

Epistemic Answer Set Programming

Ezgi Iraz Su^{1,2}(✉)

¹ Department of Computer Engineering, Istinye University, Istanbul, Turkey
Ezgi-Iraz.Su@irit.fr

² IRIT, University of Toulouse, Toulouse, France

Abstract. This paper introduces a new epistemic extension of answer set programming (ASP) called *epistemic ASP* (E-ASP). Then, it compares E-ASP with existing approaches, showing the advantages and the novelties of the new semantics and discusses which formalisms provide more intuitive results: compared to Gelfond’s epistemic specifications (ES), E-ASP defines a simpler, but sufficiently strong language. Its *epistemic view* semantics is a natural and more standard generalisation of ASP’s original answer set semantics, so it allows for ASP’s previous language extensions. Moreover, compared to all semantics proposals in the literature, epistemic view semantics facilitates understanding of the intuitive meaning of epistemic logic programs and solves unintended results discussed in the literature, especially for epistemic logic programs including constraints.

Keywords: Answer set programming · Epistemic specifications · Modal logic S5 · Stable models · Answer sets · World views · Autoepistemic equilibrium models

1 Introduction

Logic programming (LP) [21] unifies different areas of computation by exploiting the greater generality of logic. Answer set programming (ASP) [13, 14, 24, 25] is an approach to declarative programming, and it relates LP to declarative problem solving by *answer sets*—consistent sets A of literals¹ in which $p \notin A$ or $\sim p \notin A$. In a sense, they are *partial* valuations. Some researchers prefer to call them *3-valued*. For instance, empty valuation assigns neither *true* nor *false* to a propositional variable p , leaving it *undetermined*, which is characterised by *negation as failure* (NAF) [6, 7] in ASP. Answer set semantics [10] of ASP has

I want to thank Andreas Herzig, Luis Fariñas del Cerro, Michael Gelfond, Patrick Thor Kahl, Thomas Eiter, Yi-Dong Shen, Pedro Cabalar, and Jorge Fandinno for their research related to this paper and the anonymous reviewers for their valued comments on the drafts of this work.

¹ In ASP, a literal is a propositional variable p or a *strongly-negated* propositional variable $\sim p$.

Soon after, [4] pointed at another program, $\Pi = \{p \leftarrow Mq, \text{not } q \mid q \leftarrow Mp, \text{not } p\}$ giving unintended world views under Kahl’s refined version. Indeed, [18] proposes two world views $\{\emptyset\}$ and $\{\{p\}, \{q\}\}$, of which the former seems to be unintended. Following this example, Kahl et al. [20] came up with another update to address the issue, with semantics supporting only the latter: inspired by [29] (see Sect. 4), they first define

$$\text{Ep}(\Pi) = \{G : G = \text{not } \mathbf{K} L \text{ for some extended objective literal } L \text{ and } G \text{ appears in } \Pi\}.$$

Note that the set $\text{Ep}(\Pi)$ checks all extended subjective literals occurring in Π and by using the equivalences between $\text{notnot}\mathbf{K}p$ and $\mathbf{K}p$, as well as $\text{not}\mathbf{K}\text{not}p$ and Mp , picks the forms $\text{not}\mathbf{K}L$ (for an extended objective literal L) in their structure. To illustrate this set, consider the program $\Gamma = \{t \leftarrow \mathbf{K}p, Mq, \text{not}\mathbf{K}r, \text{not}Ms\}$. Thus, $\text{Ep}(\Gamma) = \{\text{not}\mathbf{K}p, \text{not}\mathbf{K}\text{not}q, \text{not}\mathbf{K}r, \text{not}\mathbf{K}\text{not}s\}$. Then, they take the subset $\Phi_{\mathcal{A}} = \{G \in \text{Ep}(\Pi) : \mathcal{A} \models_{\text{ES}} G\}$ w.r.t. a candidate model \mathcal{A} . Finally, \mathcal{A} is a world view of Π if:

$$\mathcal{A} = \text{AS}(\Pi^{\mathcal{A}}), \text{ and there is no } \mathcal{A}' \text{ such that } \mathcal{A}' = \text{AS}(\Pi^{\mathcal{A}'}) \text{ and } \Phi_{\mathcal{A}'} \supset \Phi_{\mathcal{A}}.$$

On the one hand, as mentioned in [3, 19], researchers have now discovered another problem: different from their effect on answer sets in ASP⁴, in ES inserting a constraint into a program may now bring out completely new world views. The reason is because here constraints show their effect on its belief sets rather than a world view as a whole. So, not only Kahl’s all versions, but also [4, 29] suffer from new counterintuitive results produced over acyclic programs while they are trying to obtain the intuitive understanding of the behaviour of cycles. Interestingly, only in Gelfond’s first version [8, 9], and its generalisation [34] by Truszczyński, constraints function to rule out world views, violating that constraint (as desired). To sum up, as a negative outcome of added complexity, the recent semantics approaches seem to have lost this property. On the other hand, as argued in [29], Kahl’s reduct definition offers a complex program transformation, lacking an intuitive explanation for the replacement of subjective literals.

Let us terminate our discussion with two examples: $\Pi_1 = \{p \leftarrow \mathbf{K}p \mid p \leftarrow \text{not } \mathbf{K}p\}$ and $\Pi_2 = \{p \leftarrow Mp \mid p \leftarrow \text{not } Mp\}$. As Kahl obtains no world view for Π_1 , he gets a unique world view $\{\{p\}\}$ for Π_2 . However, the body parts of these programs produce a tautology (see the equivalence relations given in Table 5), so it is strange to see two different solutions according to his semantics approach, but not $\{\{p\}\}$ only for both.

⁴ In ASP, constraints show their effect on programs by eliminating or keeping their answer sets.

3 Fariñas et al.’s Approach: Autoepistemic Equilibrium Models

In 2015, Fariñas et al. [4] proposed the *autoepistemic equilibrium models* (AEEMs) approach as an alternative semantics for ES. This section briefly recalls their approach.

3.1 Epistemic Here-and-There Logic (EHT) and Its Equilibrium Models

EHT extends the logic of here-and-there (HT) [15] by (nondual) epistemic modal operators K and \hat{K} , of which K is the same as K in ES, but [4] never explains the meaning and reading of \hat{K} . Note that as shown later via an example, the modal operator M in ES is translated to $\neg K \neg$ in EHT. An EHT model is a collection of HT models. It can also be described as a refinement of S5 models (sets of valuations) [2] in which valuations are replaced by HT models. Formally, an *EHT model* is an ordered pair $\langle \mathcal{T}, \hat{h} \rangle$ in which

- $\mathcal{T} \subseteq 2^{\mathbb{P}}$ is a nonempty set of valuations (i.e., a classical S5 model);
- $\hat{h} : \mathcal{T} \rightarrow 2^{\mathbb{P}}$ is a map, assigning to each *there*-world $T \in \mathcal{T}$ a *here*-world $\hat{h}(T) \subseteq T$.

