editors. *Logic Programming and Nonmonotonic Reasoning, 8th International Conference, LPNMR 2005, Diamante, Italy, September 5-8, 2005, Proceedings*, volume 3662 of *Lecture Notes in Computer Science*. Springer, 2005.

---

<sup>8</sup><http://www.cs.uni-potsdam.de/clasp>

- [3] M. Bartholomew, J. Lee, and Y. Meng. First-order extension of the FLP stable model semantics via modified circumscription. In *Proc. 22nd Int'l Joint Conference on Artificial Intelligence (IJCAI-11)*, pages 724–730, 2011.
- [4] G. Brewka, T. Eiter, and M. Truszczyński. Answer set programming at a glance. *Communications of the ACM*, 54(12):92–103, 2011.
- [5] J. de Bruijn, T. Eiter, and H. Tompits. Embedding approaches to combining rules and ontologies into autoepistemic logic. In *Proc. 11th International Conference on Principles of Knowledge Representation and Reasoning (KR 2008)*, pages 485–495. AAAI Press, 2008.
- [6] L. F. del Cerro, A. Herzig, and Ezgi Iraz Su. Epistemic equilibrium logic. In *Proc. 24th Int'l Joint Conference on Artificial Intelligence (IJCAI-15)*, pages 2964–2970. AAAI Press/IJCAI, 2015.
- [7] T. Eiter and G. Gottlob. On the computational cost of disjunctive logic programming: Propositional case. *Annals of Mathematics and Artificial Intelligence*, 15(3-4):289–323, 1995.
- [8] T. Eiter, G. Ianni, T. Lukasiewicz, R. Schindlauer, and H. Tompits. Combining answer set programming with description logics for the semantic web. *Artificial Intelligence*, 172(12-13):1495–1539, 2008.
- [9] W. Faber, G. Pfeifer, and N. Leone. Semantics and complexity of recursive aggregates in answer set programming. *Artificial Intelligence*, 175(1):278–298, 2011.
- [10] W. Faber and S. Woltran. Manifold answer-set programs and their applications. In Balduccini and Son [1], pages 44–63.
- [11] P. Ferraris, J. Lee, and V. Lifschitz. Stable models and circumscription. *Artificial Intelligence*, 175(1):236–263, 2011.
- [12] P. Ferraris. Answer sets for propositional theories. In Baral et al. [2], pages 119–131.
- [13] P. Ferraris and V. Lifschitz. Mathematical foundations of answer set programming. In Sergei N. Artëmov, Howard Barringer, Artur S. d'Avila Garcez, Luís C. Lamb, and John Woods, editors, *We Will Show Them! Essays in Honour of Dov Gabbay, Volume One*, pages 615–664. College Publications, 2005.
- [14] M. Gelfond and V. Lifschitz. Classical negation in logic programs and disjunctive databases. *New Generation Computing*, 9:365–385, 1991.
- [15] M. Gelfond. Strong introspection. In Thomas L. Dean and Kathleen McKeown, editors, *Proceedings of the 9th National Conference on Artificial Intelligence, Anaheim, CA, USA, July 14-19, 1991, Volume 1.*, pages 386–391. AAAI Press / The MIT Press, 1991.

- [16] M. Gelfond. New semantics for epistemic specifications. In James P. Delgrande and Wolfgang Faber, editors, *Logic Programming and Nonmonotonic Reasoning - 11th International Conference, LPNMR 2011, Vancouver, Canada, May 16-19, 2011. Proceedings*, volume 6645 of *Lecture Notes in Computer Science*, pages 260–265. Springer, 2011.
- [17] M. Gelfond and V. Lifschitz. The stable model semantics for logic programming. In Robert A. Kowalski and Kenneth A. Bowen, editors, *Logic Programming, Proceedings of the Fifth International Conference and Symposium, Seattle, Washington, August 15-19, 1988 (2 Volumes)*, pages 1070–1080. MIT Press, 1988.
- [18] P. T. Kahl. *REFINING THE SEMANTICS FOR EPISTEMIC LOGIC PROGRAMS*. PhD thesis, Texas Tech University, USA, 2014.
- [19] V. Lifschitz, L. R. Tang, and H. Turner. Nested expressions in logic programs. *Annals of Mathematics and Artificial Intelligence*, 25(1-2):369–389, 1999.
- [20] V. Lifschitz. Nonmonotonic databases and epistemic queries. In John Mylopoulos and Raymond Reiter, editors, *Proceedings of the 12th International Joint Conference on Artificial Intelligence. Sydney, Australia, August 24-30, 1991*, pages 381–386. Morgan Kaufmann, 1991.
- [21] V. Lifschitz. Thirteen definitions of a stable model. In Andreas Blass, Nachum Dershowitz, and Wolfgang Reisig, editors, *Fields of Logic and Computation, Essays Dedicated to Yuri Gurevich on the Occasion of His 70th Birthday*, volume 6300 of *Lecture Notes in Computer Science*, pages 488–503. Springer, 2010.
- [22] W. Marek and M. Truszczyński. Autoepistemic logic. *Journal of the ACM*, 38(3):588–619, 1991.
- [23] B. Motik and R. Rosati. Reconciling description logics and rules. *J. ACM*, 57(5), 2010.
- [24] D. Pearce. Equilibrium logic. *Annals of Mathematics and Artificial Intelligence*, 47(1-2):3–41, 2006.
- [25] D. Pearce. A new logical characterisation of stable models and answer sets. In Jürgen Dix, Luís Moniz Pereira, and Teodor C. Przymusiński, editors, *Non-Monotonic Extensions of Logic Programming, NMELP '96, Bad Honnef, Germany, September 5-6, 1996, Selected Papers*, volume 1216 of *Lecture Notes in Computer Science*, pages 57–70. Springer, 1996.
- [26] W. Pelov, M. Denecker, and M. Bruynooghe. Well-founded and stable semantics of logic programs with aggregates. *Theory and Practice of Logic Programming*, 7(3):301–353, 2007.
- [27] R. Reiter. On closed world data bases. In H. Gallaire and J. Minker, eds., *Logic and Data Bases*, pages 119–140. Plenum, New York, 1978.
- [28] Y. D. Shen, K. Wang, T. Eiter, M. Fink, C. Redl, T. Krennwallner, and J. Deng. FLP answer set semantics without circular justifications for general logic programs. *Artificial Intelligence*, 213:1–41, 2014.



- [29] T. C. Son, E. Pontelli, and P. H. Tu. Answer sets for logic programs with arbitrary abstract constraint atoms. *Journal of Artificial Intelligence Research*, 29:353–389, 2007.
- [30] E. Sperner. Ein Satz über Untermengen einer endlichen Menge. *Mathematische Zeitschrift*, XXVII:544–548, 1928.
- [31] M. Truszczyński. Reducts of propositional theories, satisfiability relations, and generalizations of semantics of logic programs. *Artificial Intelligence*, 174(16-17):1285–1306, 2010.
- [32] M. Truszczyński. Revisiting epistemic specifications. In Balduccini and Son [1], pages 315–333.
- [33] H. Vlaeminck, J. Vennekens, M. Bruynooghe, and M. Denecker. Ordered epistemic logic: Semantics, complexity and applications. In Gerd Brewka, Thomas Eiter, and Sheila McIlraith, editors, *Principles of the Thirteenth International Conference on Knowledge Representation and Reasoning: Proceedings of (KR 2012), Rome, Italy, June 10-14, 2012*, pages 369–379, 2012.
- [34] K. Wang and Y. Zhang. Nested epistemic logic programs. In Baral et al. [2], pages 279–290.

## A Proofs

### *Proof of Theorem 2.*

Let  $\mathcal{A} = \bigcup_i \mathcal{A}_i$  be the union of all collections  $\mathcal{A}_i$  of models of  $\Pi$ ; then  $\mathcal{A}$  consists of all models of  $\Pi$ . By the condition that for every epistemic negation  $\text{not } F$  in  $\text{ground}(\Pi)$ ,  $F$  is true in every model of  $\Pi$ , for every interpretation  $I \in \mathcal{A}$ , as  $I$  satisfies all rules in  $\text{ground}(\Pi)$  w.r.t. some  $\mathcal{A}_i$ ,  $I$  also satisfies these rules w.r.t.  $\mathcal{A}$ . This means  $\mathcal{A}$  itself is a collection of models of  $\Pi$ .

For any interpretation  $I$ , if  $\mathcal{A}_I = \{I\}$  is a collection of models of  $\Pi$ , then  $I \in \mathcal{A}$ ; conversely, if  $I \in \mathcal{A}$ , i.e.,  $I$  satisfies all rules in  $\text{ground}(\Pi)$  w.r.t.  $\mathcal{A}$ , then  $I$  also satisfies these rules w.r.t.  $\mathcal{A}_I = \{I\}$  and thus  $\mathcal{A}_I = \{I\}$  is a collection of models of  $\Pi$ . Hence for any interpretation  $I$ ,  $\mathcal{A}_I = \{I\}$  is a collection of models of  $\Pi$  iff  $I \in \mathcal{A}$  iff  $I$  is a model of  $\Pi$ . Moreover, by Proposition 1, for any interpretation  $I$ ,  $\mathcal{A}_I = \{I\}$  is a collection of models of  $\Pi$  iff  $I$  is a model of  $\Pi^\neg$ . Thus  $I$  is a model of  $\Pi$  iff  $I$  is a model of  $\Pi^\neg$ . This concludes the proof.  $\square$

### *Proof of Proposition 3.*

Consider an EFLP answer set  $I \in \mathcal{A}$  w.r.t. guess  $\Phi$ . As  $I$  is an FLP answer set of the epistemic reduct  $\Pi^\Phi$ , it is a model of  $\Pi^\Phi$ . Let  $\Pi^{\Phi^{\text{not}}}$  be  $\Pi^\Phi$  with all  $\neg F$  replaced by  $\text{not } F$ . Since for every  $\text{not } F$  in  $\Pi^{\Phi^{\text{not}}}$ ,  $F$  is true in all  $I \in \mathcal{A}$ , by Theorem 2,  $\Pi^{\Phi^{\text{not}}}$  and  $\Pi^\Phi$  have the same models. This means  $I$  is a model of  $\Pi^{\Phi^{\text{not}}}$  w.r.t.  $\mathcal{A}$ . Note that  $\Pi^{\Phi^{\text{not}}}$  is  $\text{ground}(\Pi)$  with every  $\text{not } F \in \Phi$  replaced by  $\top$ . Then  $r$  is a rule in  $\text{ground}(\Pi)$  iff  $r^\top$  is a rule in  $\Pi^{\Phi^{\text{not}}}$ , where  $r^\top$  is  $r$  with every  $\text{not } F \in \Phi$  replaced by  $\top$ . Since for every  $\text{not } F \in \Phi$ ,  $F$  is false in some  $J \in \mathcal{A}$ ,  $I$  satisfies  $r$  w.r.t.  $\mathcal{A}$  iff  $I$  satisfies  $r^\top$  w.r.t.  $\mathcal{A}$ . This shows  $I$  is also a model of  $\text{ground}(\Pi)$  w.r.t.  $\mathcal{A}$ . Hence  $\mathcal{A}$  is a collection of models of  $\Pi$ .  $\square$