Epistemic equilibrium models (EEMs) of a formula $\varphi \in \mathcal{L}_{\text{EHT}}$ are then defined as particular S5 models satisfying a minimality condition [4] (similar to that of [27])⁵:

$$\text{EEM}(\varphi) = \{ \mathcal{T} \subseteq 2^{\mathbb{P}} : \mathcal{T}, \mathcal{T} \models_{\text{S5}} \varphi \text{ and there is no } h \neq \text{id} \text{ such that } \langle \mathcal{T}, \hat{h} \rangle, \mathcal{T} \models_{\text{EHT}} \varphi \}.$$

A typical ES program Π is translated into an EHT theory Π^* via a map $(\cdot)^*$ as in:

$$\begin{aligned} \Pi &= \{ p \text{ or } \sim q \leftarrow M r, \text{not } s \mid q \leftarrow \text{not } K p \\ \Pi^* &= ((\neg K \neg r \wedge \neg s) \rightarrow (p \vee \tilde{q})) \wedge (\neg K p \rightarrow q) \wedge \neg(q \wedge \tilde{q}). \end{aligned}$$

The EEM approach fails to give intuitive results, especially in the presence of disjunction. Gelfond’s example Π_G (see (1)) immediately supports this fact (see Table 6 for more examples): Π_G has a unique world view $\{\{h, e, i\}, \{f, i\}\}$, but Π_G^* has three EEMs: $\mathcal{T}_1 = \{\{h, e, i\}, \{f, i\}\}$, $\mathcal{T}_2 = \{\{h, e\}\}$ and $\mathcal{T}_3 = \{\{f, i\}\}$, among which \mathcal{T}_2 is unintended. To overcome this problem, [4] uses a *selection process* over EEMs and proposes *autoepistemic equilibrium models* (AEEMs)⁶.

⁵ For the truth conditions of EHT, you can refer to [4].

⁶ However, as in Kahl’s approach, adding a constraint into a program may also give here unexpected results. For instance, take the eligibility program Π_G and a constraint $\leftarrow i$. Then, the resulting EHT theory $\Pi_G^* \cup \{\neg i\}$ has a unique AEEM $\mathcal{T}_2 = \{\{h, e\}\}$, instead of having no AEEM.

The AEEM approach can handle a more general language, but its way of choosing intuitive models is highly complex: the AEEM semantics depends on two orderings, set inclusion \subseteq and a preference ordering \leq_φ , which function simultaneously. So, it is possible in principle that two EEMs can eliminate each other w.r.t. their different orderings. To spell it out, it could happen that $\mathcal{T}_1 \subset \mathcal{T}_2$, but also $\mathcal{T}_2 <_\varphi \mathcal{T}_1$ for $\mathcal{T}_1, \mathcal{T}_2 \in \mathbf{EEM}(\varphi)$. For instance, let $\varphi = p \vee r \vee \mathbf{K}(p \vee q)$. Then, $\mathbf{EEM}(\varphi) = \{\{\{p\}\}, \{\{r\}\}, \{\{q\}\}, \{\{p\}, \{q\}\}, \{\{p\}, \{r\}\}\}$, among which there are two satisfying the condition above: $\{\{p\}, \{q\}\} <_\varphi \{\{p\}\}$ and $\{\{p\}\} \subset \{\{p\}, \{q\}\}$. Fortunately, in this case, $\mathbf{AEEM}(\varphi) = \{\{\{p\}, \{r\}\}\}$ since $\{\{p\}\} \subset \{\{p\}, \{r\}\}$ and $\{\{p\}, \{q\}\} <_\varphi \{\{p\}, \{r\}\}$. One immediate question is if $\{\{p\}, \{q\}\}$ is indeed unintended. Briefly, this approach may be suffering from such a clash in the selection process although none has been found so far.

4 Shen and Eiter’s Approach: Epistemic Negation

In 2016, Shen et al. [29, 30] proposed a new semantics for ES. The idea is to use \mathbf{notK} (which they call *epistemic negation*) to minimise knowledge in the set of all belief sets. Given an epistemic logic program Π and a nonempty collection $\mathcal{A} \subseteq 2^{\mathbb{P}}$ of consistent sets of objective literals, let $\mathbf{Ep}(\Pi)$ (see Sect. 2.2) be the set of all epistemic negations appearing in Π , and let $\Phi \subseteq \mathbf{Ep}(\Pi)$ be its subset (which they call a *guess*). Let $\Phi_{\mathcal{A}} = \{G \in \mathbf{Ep}(\Pi) : \mathcal{A} \models G\}$ be the set of all epistemic negations in Π , satisfied by \mathcal{A} . Then, we transform Π into an epistemic reduct Π^Φ w.r.t. Φ by replacing every $\mathbf{notK} L \in \Phi$ with \top and every $\mathbf{notK} L \in \mathbf{Ep}(\Pi) \setminus \Phi$ with $\mathbf{not} L$. Finally, \mathcal{A} is a world view of Π if

1. $\mathcal{A} = \mathbf{AS}(\Pi^\Phi) = \{A : A \text{ is an answer set of } \Pi^\Phi\}$;
2. $\Phi_{\mathcal{A}}$ agrees with Φ , i.e., $\Phi_{\mathcal{A}} = \Phi$;
3. Φ is maximal, i.e., there is no bigger guess $\Phi' \supset \Phi$ such that $\mathcal{A}' = \mathbf{AS}(\Pi^{\Phi'})$ and $\Phi_{\mathcal{A}'} = \Phi'$ for some nonempty collection \mathcal{A}' of consistent sets of objective literals.

Let us illustrate their approach by an important application and motivation of ES: *Closed World Assumption* (CWA), which says that “*p is assumed to be false if there is no evidence to the contrary*” and is expressed in ASP by $\sim p \leftarrow \mathbf{not} p$ ⁷. However, it is formalised more adequately in ES as $\sim p \leftarrow \mathbf{notM} p$ by [12] or $\sim p \leftarrow \mathbf{notK} p$ by [29]. For instance, let $\Pi = \{\tilde{p} \leftarrow \mathbf{notK} p \mid \perp \leftarrow p, \tilde{p}\}$. Then, take the guess $\Phi = \{\mathbf{notK} p\}$. Thus, $\Pi^\Phi = \{\tilde{p} \leftarrow \top \mid \perp \leftarrow p, \tilde{p}\}$. Clearly, $\mathbf{AS}(\Pi^\Phi) = \{\{\tilde{p}\}\}$ and $\{\{\tilde{p}\}\} \models_{\mathbf{ES}} \Phi$. Since Φ is the maximal guess possible (see item 3 above), $\{\{\tilde{p}\}\}$ is the unique world view of Π .

⁷ However, this formalisation was then discovered to cause problems [12]. Consider $\Pi = \{p \text{ or } q \mid \sim p \leftarrow \mathbf{not} p\}$. Then, $\mathbf{AS}(\Pi) = \{\{p\}, \{q, \sim p\}\}$, and it answers the query $\sim p$? unknown (as it does not appear in both answer sets) while p is undetermined. This result is unintended.

5 Our Approach: Epistemic ASP (E-ASP) and Its Epistemic Views

This section introduces an epistemic extension of ASP called *epistemic ASP*. We begin with a discussion on our motivation and the main differences with other approaches.

5.1 Motivation and Novelty

The problem of incomplete information in ASP still matters after more than two decades of research on the subject. Despite their successes, the approaches [4, 20, 29] are not fully satisfactory, and some of their seemingly intuitive results are still under discussion.