*Proof of Theorem 4.*

(1) Let  $\mathcal{A}$  be a world view of  $\Pi$  w.r.t. a guess  $\Phi$ . Then  $\mathcal{A}$  is also the set of all FLP answer sets of the epistemic reduct  $\Pi^\Phi$ . Rules in  $\Pi^\Phi$  are of the form

$$A_0 \leftarrow A_1 \wedge \dots \wedge A_m \wedge \neg B_1 \wedge \dots \wedge \neg B_n$$

where every  $B_i$  is true in every  $I \in \mathcal{A}$  because  $\mathbf{not} B_i \in Ep(\Pi) \setminus \Phi$ . This means  $\Pi^\Phi$  and  $\Pi^{\Phi^+}$  have the same set of FLP answer sets, where  $\Pi^{\Phi^+}$  is  $\Pi^\Phi$  with all rules containing a negative literal  $\neg B_i$  removed. Since  $\Pi^{\Phi^+}$  is a positive logic program, it has a unique FLP answer set. Hence  $\mathcal{A}$  has only one FLP answer set, i.e.,  $|\mathcal{A}| = 1$ .

(2) For normal logic programs, FLP semantics coincides with the standard answer set semantics. So it suffices to show that  $\mathcal{A} = \{I\}$  is a world view of  $\Pi$  iff  $I$  is a standard answer set of  $\Pi^\neg$ .

We first show that  $\mathcal{A} = \{I\}$  is a candidate world view of  $\Pi$  iff  $I$  is a standard answer set of  $\Pi^\neg$ . ( $\implies$ ) Assume  $\mathcal{A} = \{I\}$  is a candidate world view of  $\Pi$  w.r.t. a guess  $\Phi$ . Then  $I$  is also the FLP answer set of the epistemic reduct  $\Pi^\Phi$ , and for every  $\mathbf{not} B$  in  $ground(\Pi)$ ,  $\mathbf{not} B \in \Phi$  iff  $B$  is false in  $I$  iff  $B \notin I$ . Let  $r$  be a rule in  $\Pi^\Phi$ . For every negative literal  $\neg B_i$  in the rule body of  $r$ ,  $B_i$  is true in  $I$ , i.e.,  $B_i \in I$ . Let  $\Pi^{\Phi^+}$  be  $\Pi^\Phi$  with all rules containing a negative literal  $\neg B_i$  removed. Then  $I$  is also the FLP answer set of  $\Pi^{\Phi^+}$ , which is the least model of  $\Pi^{\Phi^+}$ . Note that  $\Pi^{\Phi^+}$  is in fact the GL-reduct of  $\Pi^\neg$ . This means  $I$  is also a standard answer set of  $\Pi^\neg$ .

( $\impliedby$ ) Assume  $I$  is a standard answer set of  $\Pi^\neg$ . Let  $\Phi$  be a guess such that for every  $\mathbf{not} B$  in  $ground(\Pi)$ ,  $\mathbf{not} B \in \Phi$  iff  $B$  is false in  $I$  iff  $B \notin I$ . Let  $r$  be a rule in  $ground(\Pi)$  such that  $I$  satisfies  $body(r)$  w.r.t.  $\mathcal{A} = \{I\}$ .  $r$  must be of the form

$$H \leftarrow A_1 \wedge \dots \wedge A_m \wedge \mathbf{not} B_1 \wedge \dots \wedge \mathbf{not} B_n$$

with every  $A_i \in I$  and every  $B_i \notin I$  (i.e.,  $\mathbf{not} B_i \in \Phi$ ). Then  $\Pi^\neg$  must have a rule  $r'$  of the form

$$H \leftarrow A_1 \wedge \dots \wedge A_m \wedge \neg B_1 \wedge \dots \wedge \neg B_n$$

Note that  $I$  satisfies  $body(r')$ . As  $I$  is a model of  $\Pi^\neg$ ,  $H$  is in  $I$ . This means  $I$  satisfies  $r$  w.r.t.  $\mathcal{A}$ . Hence  $\mathcal{A} = \{I\}$  is a candidate world view w.r.t.  $\Phi$ .

We have proved that  $\mathcal{A} = \{I\}$  is a candidate world view of  $\Pi$  iff  $I$  is a standard answer set of  $\Pi^\neg$ . To show that  $\mathcal{A} = \{I\}$  is a world view of  $\Pi$  iff  $I$  is a standard answer set of  $\Pi^\neg$ , it suffices to show that every candidate world view of  $\Pi$  is a world view of  $\Pi$ .

Assume on the contrary that there are two guesses  $\Phi' \supset \Phi$  such that  $\mathcal{A}' = \{I'\}$  and  $\mathcal{A} = \{I\}$  are candidate world views w.r.t.  $\Phi'$  and  $\Phi$ , respectively. For each  $\mathbf{not} B \in \Phi' \setminus \Phi$ ,  $B \notin I'$  and  $B \in I$ . This means  $I' \neq I$ . By the above proof,  $I'$  and  $I$  are standard answer sets of  $\Pi^\neg$ . We have the following two GL-reducts w.r.t.  $I'$  and  $I$  respectively:

$$(\Pi^\neg)^{I'} = \{A_0 \leftarrow A_1 \wedge \dots \wedge A_m \mid A_0 \leftarrow A_1 \wedge \dots \wedge A_m \wedge \neg B_1 \wedge \dots \wedge \neg B_n \\ \text{is a rule in } ground(\Pi^\neg) \text{ and for every } B_i, \mathbf{not} B_i \text{ is in } \Phi'\},$$

$$(\Pi^\neg)^I = \{A_0 \leftarrow A_1 \wedge \dots \wedge A_m \mid A_0 \leftarrow A_1 \wedge \dots \wedge A_m \wedge \neg B_1 \wedge \dots \wedge \neg B_n \\ \text{is a rule in } ground(\Pi^\neg) \text{ and for every } B_i, \mathbf{not} B_i \text{ is in } \Phi\}.$$

As the two GL-reducts are positive logic programs,  $I'$  and  $I$  are their least fixpoints, respectively. As  $\Phi' \supset \Phi$ , every rule in  $(\Pi^\neg)^I$  must be in  $(\Pi^\neg)^{I'}$ . This implies  $I \subset I'$  ( $I' \neq I$ , as shown above),

which contradicts the fact that every standard answer set of a normal logic program is a minimal model of the program. We then conclude the proof.  $\square$

*Proof of Theorem 6.*

**Membership.** The proof of  $D_2^p$  membership is in the discussion.

**Hardness.** The  $D_2^p$ -hardness is shown by the reduction of the  $D_2^p$ -complete problem where, given a pair  $(\Pi_1, \Pi_2)$  of programs, we must decide whether  $\Pi_1$  has some FLP answer set and  $\Pi_2$  has no FLP answer set. (This result in turn is easily obtained by a reduction of evaluating QBFs of the form  $\exists X \forall Y \phi$  to FLP answer set existence of disjunctive logic programs [7].)

Assume w.l.o.g. that  $\Pi_1$  and  $\Pi_2$  are on disjoint signatures and let  $A, \bar{A}, C$  be fresh atoms. Then we claim that

$$\begin{aligned} \Pi = \Pi_1 \cup \{ & H \leftarrow B \wedge \neg A \mid H \leftarrow B \in \Pi_2 \} \\ & \cup \{ A \leftarrow \neg \bar{A}, \bar{A} \leftarrow \neg A, C \leftarrow \mathbf{not} A \} \end{aligned}$$

has a candidate world view w.r.t.  $\Phi = \emptyset$  iff  $\Pi_1$  has some FLP answer set and  $\Pi_2$  has no FLP answer set. Notice that the epistemic reduct is

$$\Pi^\Phi = (\Pi \setminus \{C \leftarrow \mathbf{not} A\}) \cup \{C \leftarrow \neg A\}.$$

( $\implies$ ) Suppose  $\Pi$  has a candidate world view  $\mathcal{A}$  w.r.t.  $\Phi = \emptyset$ . Then  $\Pi_1$  must have an FLP answer set. Furthermore,  $A$  must be true in every FLP answer set of  $\Pi^\Phi$ . In particular, this means that the guess  $\bar{A}$  from  $A \leftarrow \bar{A}, \bar{A} \leftarrow A$ , does not lead to an FLP answer set of  $\Pi$ ; this means that the program  $\{H \leftarrow B \wedge \neg A \mid H \leftarrow B \in \Pi_2\}$  has no FLP answer set, from which in turn it follows that  $\Pi_2$  has no FLP answer set.

( $\impliedby$ ) For each FLP answer set  $I_1$  of  $\Pi_1$ , the set  $I_1 \cup \{A\}$  is an FLP answer set of  $\Pi^\Phi$ , and if since  $\Pi_2$  has no FLP answer set, each answer set in the collection  $\mathcal{A}$  of FLP answer sets of  $\Pi^\Phi$  has this form. As  $\Pi_1$  has some FLP answer set  $I_1$ , it follows that  $\mathcal{A} \neq \emptyset$  and  $\mathbf{not} A$  is false in  $\mathcal{A}$ ; hence  $\mathcal{A}$  is a candidate world view of  $\Pi$  w.r.t.  $\Phi = \emptyset$   $\square$

*Proof of Theorem 7.*

(i) **Membership.** A nondeterministic Turing machine can make a guess  $\Phi \subseteq EP(\Pi)$  such that  $\Pi$  has some candidate world view w.r.t.  $\Phi$  in polynomial time, and then continue to check that  $\Phi$  is indeed proper with an oracle for  $D_2^p$  in polynomial time, by Theorem 6. As each problem in  $D_2^p$  can be decided with two calls of an  $\Sigma_2^p$  oracle, it follows that the problem is in  $\Sigma_3^p$ .