Compared to Kahl’s language [18], we introduce a simpler language. We propose a modest syntax character, allowing only one epistemic operator K . However, different from ES, K may also appear in the head of a rule. We find our language strong enough to solve the problem of incomplete information in ASP because most of the critical examples in the literature, including Gelfond’s Π_G (1) and the new formalisation of CWA, use $\text{not}K$ only to solve the quantification problem. Note that Kahl and others use also M as dual of K . Besides, we more naturally extend the syntax of ASP through the same structure of program rules, allowing not to appear only in front of literals. Note that Kahl and others use it in a literal formation as $\text{not}K$, $\text{not}K\text{not}$, $K\text{not}$ etc.

The semantics of the new language is via an *epistemic view*, which is a straightforward generalisation of the answer set notion in ASP. Different from world view semantics, our semantics approach exploits a two-fold computation procedure, by splitting the program into two levels: we first look for if the candidate model, which is involved in the reduction process, is a maximal minimal model of the first level. Our reduct definition is oriented to eliminate NAF in a similar way with that of ASP. Existing reduct definitions simplify the program by removing subjective literals in the form of Kl , Ml , $\text{not}Kl$ and $\text{not}Ml$ (but, not ‘ $\text{not}l$ ’) for an objective literal l . So, our reduct is always a positive program containing no NAF in it. The minimality condition is understood in the sense of set inclusion. It is given by checking the minimality of each set making up the (biggest possible) collection. This is similar to the method, searching for answer sets. Second, we check if such minimal models of the first level are compatible with the second level, composed of only the constraints of the main program. So, we aim to solve the recent constraint problem discussed in Sect. 2.2. We find the semantics approaches of [18, 29] a bit nonstandard: first, is it a right attitude to eliminate the “positive” constructs in the form of Kl and Ml especially while world views of ES are given as kind of S5 models? Second, why do we force each valuation in such S5 models to be an answer set of the resulting program? In ES, we ask the query to the collection as a whole rather than its elements separately. At least, there is something going wrong in these approaches as always a new unintended model is being discovered, and then the reduction definitions have to be changed. To end with, our semantics approach can also be more smoothly

adapted to E-ASP programs with NAF in the head [16] and to E-ASP programs with nested expressions [23], also including subjective literal Kl .

5.2 The Language of Epistemic ASP ($\mathcal{L}_{\text{E-ASP}}$)

Literals (λ) of $\mathcal{L}_{\text{E-ASP}}$ are of two types: *objective literals* (l) and *subjective literals* (g).

l		g	
p	$\sim p$	Kp	$K\sim p$

in which $p \in \mathbb{P}$, and \sim denotes strong negation. Kl is read “ l is known”. Different from ES, we do not allow NAF to appear in a literal formation. However, NAF can precede any literal λ in the body of a rule, and **not** λ means that: *there is no evidence for λ , and so, the query $\lambda?$ is undetermined*. Again, different from ES, we allow Kl to appear in the head of a rule. Thus, an E-ASP program is defined as a finite collection of rules

$$\lambda_1 \text{ or } \dots \text{ or } \lambda_k \leftarrow \lambda_{k+1}, \dots, \lambda_m, \text{not } \lambda_{m+1}, \dots, \text{not } \lambda_n$$

in which λ_i 's are arbitrary (objective or subjective) literals. When we restrict λ_i 's to objective literals, the resulting program is a disjunctive logic program [14]. Hence, E-ASP rules are conservative extensions of ASP's disjunctive rules. As we follow the same structure, extensions to richer languages are straightforward via the main ASP track.

5.3 The Semantics of Epistemic ASP

The semantics of E-ASP is given by an *epistemic view*. Similar to a world view, it is a nonempty collection of consistent sets of objective literals. What we substantially differ is how we pick such intuitive models from among all models of an epistemic program. Let Π be an E-ASP program. We first split Π into two disjoint parts. The set of all constraints $r_c \in \Pi$ constitutes the upper layer (*'top'*), symbolised by $\overline{\Pi}$. This is the part of the program where we decide the ultimate epistemic views of Π through the process: *refute*, *accept* or *reorganise*. The rest, i.e., the set $\Pi \setminus \overline{\Pi}$ forms the lower layer (*'bottom'*), where we determine the collections of possible belief sets. We denote it by $\underline{\Pi}$.

Example 1. Given a program

$$\Sigma = \left\{ p \leftarrow \text{not } \sim q \mid \sim q \leftarrow \text{not } p \mid r \leftarrow \text{not } Kp \mid \leftarrow \text{not } r \right. \quad (2)$$

we have $\underline{\Sigma} = \left\{ p \leftarrow \text{not } \sim q \mid \sim q \leftarrow \text{not } p \mid r \leftarrow \text{not } Kp \right.$ and $\overline{\Sigma} = \left\{ \leftarrow \text{not } r \right.$

We start by computing the epistemic views of $\underline{\Pi}$, each of which are then involved in an evaluation process carried out in $\overline{\Pi}$. However, if $\underline{\Pi} = \emptyset$, then $\text{EV}(\Pi) = \text{EV}(\overline{\Pi})$, where $\text{EV}(\Pi)$ denotes the set of all epistemic views of Π . In this case, the

epistemic view of the program is either $\{\emptyset\}$ or none. For instance, $\text{EV}(\{\leftarrow p\}) = \{\{\emptyset\}\}$. Recall that $\leftarrow p$ has a unique answer set, namely \emptyset . If $\text{EV}(\underline{II}) = \emptyset$, then $\text{EV}(\overline{II}) = \emptyset$. When $\overline{II} = \emptyset$, $\text{EV}(\overline{II}) = \text{EV}(\underline{II})$.

Our reduct based semantics is oriented to eliminate only NAF as in ASP. Remember that NAF appears as part of a construct $\text{not } \lambda$ in an E-ASP program in which λ is an arbitrary literal. We here follow a “guess-and-check” method: let \mathcal{A} be a nonempty collection of consistent sets of objective literals, and let $A \in \mathcal{A}$. Then, $\langle \mathcal{A}, A \rangle$ is a sort of pointed (3-valued) S5 model with A being the *actual* world. In an explicit representation, we simply underline the actual world A in a collection \mathcal{A} . The partial valuation of A assigns *true* to p if $p \in A$ and *false* if $\sim p \in A$ (*undefined* otherwise). The reduct $\underline{II}^{\langle \mathcal{A}, A \rangle}$ of \underline{II} w.r.t. $\langle \mathcal{A}, A \rangle$ is given by replacing every occurrence of $\text{not } \lambda$ with⁸

- R.1** \perp if $\mathcal{A}, A \models_{\text{E-ASP}} \lambda$ (simply, for $\lambda = l$ if $A \models_{\text{E-ASP}} l$; for $\lambda = \text{K}l$ if $\mathcal{A} \models_{\text{E-ASP}} \text{K}l$);
R.2 \top if $\mathcal{A}, A \not\models_{\text{E-ASP}} \lambda$ (simply, for $\lambda = l$ if $A \not\models_{\text{E-ASP}} l$; for $\lambda = \text{K}l$ if $\mathcal{A} \not\models_{\text{E-ASP}} \text{K}l$).

Example 2. Given a pointed model $\{\underline{\{p\}}, \{\sim q\}\}$, consider (2) above. Then, $\underline{\Sigma}^{\{\underline{\{p\}}, \{\sim q\}\}} = \{p \leftarrow \top \mid \sim q \leftarrow \perp \mid r \leftarrow \top$ since $\{p\} \not\models_{\text{E-ASP}} \sim q$, $\{p\} \models_{\text{E-ASP}} p$ and $\{\{p\}, \{\sim q\}\} \not\models_{\text{E-ASP}} \text{K}p$. Now, we replace $\text{notK}p$ by $\text{K}p$ and $\text{not}r$ by $\text{notK}r$ in Σ and call the resulting program Γ :

$$\Gamma = \{p \leftarrow \text{not } \sim q \mid \sim q \leftarrow \text{not } p \mid r \leftarrow \text{K}p \mid \leftarrow \text{not } \text{K}r. \quad (3)$$

Then, $\underline{\Gamma}^{\{\underline{\{p\}}, \{\sim q\}\}} = \{p \leftarrow \perp \mid \sim q \leftarrow \top \mid r \leftarrow \text{K}p$ since $\{\sim q\} \models_{\text{E-ASP}} \sim q$, but $\{\sim q\} \not\models_{\text{E-ASP}} p$.

Thus, our reduct definition simplifies a program, removing only NAF w.r.t. R.1 and R.2.

First of all, we introduce a *truth-minimality* criterion, based on set inclusion over each set A making up a collection \mathcal{A} : let $\mathbb{O}\text{-Lit}$ be the set of all objective literals of $\mathcal{L}_{\text{E-ASP}}$, and let $\mathbf{s} : \mathcal{A} \rightarrow 2^{\mathbb{O}\text{-Lit}}$ be a (subset) map such that $\mathbf{s}(A) \subseteq A$ for every $A \in \mathcal{A}$. (When \mathbf{s} equals the identity map id , we obtain \mathcal{A} itself.) Then, a *weakening* of \mathcal{A} at a point $A \in \mathcal{A}$ is identified with $\langle \mathbf{s}[\mathcal{A}], \mathbf{s}(A) \rangle$ such that $\mathbf{s} \neq id$ and $\mathbf{s}|_{\mathcal{A} \setminus \{A\}} = id$, by which we take a strict subset of $A \in \mathcal{A}$ and do not modify the rest. We say that $\langle \mathbf{s}[\mathcal{A}], \mathbf{s}(A) \rangle$ is *weaker* than $\langle \mathcal{A}, A \rangle$ and denote it by $\langle \mathbf{s}[\mathcal{A}], \mathbf{s}(A) \rangle \triangleleft \langle \mathcal{A}, A \rangle$. For example, the weakenings of $\{\underline{\{p, \sim q\}}, \{r\}\}$ are $\{\underline{\{p\}}, \{r\}\}$, $\{\underline{\{\sim q\}}, \{r\}\}$ and $\{\underline{\emptyset}, \{r\}\}$. Finally, we define a nonmonotonic satisfaction relation \models^* for pointed (three-valued) S5 models: $\mathcal{A}, A \models^* II$ if and only if

$$\mathcal{A}, A \models_{\text{E-ASP}} II \text{ and } \mathbf{s}[\mathcal{A}], \mathbf{s}(A) \not\models_{\text{E-ASP}} II \text{ for every } \mathbf{s} \text{ viz. } \langle \mathbf{s}[\mathcal{A}], \mathbf{s}(A) \rangle \triangleleft \langle \mathcal{A}, A \rangle$$

where the latter condition says that none of the weakenings of $\langle \mathcal{A}, A \rangle$ is a model of II .

⁸ The satisfaction relation $\models_{\text{E-ASP}}$ of E-ASP is the same as the relation \models_{ES} (see Sect. 2.2).

Definition 1. Let $\mathcal{A} \subseteq 2^{\mathbb{0}\text{-Lit}}$ be a nonempty set of consistent sets of objective literals. Then, \mathcal{A} is a minimal model of Π if $\mathcal{A}, A \models^* \Pi^{\langle \mathcal{A}, A \rangle}$ for every $A \in \mathcal{A}$.

Example 3. $\{\{p\}, \{\sim q\}\}$ is a minimal model of $\underline{\Gamma}$ (3): $\underline{\Gamma}^{\{\{p\}, \{\sim q\}\}} = \{p \leftarrow \mid r \leftarrow \text{K}p$ and $\{\{p\}, \{\sim q\}\} \models_{\text{E-ASP}} \underline{\Gamma}^{\{\{p\}, \{\sim q\}\}}$ while its weakening $\{\emptyset, \{\sim q\}\}$ refutes it. Likewise, $\underline{\Gamma}^{\{\{p\}, \{\sim q\}\}} = \{\sim q \leftarrow \mid r \leftarrow \text{K}p$ and $\{\{p\}, \{\sim q\}\} \models_{\text{E-ASP}} \underline{\Gamma}^{\{\{p\}, \{\sim q\}\}}$ while its only weakening $\{\{p\}, \emptyset\}$ does not satisfy it. Clearly, $\{\{p, r\}\}$ and $\{\{\sim q\}\}$ are the other minimal models of $\underline{\Gamma}$. Similarly, $\{\{p, r\}, \{\sim q, r\}\}$ is a minimal model of $\underline{\Sigma}$ (2): indeed, $\underline{\Sigma}^{\{\{p, r\}, \{\sim q, r\}\}} = \{p \leftarrow \mid r \leftarrow$ and it is obvious that $\{\{p, r\}, \{\sim q, r\}\}$ satisfies it while all its weakenings refute it. We also have $\underline{\Sigma}^{\{\{p, r\}, \{\sim q, r\}\}} = \{\sim q \leftarrow \mid r \leftarrow$ and $\{\{p, r\}, \{\sim q, r\}\} \models_{\text{E-ASP}} \underline{\Sigma}^{\{\{p, r\}, \{\sim q, r\}\}}$ while any of its weakenings violates it. The other two minimal models of $\underline{\Sigma}$ are $\{\{\sim q, r\}\}$ and $\{\{p\}\}$. In each program, the last two minimal models (i.e., the singleton models⁹) are unintended.

As seen above, minimality of truth does not always guarantee intuitive results. Therefore, we will now introduce a criterion to choose intended models among all such minimal models. Given an E-ASP program Π , we first define a Π -indexed partial preorder (denoted by \preceq_{Π}) over three-valued S5 models by: $\mathcal{A} \preceq_{\Pi} \mathcal{A}'$ if and only if

$$\mathcal{A} \cup \mathcal{A}', A \models_{\text{E-ASP}} \Pi \text{ for all } A \in \mathcal{A} \text{ implies } \mathcal{A} \cup \mathcal{A}', A' \models_{\text{E-ASP}} \Pi \text{ for all } A' \in \mathcal{A}'. \quad (4)$$

The strict version of \preceq_{Π} is given as usual: $\mathcal{A} \prec_{\Pi} \mathcal{A}'$ iff $\mathcal{A} \preceq_{\Pi} \mathcal{A}'$ and $\mathcal{A}' \not\preceq_{\Pi} \mathcal{A}$. If $\mathcal{A} \preceq_{\Pi} \mathcal{A}'$ and $\mathcal{A}' \preceq_{\Pi} \mathcal{A}$, then \mathcal{A} is *equivalent* to \mathcal{A}' w.r.t. \preceq_{Π} (denoted by $\mathcal{A} \approx_{\Pi} \mathcal{A}'$).