**Hardness.** The  $\Sigma_3^p$ -hardness of deciding candidate world view existence can be shown by a reduction from evaluating QBFs of the form

$$\exists X \forall Y \exists Z \phi, \tag{3}$$

where  $X = X_1, \dots, X_{n_X}$ ,  $Y = Y_1, \dots, Y_{n_Y}$ , and  $Z = Z_1, \dots, Z_{n_Z}$  are lists (viewed as sets) of distinct atoms, and  $\phi = \bigwedge_{j=1}^k (L_{j,1} \vee L_{j,2} \vee L_{j,3})$  is a CNF over atoms  $X \cup Y \cup Z$ , by lifting a reduction in [7] that proved  $\Pi_2^p$ -completeness of deciding whether an atom  $U$  is false in all answer sets of a disjunctive logic program. Without loss of generality, we assume that by assigning each  $Y_i \in Y$  the value true, the formula  $\phi$  evaluates to true for every assignment to the remaining variables, i.e.,

the formula  $\forall X, Z \phi[Y/\top]$  evaluates to true; hence, regardless of a concrete assignment  $\sigma$  to  $X$  the formula  $\phi[X/\sigma, Y/\top]$  will be satisfiable.

For each atom  $A \in X \cup Y \cup Z$ , we introduce a fresh atom  $\bar{A}$ , and we use further fresh atoms  $U$  and  $V$ . Let

$$\Pi_X = \{X_i \leftarrow \mathbf{not} \bar{X}_i; \bar{X}_i \leftarrow \mathbf{not} X_i \mid X_i \in X\}, \quad (4)$$

$$\Pi_Y = \{Y_i \leftarrow \neg \bar{Y}_i; \bar{Y}_i \leftarrow \neg Y_i \mid Y_i \in Y\}, \quad (5)$$

$$\Pi_Z = \{Z_i \vee \bar{Z}_i \leftarrow \mid Z_i \in Z\}, \quad (6)$$

$$\Pi_\phi = \{U \leftarrow L_{j,1}^* \wedge L_{j,3}^* \wedge L_{j,3}^* \mid 1 \leq j \leq k\} \cup \quad (7)$$

$$\{Z_i \leftarrow U; \bar{Z}_i \leftarrow U \mid Z_i \in Z\} \cup \quad (8)$$

$$\{V \leftarrow \mathbf{not} V, \mathbf{not} \neg U\} \quad (9)$$

where the operation  $*$  converts each positive literal  $A$  into  $\bar{A}$  and each negative literal  $\neg A$  into  $A$ . If we suppose that  $X$  is void, the program  $\Pi_Y \cup \Pi_Z \cup \Pi_\phi \setminus \{(9)\}$  amounts to a (variant of) the program  $\Pi_\Phi$  in [7] for evaluating  $\Psi = \forall Y \exists Z \phi$  such that  $U$  is false in all answer sets of  $\Pi_\Psi$  iff  $\Psi$  evaluates to true; for each assignment  $\mu$  to  $Y$ , which is guessed by the rules (5) the program  $\Pi_\Psi$  has some answer set candidate  $I_\mu = \{Y_i \in Y \mid \mu(Y_i) = \mathit{true}\} \cup \{\bar{Y}_i \in Y \mid \mu(Y_i) = \mathit{false}\} \cup Z \cup \bar{Z} \cup \{U\}$ , where  $\bar{S} = \{\bar{A} \mid A \in S\}$ , that is an answer set if  $\phi[Y/\sigma]$  is unsatisfiable; otherwise, if  $\phi[Y/\mu]$  is satisfiable, some answer set  $I \subset I_\mu$  exists and  $U$  is false in each such  $I$ . By our assumption, for  $\mu = \top$  (i.e.,  $\mu(Y_i) = \mathit{true}$  for every  $Y_i \in Y$ ) such an answer set exists, thus  $\Pi_\Psi$  has some answer set.

Now let us consider the case where  $X$  is non-void, and let  $\Pi = \Pi_X \cup \Pi_Y \cup \Pi_Z \cup \Pi_\phi$ . Intuitively, the rules in  $\Pi_X$  guess an assignment  $\sigma$  to  $X$ , and the candidate world view conditions amount to evaluating the QBF  $\forall Y \exists Z \phi[X/\sigma]$ . Here, the rule (9) plays a crucial role; informally, it checks that in all answer sets of  $(\Pi \setminus \{(9)\})^\Phi$  the atom  $U$  is false, without pruning answer sets in which  $U$  is true (note that a constraint  $\perp \leftarrow \mathbf{not} \neg U$  or  $V \leftarrow \neg V, \mathbf{not} \neg U$  would not work). Indeed, the rule (9) can be satisfied only by a guess  $\Phi$  such that  $\Phi \cap \{\mathbf{not} V, \mathbf{not} \neg U\} = \{\mathbf{not} V\}$ , which means that  $V \leftarrow \neg \neg U$  is in  $\Pi^\Phi$ ; the candidate world conditions (a)-(c) are then met iff  $U$  is false in every answer set of  $(\Pi \setminus \{(9)\})^\Phi$ . Otherwise, if  $\mathbf{not} V \notin \Phi$ , we have a rule  $V \leftarrow \neg V, \dots$  in  $\Pi^\Phi$ , and as  $V$  occurs in no other rule head,  $\Pi^\Phi$  has no answer set in which  $V$  is true, and thus  $\Phi$  violates condition (a) or (c); if  $\mathbf{not} \neg U, \mathbf{not} V \in \Phi$ , then  $V \leftarrow$  is in  $\Pi^\Phi$ , and hence  $\Phi$  violates condition (b).