Example 4. The program $\Upsilon = \{p \text{ or } q \mid p \leftarrow \text{notK}q$ has two minimal models: $\{\{p\}\}$ and $\{\{q\}\}$, among which $\{\{q\}\} \prec_r \{\{p\}\}$ since $\{\{p\}, \{q\}\} \models_{\text{E-ASP}} \Upsilon$, but $\{\{p\}, \{q\}\} \not\models_{\text{E-ASP}} \Upsilon$.

Definition 2. $\mathcal{A} \subseteq 2^{\mathbb{0}\text{-Lit}}$ is an epistemic view of a “constraint-free” program Π if

1. \mathcal{A} is a minimal model of Π ;
2. there is no minimal model \mathcal{A}' of Π such that $\mathcal{A} \prec_{\Pi} \mathcal{A}'$;

⁹ Singleton minimal models of a program Π are sometimes source of a problem in capturing intuitive results: for a singleton set, $\text{K}p$ and p are of no difference, as well as $\text{notK}p$ and $\text{not}p$. Thus, an E-ASP program performs like an ASP program, and we may obtain “unjustified” minimal models. For instance, in $\underline{\Sigma}$, if we replace notK with not , the resulting ASP program has the answer sets $\{p\}$ and $\{\sim q, r\}$. Note that $\{\{p\}\}$ and $\{\{\sim q, r\}\}$ are minimal models of $\underline{\Sigma}$. We get a similar result if we change $\text{K}p$ with p in $\underline{\Gamma}$. Thus, singleton sets do not allow us to quantify over all possible beliefs. In order to overcome this obstacle, we need to check the behaviour of singletons in an interplay with other minimal models by using an ordering.

Example 4, cont. Thus, $\text{EV}(\Upsilon) = \{\{\{p\}\}\}$. Let $\Lambda = \{p \leftarrow \text{not } q \mid q \leftarrow \text{notK } p\}$. Clearly, $\{\{p\}\}$ and $\{\{q\}\}$ are the only minimal models of Λ . Recall that for singletons, Λ behaves as an ASP program $\Lambda' = \{p \leftarrow \text{not } q \mid q \leftarrow \text{not } p\}$ and Λ' has 2 answer sets $\{p\}$ and $\{q\}$. So, again we cannot quantify over all beliefs. Indeed, $\{\{p\}, \{q\}\} \not\models_{\text{E-ASP}} \Lambda$, but $\{\{p\}, \{q\}\} \models_{\text{E-ASP}} \Lambda$. Thus, we have $\{\{p\}\} \prec_{\Lambda} \{\{q\}\}$. Consequently, $\text{EV}(\Lambda) = \{\{\{q\}\}\}$.

Example 5. We have seen that $\underline{\Sigma}$ (2) and $\underline{\Gamma}$ (3) have 3 minimal models. Among these, we have the order, $\{\{p\}\} \prec_{\underline{\Sigma}} \{\{\sim q, r\}\} \approx_{\underline{\Sigma}} \{\{p, r\}, \{\sim q, r\}\}$ and $\{\{\sim q\}\} \approx_{\underline{\Gamma}} \{\{p, r\}\} \approx_{\underline{\Gamma}} \{\{p\}, \{\sim q\}\}$. So, the ordering \preceq_{Π} is not strong enough to rule out all unintended models. When this is the case, we need to apply a third condition to compare equivalent models w.r.t. \preceq_{Π} .

We now introduce a *knowledge-minimising* condition: Let $\mathbb{L}(\cdot)$ represent the set of objective literals occurring in any syntactic construct (head, body, etc). We first consider the set $\mathcal{H}_{\Pi} = \bigcup_{r \in \Pi} \mathbb{L}(\text{head}(r))$ of all objective literals occurring in the head parts of a program Π . For example, $\mathcal{H}_{\underline{\Sigma}} = \mathcal{H}_{\underline{\Sigma}} = \mathcal{H}_{\underline{\Gamma}} = \mathcal{H}_{\underline{\Gamma}} = \{p, \sim q, r\}$ (see (2) and (3)). Note that belief sets \mathcal{A} 's of an epistemic view \mathcal{A} of Π can only contain literals from \mathcal{H}_{Π} ¹⁰. Inspired by [29] (but, in a different way), we define the set of all *unknowns* among the literals in \mathcal{H}_{Π} w.r.t. \mathcal{A} and denote it by $\Phi_{\mathcal{A}}^{\Pi} = \{l \in \mathcal{H}_{\Pi} : \mathcal{A} \models_{\text{E-ASP}} \text{notK } l\}$.

Definition 2, cont.

3. $\Phi_{\mathcal{A}}^{\Pi}$ is maximal, i.e., there is no minimal model \mathcal{A}' of Π such that $\Phi_{\mathcal{A}}^{\Pi} \subset \Phi_{\mathcal{A}'}^{\Pi}$.

Intuitively, item 3 means \mathcal{A} to answer maximum possible head-literals *undetermined*.

Example 5, cont. If we reconsider the above $\preceq_{\underline{\Sigma}}$ -equivalent and $\preceq_{\underline{\Gamma}}$ -equivalent minimal models, then we see that $\{p, \sim q\} = \Phi_{\{\{p, r\}, \{\sim q, r\}\}}^{\underline{\Sigma}} \supset \Phi_{\{\{\sim q, r\}\}}^{\underline{\Sigma}} = \{p\}$. As a result, $\text{EV}(\underline{\Sigma}) = \{\{\{p, r\}, \{\sim q, r\}\}\}$. Similarly, $\Phi_{\{\{p\}, \{\sim q\}\}}^{\underline{\Gamma}} = \{p, \sim q, r\}$, $\Phi_{\{\{\sim q\}\}}^{\underline{\Gamma}} = \{p, r\}$ and $\Phi_{\{\{p, r\}\}}^{\underline{\Gamma}} = \{\sim q\}$. Then, we have: $\Phi_{\{\{p\}, \{\sim q\}\}}^{\underline{\Gamma}} \supset \Phi_{\{\{\sim q\}\}}^{\underline{\Gamma}}$ and $\Phi_{\{\{p\}, \{\sim q\}\}}^{\underline{\Gamma}} \supset \Phi_{\{\{p, r\}\}}^{\underline{\Gamma}}$. Thus, $\text{EV}(\underline{\Gamma}) = \{\{\{p\}, \{\sim q\}\}\}$.

Remark 1. Note that there is also an order between the orders of item 2 and item 3: we only use item 3 over minimal models of Π that are maximal, but equivalent w.r.t. \preceq_{Π} .

Example 6. We now consider Gelfond's program Π_G (1): Π_G has 3 minimal models, namely $\mathcal{A}_1 = \{\{f, i\}\}$, $\mathcal{A}_2 = \{\{h, e\}\}$ and $\mathcal{A}_3 = \{\{f, i\}, \{h, e, i\}\}$, among which $\mathcal{A}_2 \prec_{\Pi} \mathcal{A}_3 \approx_{\Pi} \mathcal{A}_1$. Thus, we need to check the unknowns of \mathcal{A}_3 and \mathcal{A}_1 . Since $\{e, \sim e, h, f\} = \Phi_{\mathcal{A}_3}^{\Pi_G} \supset \Phi_{\mathcal{A}_1}^{\Pi_G} = \{e, \sim e, h\}$, we have $\text{EV}(\Pi_G) = \{\mathcal{A}_3\}$. However, we may not always compare maximal \preceq_{Π} -equivalent minimal models: let $\Omega = \{p \leftarrow \text{notK } q \mid q \leftarrow \text{notK } p\}$. Ω has two minimal models $\{\{p\}\}$ and $\{\{q\}\}$ such that

¹⁰ Fact [in ASP]: if $A \in \text{AS}(\Pi)$, then every $l \in A$ belongs to the head of one of the rules in Π .