Having clarified the working of (9) for guesses  $\Phi$  that have some candidate world views, we can verify that any such guess encodes a truth assignment  $\sigma$  to  $X$ . For  $\mathbf{not} \bar{X}_i$  and  $\mathbf{not} X_i$ , we have four possible cases for  $\Phi_j = \Phi \cap \{\mathbf{not} \bar{X}_i, \mathbf{not} X_i\}$ : 1.  $\Phi_1 = \emptyset$ , 2.  $\Phi_2 = \{\mathbf{not} \bar{X}_i, \mathbf{not} X_i\}$ , 3.  $\Phi_3 = \{\mathbf{not} \bar{X}_i\}$ , and 4.  $\Phi_4 = \{\mathbf{not} X_i\}$ . Of these,  $\Phi_1$  means that condition (c) is violated, as the program  $(\Pi \setminus \{(9)\})^\Phi$  will contain rules

$$X_i \leftarrow \neg \bar{X}_i; \bar{X}_i \leftarrow \neg X_i,$$

i.e., a choice between  $X_i$  and  $\bar{X}_i$ ; as these atoms do not occur in other rule heads, our assumption about  $\phi$  implies that  $(\Pi \setminus \{(9)\})^\Phi$  has answer sets in which  $X_i$  resp.  $\bar{X}_i$  is false. Likewise,  $\Phi_2$

does not yield a candidate world view:  $\Pi^\Phi$  contains the facts  $X_i$  and  $\bar{X}_i$ , and thus condition (b) is violated. Thus, only  $\Phi_3$  and  $\Phi_4$  remain. In case of  $\Phi_3$ , the reduct  $\Pi^\Phi$  contains

$$X_i \leftarrow ; \quad \bar{X}_i \leftarrow \neg X_i.$$

and thus each FLP answer set of  $\Pi^\Phi$  must contain  $X_i$  but not  $\bar{X}_i$ ; that is,  $\Phi_3$  encodes that  $X_i$  is true in it. Similarly, for  $\Phi_4$  each FLP answer set of  $\Pi^\Phi$  must contain  $\bar{X}_i$  but not  $X_i$ ; that is  $\Phi_4$  encodes that  $X_i$  is false.

Putting things together, if  $\Phi$  is guess that has some world view, then  $\Phi$  encodes a truth assignment  $\sigma$  to  $X$  such that the formula  $\forall Y \exists Z \phi[X/\sigma]$  evaluates to true; that is,  $\exists X \forall Y \exists Z \phi$  evaluates to true. Conversely, if  $\sigma$  is a truth assignment to  $X$  such that  $\forall Y \exists Z \phi[X/\sigma]$  evaluates to true, then the guess  $\Phi = \{\mathbf{not} \bar{X}_i \mid \sigma(X_i) = \text{true}\} \cup \{\mathbf{not} X_i \mid \sigma(X_i) = \text{false}\} \cup \{\mathbf{not} \neg U\}$  has some candidate world view. As  $\Pi$  is clearly constructible in polynomial time from (3), this proves  $\Sigma_3^p$ -hardness of deciding candidate world view existence; furthermore, note that the rules in  $\Pi$  match the syntax of (1).

(ii) Follows from (i), as some world view exists iff some candidate world view exists.  $\square$

### *Proof of Theorem 8*

(i) Membership. As in the discussion, it suffices to guess  $\Phi \subseteq EP(\Pi)$  such that  $\Pi$  has a candidate world view  $\mathcal{A}$  w.r.t.  $\Phi$  and  $\Pi^\Phi$  cautiously entails  $F$ . From Theorem 6 and [28] it follows that the guess can be verified with an  $\Sigma_2^p$  oracle in polynomial time; this proves membership in  $\Sigma_3^p$ .

Hardness.  $\Sigma_3^p$ -hardness is a trivial from Theorem 7: a given program  $\Pi$  has some candidate world view iff  $A$  is true in  $\Pi \cup \{A \leftarrow\}$  w.r.t. candidate world views, i.e., in some candidate world view of  $\Pi$ , where  $A$  is a fresh atom.

(ii) Membership. For a guess  $\Phi$  as in (i), we must in addition test that no  $\Phi' \supset \Phi$  exists that has a candidate world view; by Theorem 7, the latter is in  $\Pi_3^p$ . Thus in summary, a guess  $\Phi$  that has a world view  $\mathcal{A}$  in which  $F$  is true can be verified with a  $\Sigma_3^p$  oracle in polynomial; this proves  $\Sigma_4^p$  membership.

Hardness.  $\Sigma_4^p$ -hardness is shown by generalizing the reduction in Theorem 7 to encode the evaluation of a QBF

$$\exists W \forall X \exists Y \forall Z \psi, \tag{10}$$

where  $\psi = \bigvee_{j=1}^k L_{j,1} \wedge L_{j,2} \wedge L_{j,3}$  is a DNF over atoms  $W \cup X \cup Y \cup Z$ . Without loss of generality, we assume that  $\psi$  is unsatisfiable if we assign each  $Y_i \in Y$  the value true, i.e.,  $\forall W, Y, Z \neg \psi[Y/\top]$  evaluates to true.