$\{\{p\}\} \approx_{\Omega} \{\{q\}\}$ since $\{\{p\}, \{q\}\} \not\models_{\text{E-ASP}} \Omega$ and $\{\{p\}, \{q\}\} \not\models_{\text{E-ASP}} \Omega$. Moreover, $\Phi_{\{\{p\}\}}^{\Omega} = \{q\}$ and $\Phi_{\{\{q\}\}}^{\Omega} = \{p\}$. As a result, $\text{EV}(\Omega) = \{\{\{p\}\}, \{\{q\}\}\}$.

When a program Π contains constraints, i.e., $\overline{\Pi} \neq \emptyset$, we first compute $\text{EV}(\underline{\Pi})$ as explained above. Then, we evaluate each $\mathcal{A} \in \text{EV}(\underline{\Pi})$ w.r.t. their behaviour on $\overline{\Pi}$: take $\varphi = \bigvee_{r_c \in \overline{\Pi}} \text{body}(r_c)$. For every $A \in \mathcal{A}$, if $\mathcal{A}, A \not\models_{\text{E-ASP}} \varphi$, then we *accept* \mathcal{A} and call it $\mathcal{A}_{\text{accept}}$; else if $\mathcal{A}, A \models_{\text{E-ASP}} \varphi$, then we *eliminate* \mathcal{A} and call it $\mathcal{A}_{\text{refute}}$. Finally, we reorganise the rest in such a way that we take the biggest possible subset $\mathcal{A}_{\text{new}} \subseteq \mathcal{A}$ such that \mathcal{A}_{new} is still a minimal model of $\underline{\Pi}$ and $\mathcal{A}_{\text{new}}, A \not\models_{\text{E-ASP}} \varphi$, for every $A \in \mathcal{A}_{\text{new}}$. As a result, $\text{EV}(\Pi)$ is the collection of all $\mathcal{A}_{\text{accept}}$'s and \mathcal{A}_{new} 's. Note that when $\overline{\Pi}$ exclusively contains the constraints composed of only (negated) subjective literals, we either refute or accept the epistemic views of $\underline{\Pi}$.

Example 5, cont. We have seen that $\text{EV}(\underline{\Sigma}) = \{\{\{p, r\}, \{\sim q, r\}\}\}$ and $\text{EV}(\underline{\Gamma}) = \{\{\{p\}, \{\sim q\}\}\}$. As $\{\{p\}, \{\sim q\}\}$ violates $\leftarrow \text{notK}r$, it fails to be the epistemic view of the program Γ (*refute!*). Hence, $\text{EV}(\Gamma) = \emptyset$. However, as $\{\{p, r\}, \{\sim q, r\}\} \models_{\text{E-ASP}} r$, it satisfies $\leftarrow \text{not}r$, and so, it passes the test (*accept!*). Thus, $\text{EV}(\Sigma) = \{\{\{p, r\}, \{\sim q, r\}\}\}$.

Example 7. Let $\Delta = \{p \text{ or } q \leftarrow \mid r \text{ or } s \leftarrow \text{notK}p \mid \leftarrow r$. It is easy to see that

$$\text{EV}(\underline{\Delta}) = \{\{\{p, r\}, \{q, r\}, \{p, s\}, \{q, s\}\}\}.$$

Then, since $\{\{p, r\}, \{q, r\}, \{p, s\}, \{q, s\}\} \models_{\text{E-ASP}} r$, we have to remove the actual worlds $\{p, r\}$ and $\{q, r\}$, resulting in a new collection $\{\{p, s\}, \{q, s\}\}$ (*reorganise!*). As a result, $\text{EV}(\Delta) = \{\{\{p, s\}, \{q, s\}\}\}$. However, while $\text{EV}(\underline{\Delta} \cup \{\leftarrow \text{K}s\}) = \text{EV}(\underline{\Delta})$, $\text{EV}(\Delta \cup \{\leftarrow \text{K}s\}) = \emptyset$.

Example 8. Let $\Psi_1 = \{\leftarrow \text{K}p, \text{not}q\}$ and $\Psi_2 = \{\leftarrow \text{notK}p\}$ be the one-rule (constraint) E-ASP programs. As mentioned above, the only candidate epistemic view is $\{\emptyset\}$ for Ψ_1 and Ψ_2 . Since $\{\emptyset\} \not\models_{\text{E-ASP}} q$ and $\{\emptyset\} \not\models_{\text{E-ASP}} \text{K}p$, we have $\Psi_1^{\{\emptyset\}} = \{\leftarrow \text{K}p, \top\}$ and $\Psi_2^{\{\emptyset\}} = \{\leftarrow \top\}$. Clearly, $\text{EV}(\Psi_1) = \{\{\emptyset\}\}$ and $\text{EV}(\Psi_2) = \emptyset$.

5.4 Comparison of Epistemic Views with World Views and AEEMs

We here compare epistemic views with world views and AEEMs over some examples. Table 6 illustrates all these approaches. Overall, epistemic views of an E-ASP program perform well, aligning with its world views and AEEMs. However, one striking advantage of our method over existing semantics is its reasonable behaviour with programs including constraints: we have seen that $\text{EV}(\underline{\Gamma}) = \{\{\{p\}, \{\sim q\}\}\}$ and $\text{EV}(\Gamma) = \emptyset$. However, while $\underline{\Gamma}$ has a unique world view (AEEM) $\{\{p\}, \{\sim q\}\}$, when we add a constraint $\leftarrow \text{notK}r$ into $\underline{\Gamma}$, the resulting program Γ has another world view (AEEM) $\{\{p, r\}\}$, violating the above property (see the last two examples of Table 6 as well).

To end with, Shen et al. [29] discuss that Pearce's equilibrium semantics suffers from circular justifications and relatedly claim that [4] inherits the same

Table 6. World views by both [20] and [29], AEEMs (bold), and epistemic views (bold)

program Π	world views	EEEMs of Π^*	epistemic views
$p \text{ or } q \leftarrow$	$\{\{p\}, \{q\}\}$	$\{\{p\}\} \subset \{\{p\}, \{q\}\}$ $\{\{q\}\} \subset \{\{p\}, \{q\}\}$ $\{\{p\}\} \approx_{\Pi^*} \{\{q\}\} \approx_{\Pi^*} \{\{p\}, \{q\}\}$	$\{\{p\}, \{q\}\}$ minimal models (mm): $\{\{p\}\}, \{\{q\}\}$ and $\{\{p\}, \{q\}\}$
$p \text{ or } q \leftarrow$ $p \leftarrow Kq$	$\{\{p\}, \{q\}\}$	$\{\{p\}\} \subset \{\{p\}, \{q\}\}$ $\{\{p\}\} \approx_{\Pi^*} \{\{p\}, \{q\}\}$	$\{\{p\}, \{q\}\}$ mm: $\{\{p\}\}$ and $\{\{p\}, \{q\}\}$
$p \text{ or } q \leftarrow$ $p \leftarrow \text{not } Kq$	$\{\{p\}\}$	$\{\{q\}\} <_{\Pi^*} \{\{p\}\}$ (incomparable w.r.t. \subseteq)	$\{\{p\}\}$ mm: $\{\{p\}\}$
CWA: $p \text{ or } q \leftarrow$ $\bar{p} \leftarrow \text{not } Kp$ $\perp \leftarrow p, \bar{p}$	$\{\{q, \bar{p}\}\}$	$\{\{p\}\} <_{\Pi^*} \{\{q, \bar{p}\}\}$ (incomparable w.r.t. \subseteq)	$\{\{q, \bar{p}\}\}$ mm: $\{\{p\}\}$ and $\{\{q, \bar{p}\}\}$
$p \text{ or } q \leftarrow$ $r \leftarrow Kp$	$\{\{p\}, \{q\}\}$	$\{\{p, r\}\} <_{\Pi^*} \{\{p\}, \{q\}\}$ $\{\{q\}\} \subset \{\{p\}, \{q\}\}$ $\{\{q\}\} \approx_{\Pi^*} \{\{p\}, \{q\}\}$	$\{\{p\}, \{q\}\}$ mm: $\{\{q\}\}, \{\{p, r\}\},$ and $\{\{p\}, \{q\}\}$
$p \leftarrow$ $q \text{ or } r \leftarrow Kp$	$\{\{p, q\}, \{p, r\}\}$	$\{\{p, q\}\} \subset \{\{p, q\}, \{p, r\}\}$ $\{\{p, r\}\} \subset \{\{p, q\}, \{p, r\}\}$ $\{\{p, q\}\} \approx_{\Pi^*} \{\{p, r\}\} \approx_{\Pi^*} \{\{p, q\}, \{p, r\}\}$	$\{\{p, q\}, \{p, r\}\}$ mm: $\{\{p, r\}\}, \{\{p, q\}\},$ and $\{\{p, q\}, \{p, r\}\}$
$p \leftarrow Kp$	$\{\emptyset\}$	$\{\emptyset\}$	$\{\emptyset\}$
$p \leftarrow Kp$ $p \leftarrow \text{not } Kp$	none	none	none no minimal models
$p \leftarrow Kq$ $q \leftarrow Kp$	$\{\emptyset\}$	$\{\emptyset\}$	$\{\emptyset\}$ mm: $\{\emptyset\}$
$p \leftarrow q$ $q \leftarrow Kp$	$\{\emptyset\}$	$\{\emptyset\}$	$\{\emptyset\}$ mm: $\{\emptyset\}$
$p \leftarrow \text{not } q$ $q \leftarrow \text{not } Kp$	$\{\{q\}\}$	$\{\{p\}\} <_{\Pi^*} \{\{q\}\}$ (incomparable w.r.t. \subseteq)	$\{\{q\}\}$ mm: $\{\{p\}\}$ and $\{\{q\}\}$
$p \leftarrow \text{not } Kq$ $q \leftarrow \text{not } Kp$	$\{\{p\}\}$ and $\{\{q\}\}$	$\{\{p\}\}$ and $\{\{q\}\}$ (incomparable w.r.t. \subseteq and \leq_{Π^*})	$\{\{p\}\}$ and $\{\{q\}\}$ mm: $\{\{p\}\}$ and $\{\{q\}\}$
$p \text{ or } q \leftarrow$ $\leftarrow \text{not } Kp$	$\{\{p\}\}$	$\{\{p\}\}$	none no minimal models
$p \text{ or } q \leftarrow$ $r \leftarrow \text{not } Kq$ $\leftarrow p$	$\{\{q\}\}$	$\{\{q\}\}$	none no minimal models

circularity, leading to some undesired results. One supporting example is $\Pi = \{p \leftarrow \text{not } Kp \mid p \leftarrow p$. [29] argues that Π has no AEEMs, but in fact, $\{\{p\}\}$ is expected to be its unique world view since $\text{not } Kp \vee p$ constitutes a tautology. However, for $\{\{p\}\}$, this formula is of no difference than $\text{not } p \vee p$ and it is hard to believe that the latter is a tautology. Our approach agrees with [4]: since $\{\{p\}\} \not\models_{\text{E-ASP}} \text{not } Kp$, we have $\Pi^{\{\{p\}\}} = \{p \leftarrow \perp \mid p \leftarrow p$. Clearly, $\{\emptyset\}$ is the unique minimal model of $\Pi^{\{\{p\}\}}$. Thus, $\text{EV}(\Pi) = \emptyset$. Moreover, the first rule of Π intuitively says that *if there is no evidence for Kp , then p is always true*, so as already agreed by most of the approaches in the literature, this rule does

not have a world view in ES. The second rule is just a tautology, giving no information. Under these conditions, Π cannot have a world view in ES.

6 Splitting Epistemic Logic Programs

Cabalar et al. [3] have recently established a formal property called *epistemic splitting*, with which they test if a semantics proposal of ES has a reasonable behaviour when subjective literals are stratified. The idea is to separate a program Π into two disjoint subprograms (if possible), *top* and *bottom*, among which top questions bottom via its subjective literals, and bottom never refers to head literals of top. If splitting is the case w.r.t. a set of literals U , then we calculate world views of Π in four steps: first we compute the world views \mathcal{A}_b of bottom; second for each \mathcal{A}_b , we take kind of partial reduct $\Pi_U^{\mathcal{A}_b}$ by replacing subjective literals (whose literals are included in U) of top with their truth values in \mathcal{A}_b ; third we find the world views \mathcal{A}_t of $\Pi_U^{\mathcal{A}_b}$, and end with a solution $\langle \mathcal{A}_b, \mathcal{A}_t \rangle$ for Π ; finally we concatenate the elements of \mathcal{A}_b and \mathcal{A}_t , and result in new world views $\mathcal{A}_b \sqcup \mathcal{A}_t = \{A_b \cup A_t : A_b \in \mathcal{A}_b \text{ and } A_t \in \mathcal{A}_t\}$, answering the queried information.

All proposed semantics trials in the literature fail to satisfy this candidate property, but Gelfond's first version [8], which suffers most, among others, the counterintuitive behaviour of cyclic programs. (Recall that [8] computes two world views $\{\emptyset\}$ and $\{\{p\}\}$ for both $p \leftarrow Kp$ and $p \leftarrow Mp$. For the former rule, while $\{\{p\}\}$ is counterintuitive, for the latter, $\{\emptyset\}$ is counterintuitive, which has been justified by almost all semantics proposals in the literature.) The other semantics that passes epistemic splitting test is Truszczyński's approach [34]. (Remember that [34] produces a world view $\{\emptyset\}$ for the program $p \leftarrow p, \text{not}p$, which departs it even from ASP.) Our approach is also compatible with epistemic splitting property because first, in a splittable program we can always put all (and only) constraints composed of just (negated) subjective literals into topmost layer since they are headless, and they do not contain objective literal conjuncts, and the rest of the constraints will appear in the below layers; second, we can compute the epistemic views of the lower layers as defined in Sect. 2.2 by dividing each layer into two parts where constraints are located at the top; finally, we evaluate the final epistemic views according to the truth values of topmost subjective literals conjuncts in the candidate world view by keeping or eliminating candidate epistemic views.