Assume for the moment that  $W$  is void; then the negation of the QBF  $\forall X \exists Y \forall Z \psi$ , i.e.,  $\exists X \forall Y \exists Z \phi$  were  $\phi = \neg \psi$ , is encoded by the program  $\Pi$  in the proof of Theorem 7. That is,  $\Pi$  has some candidate world view iff  $\forall X \exists Y \forall Z \psi$  evaluates to false. A modification of  $\Pi$  yields that maximality testing of a guess  $\Phi$  is  $\Pi_3^p$ -hard. To this end, let

$$\Pi_0 = \{H \leftarrow B \wedge A \mid H \leftarrow B \in \Pi\} \cup \tag{11}$$

$$\{X_i \leftarrow \neg A; \bar{X}_i \leftarrow \neg A \mid X_i \in X\} \cup \{V \leftarrow \neg A\} \cup \tag{12}$$

$$\{A \leftarrow \mathbf{not} \neg A\}, \tag{13}$$

where  $A$  is a fresh atom, and  $\Phi = \emptyset$ . It is easy to see that  $\Phi$  has a candidate world view:  $I_0 = X \cup \bar{X} \cup \{V\}$  is the single answer set of  $\Pi_0^\Phi$ , and for each epistemic negation  $\mathbf{not} F$  in  $EP(\Pi_0) = \{\mathbf{not} V, \mathbf{not} \neg U, \mathbf{not} \neg A, \mathbf{not} X_i, \mathbf{not} \bar{X}_i \mid X_i \in X\}$  the formula  $F$  is true in  $I$ .

Furthermore, it holds that no guess  $\Phi' \supset \Phi$  has a candidate world view iff  $\Pi$  has no candidate world view. To see the only if part, suppose some  $\Phi' \supset \Phi$  has a candidate world view. Then  $\mathbf{not} \neg A \in \Phi'$  must hold, as otherwise  $I_0$  is the single answer set of  $\Pi_0^{\Phi'}$ ; this would imply  $\Phi' = \Phi$ , a contradiction. Now  $\Pi_0^{\Phi'}$  contains  $A \leftarrow$  and amounts to  $\Pi^{\Phi'}$  (simply eliminate  $A$ ); it follows that  $\Pi$  has a candidate world view w.r.t.  $\Phi' \cap EP(\Pi)$ . Conversely, if  $\Pi$  has a candidate world view w.r.t.  $\Phi$ , then it is easy to see that  $\Pi_0$  has a candidate world view w.r.t.  $\Phi' = \Phi \cup \{\mathbf{not} \neg A\}$ .

Thus,  $\Pi_0$  has a world view w.r.t.  $\Phi = \emptyset$  iff the QBF  $\forall X \exists Y \forall Z \psi$  evaluates to true (which proves  $\Pi_3^p$ -hardness of world view checking, i.e., deciding given a program  $\Pi$  and  $\Phi \subseteq EP(\Pi)$  whether  $\Pi$  has some world view w.r.t.  $\Phi$ ). Note that since any other guess  $\Phi' \supset \Phi$  that has a candidate world view w.r.t.  $\Pi_0$  must contain  $\mathbf{not} \neg A$ , we can equivalently ask whether  $\neg A$  is true in  $\Pi_0$  under EFLP semantics.

Now we generalize  $\Pi_0$  to accommodate non-void  $W$ , i.e., to encode evaluating the QBF (10). To this end, we let

$$\Pi_1 = \Pi_0^* \cup \{W_i \leftarrow \mathbf{not} \bar{W}_i; \bar{W}_i \leftarrow \mathbf{not} W_i \mid W_i \in W\}$$

where  $\bar{W}_i$  is a fresh atom for  $W_i$  and  $\Pi_0^*$  is constructed like  $\Pi_0$  where in the construction of the rules (7) of  $\Pi$  each  $W_i$  (resp.  $\neg W_i$ ) literal is replaced by  $\bar{W}_i$  (resp.  $W_i$ ).<sup>9</sup> Note that  $W_i$  and  $\bar{W}_i$  do not occur in  $\Pi_0^*$  in rule heads; combined with the assumption that  $\forall W, X, Z \neg \psi[Y/\top]$  evaluates to true, by similar reasoning as in the proof of Theorem 7 for  $W_i$  and  $\bar{W}_i$  in place of  $X_i$  and  $\bar{X}_i$ , we obtain that every guess  $\Phi \subseteq EP(\Pi_1)$  such that  $\Pi_1$  has a candidate world view w.r.t.  $\Phi$  must contain exactly one of  $\mathbf{not} W_i$  and  $\mathbf{not} \bar{W}_i$ , and thus  $\Phi$  encodes a truth assignment  $\nu$  to  $W$  where  $\nu(W_i) = \mathit{true}$  if  $\mathbf{not} \bar{W}_i \in \Phi$  (as  $W_i \leftarrow; \bar{W}_i \leftarrow \neg W_i$  is in  $\Pi_1^\Phi$ ) and  $\nu(W_i) = \mathit{false}$  if  $\mathbf{not} W_i \in \Phi$  (as  $\bar{W}_i \leftarrow; W_i \leftarrow \neg \bar{W}_i$  is in  $\Pi_1^\Phi$ ). Furthermore, if  $\mathbf{not} \neg A \notin \Phi$ , then  $\Phi$  contains no other epistemic negations, and we denote it by  $\Phi_\nu$ ; if  $\mathbf{not} \neg A \in \Phi$ , then the QBF  $\forall X \exists Y \forall Z \psi[W/\nu]$  must evaluate to false. Thus, if  $\Phi_\nu$  has a world view, then  $\forall X \exists Y \forall Z \psi[W/\nu]$  evaluates to true. On the other hand, every truth assignment to  $W$  is encoded by some  $\Phi_\nu$ ; thus, it follows that some  $\Phi_\nu$  has a world view iff the QBF (10) evaluates to true. As  $\neg A$  is true in the candidate world view of a guess  $\Phi$  iff  $\Phi = \Phi_\nu$  for some  $\nu$ , it follows that  $\neg A$  is true in  $\Pi_1$  under EFLP semantics iff the QBF (10) evaluates to true. As  $\Pi_1$  is constructible in polynomial time from (10), this proves  $\Sigma_4^p$ -hardness.  $\square$

## B Programs Matching Normal Epistemic Specifications

We show that the complexity of world view existence drops by one level in the polynomial hierarchy for epistemic specifications with rules of the form

$$L_0 \leftarrow L_1 \wedge \dots \wedge L_m \wedge (\neg) \mathbf{not} L_{m+1} \wedge \dots \wedge (\neg) \mathbf{not} L_n \quad (14)$$

<sup>9</sup>As  $\phi = \neg \psi$ , the final rules (11) that emerge from (7) are of the form  $U \leftarrow L_{j,1}^\dagger \wedge L_{j,3}^\dagger \wedge L_{j,3}^\dagger \wedge A$  where for any atom  $C$  we have  $C^\dagger = C$  and  $\neg C^\dagger = \bar{C}$ .

where  $L_0$  is an atom and each  $L_i$  is an atom  $A$  or default negated atom  $\neg A$ ; i.e., the rules match the syntax of (1) with  $m = 1$  (the modal atom  $\mathbf{MA}$  amounts to  $\mathbf{not} \neg A$ , and  $\neg\neg A = A$  under FLP semantics).<sup>10</sup>

Formally, we obtain, where  $D^p = D_1^p$  is the analog of  $D_2^p$  where  $\Sigma_2^p/\Pi_2^p$  is replaced by NP/coNP:

**Theorem 9** *Given a propositional program  $\Pi$  that is an normal epistemic specification and a guess  $\Phi$  for it, deciding whether  $\Pi$  has a candidate world view w.r.t.  $\Phi$  is  $D^p$ -complete.*

*Proof.* Membership. The proof of  $D^p$  membership is analogous to the one of Theorem 6, where NP/coNP replaces  $\Sigma_2^p/\Pi_2^p$ .

Hardness. The  $D^p$ -hardness follows from the reduction in the proof of Theorem 6: if the pair  $(\Pi_1, \Pi_2)$  consists of normal logic programs, deciding whether  $\Pi_1$  has some FLP answer set and  $\Pi_2$  has no FLP answer set is  $D^p$ -complete. Furthermore, the program  $\Pi$  constructed from  $(\Pi_1, \Pi_2)$  matches a normal epistemic specification.  $\square$

**Theorem 10** *Let  $\Pi$  be a propositional program that matches normal epistemic specifications. Then deciding whether  $\Pi$  has (i) some candidate world view and (ii) some world view are both  $\Sigma_2^p$ -complete.*

*Proof.* Membership. The membership proof is analogous to the one of Theorem 7, but uses Theorem 10 instead of Theorem 6.

Hardness. The  $\Sigma_2^p$ -hardness is inherited from the reduction in Theorem 7: if we consider a QBF (3) in which  $Z$  is void, the resulting program  $\Pi$  matches normal epistemic specifications, and has some candidate world view iff  $\exists X \forall Y \neg \phi$ , where  $\phi$  is a CNF, evaluates to true. Evaluating such QBFs is  $\Sigma_2^p$ -complete, and remains  $\Sigma_2^p$ -hard even if  $\phi$  is unsatisfiable if each  $Y_i \in Y$  is assigned true, i.e., the QBF  $\forall X \neg \phi[Y/\top]$  evaluates to true.  $\square$

**Theorem 11** *Given a propositional program  $\Pi$  that matches a normal epistemic specification and a propositional formula  $F$ , deciding whether  $F$  is true in  $\Pi$  (i) w.r.t. candidate world views is  $\Sigma_s^p$ -complete, and (ii) w.r.t. world views, i.e., under EFLP semantics, is  $\Sigma_3^p$ -complete.*

*Proof.* Membership. The membership proofs are analogous to the one of Theorem 8, but use Theorem 10 instead of Theorem 6 and the fact that for any guess  $\Phi$ ,  $\Pi^\Phi$  is a normal logic program, and cautious inference from the answer sets of such programs is coNP-complete [22].

Hardness. The  $\Sigma_2^p$ -hardness of (i) is immediate from the reduction in the proof of item (i) of Theorem 8 and Theorem 10. The  $\Sigma_3^p$ -hardness of (ii) is shown by generalizing the encoding of evaluating a QBF  $\exists X \forall Y \psi$  to deciding candidate world existence for a program  $\Pi$  that matches normal epistemic specifications in Theorem 10 in the same way as in the proof of item (ii) in Theorem 8; note that all rules created match normal epistemic specifications.  $\square$

<sup>10</sup>Recall that we omit strong negation  $\sim$ ; however, strongly negated literals  $\sim A$  can be compiled away in polynomial time (cf. Example 3), and thus the results of this section extend to a respective extension.