Example 9. We can split Γ (see (3)) into 3 layers: $L_0 = \{p \leftarrow \text{not}\sim q \mid \sim q \leftarrow \text{not}p\}$, $L_1 = \{r \leftarrow Kp \text{ and } L_2 = \{\leftarrow \text{not}Kr\}$. Then, $\text{EV}(L_0) = \{\{\{p\}, \{\sim q\}\}\}$ and $\text{EV}(L_1^{\{\{p\}, \{\sim q\}\}}) = \text{EV}(r \leftarrow \perp) = \{\{\emptyset\}\}$. Next, $\text{EV}(L_0 \cup L_1) = \{\{\{p\}, \{\sim q\}\}\}$. Finally, since $\text{EV}(L_2^{\{\{p\}, \{\sim q\}\}}) = \text{EV}(\leftarrow \top) = \emptyset$, we have $\text{EV}(\Gamma) = \emptyset$.

7 Conclusion

In this paper, we propose a neat and more standard epistemic extension of ASP (E-ASP). E-ASP is a strong rival to existing approaches in the sense that: we

introduce a simpler and more intuitive semantics, which will be better suited for knowledge representation, and the design of intelligent agents. The new reduct definition, which is similar to that of ASP, will hopefully lead to an efficient implementation of an E-ASP program solver, allowing the new language to be of more practical use. We will search first if ASP technology can be exploited to compute epistemic views. E-ASP provides a solid framework for further language extensions of ASP. Therefore, we also plan to adapt previous language extensions of ASP to E-ASP. Finally, we would like to propose a new epistemic extension of equilibrium logic, embedding E-ASP as well.

References

1. Baral, C., Gelfond, M.: Logic programming and knowledge representation. *J. Log. Program.* **19**, 73–148 (1994)
2. Blackburn, P., de Rijke, M., Venema, Y.: *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge (2001)
3. Cabalar, P., Fandinno, J., Fariñas del Cerro, L.: Splitting epistemic logic programs. In: *Proceedings of the 17th International Workshop on Nonmonotonic Reasoning, NMR 2018, Tempe, Arizona, USA, 27–29 October 2018* (2018)
4. Fariñas del Cerro, L., Herzig, A., Su, E.I.: Epistemic equilibrium logic. In: Yang, Q., Wooldridge, M. (eds.) *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, pp. 2964–2970. AAAI Press (2015). <http://ijcai.org/papers15/Abstracts/IJCAI15-419.html>
5. Chen, J.: The generalized logic of only knowing (GOL) that covers the notion of epistemic specifications. *J. Log. Comput.* **7**(2), 159–174 (1997)
6. Clark, K.L.: Negation as failure. In: Gallaire, H., Minker, J. (eds.) *Logic and Data Bases, Symposium on Logic and Data Bases, Centre d’études et de recherches de Toulouse, France. Advances in Data Base Theory*, pp. 293–322. Plenum Press, New York (1977)
7. Gabbay, D.M.: What is negation as failure? In: Artikis, A., Craven, R., Kesim Çiçekli, N., Sadighi, B., Stathis, K. (eds.) *Logic Programs, Norms and Action*. LNCS (LNAI), vol. 7360, pp. 52–78. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29414-3_5
8. Gelfond, M.: Strong introspection. In: Dean, T.L., McKeown, K. (eds.) *Proceedings of the 9th National Conference on Artificial Intelligence, Anaheim, CA, USA, 14–19 July 1991*, vol. 1, pp. 386–391. AAAI Press/The MIT Press (1991)
9. Gelfond, M.: Logic programming and reasoning with incomplete information. *Ann. Math. Artif. Intell.* **12**(1–2), 89–116 (1994)
10. Gelfond, M.: Answer sets. In: *Handbook of Knowledge Representation*, vol. 1, p. 285 (2008)
11. Gelfond, M.: New semantics for epistemic specifications. In: Delgrande, J.P., Faber, W. (eds.) *LPNMR 2011*. LNCS (LNAI), vol. 6645, pp. 260–265. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20895-9_29
12. Gelfond, M.: New definition of epistemic specifications. In: *KR Seminar*. Texas Tech University, 28 April 2011. (talk)
13. Gelfond, M., Lifschitz, V.: The stable model semantics for logic programming. In: Kowalski, R.A., Bowen, K.A. (eds.) *Logic Programming, Proceedings of the Fifth International Conference and Symposium, Seattle, Washington, USA, 15–19 August 1988*, vol. 2, pp. 1070–1080. MIT Press (1988)

14. Gelfond, M., Lifschitz, V.: Classical negation in logic programs and disjunctive databases. *New Gener. Comput.* **9**(3/4), 365–386 (1991)
15. Heyting, A.: Die formalen Regeln der intuitionistischen Logik. *Sitzungsber. Preuss. Akad. Wiss.* **42–71**, 158–169 (1930)
16. Inoue, K., Sakama, C.: Negation as failure in the head. *J. Log. Program.* **35**(1), 39–78 (1998)
17. Kahl, P., Watson, R., Balai, E., Gelfond, M., Zhang, Y.: The language of epistemic specifications (refined) including a prototype solver. *J. Log. Comput.* (2015)
18. Kahl, P.T.: Refining the semantics for epistemic logic programs. Ph.D. thesis, Texas Tech University, Department of Computer Science, Lubbock, TX, USA, May 2014
19. Kahl, P.T., Leclerc, A.P.: Epistemic logic programs with world view constraints. In: Palù, A.D., Tarau, P., Saeedloei, N., Fodor, P. (eds.) *Technical Communications of the 34th International Conference on Logic Programming, ICLP 2018*, Oxford, United Kingdom, 14–17 July 2018. OpenAccess Series in Informatics OASICS, vol. 64, pp. 1:1–1:17. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2018). <https://doi.org/10.4230/OASICS.ICLP.2018.1>
20. Kahl, P.T., Leclerc, A.P., Son, T.C.: A parallel memory-efficient epistemic logic program solver: harder, better, faster. *CoRR abs/1608.06910* (2016). <http://arxiv.org/abs/1608.06910>
21. Kowalski, R.A.: Logic programming. In: Siekmann, J.H. (ed.) *Computational Logic, Handbook of the History of Logic*, vol. 9, pp. 523–569. Elsevier (2014). <https://doi.org/10.1016/B978-0-444-51624-4.50012-5>
22. Leclerc, A.P., Kahl, P.T.: A survey of advances in epistemic logic program solvers. *abs/1809.07141* (2018). <http://arxiv.org/abs/1809.07141>. (Also in the Proceedings of the 11th International Workshop on Answer Set Programming and other Computer Paradigms, ASPOCP 2018, Oxford, UK, 18 July 2018)
23. Lifschitz, V., Tang, L.R., Turner, H.: Nested expressions in logic programs. *Ann. Math. Artif. Intell.* **25**(3–4), 369–389 (1999)
24. Marek, V.W., Truszczyński, M.: Stable models and an alternative logic programming paradigm. *CoRR cs.LO/9809032* (1998). <http://arxiv.org/abs/cs.LO/9809032>
25. Niemelä, I.: Logic programs with stable model semantics as a constraint programming paradigm. *Ann. Math. Artif. Intell.* **25**(3–4), 241–273 (1999). <https://doi.org/10.1023/A:1018930122475>
26. Pearce, D.: A new logical characterisation of stable models and answer sets. In: Dix, J., Pereira, L.M., Przymusiński, T.C. (eds.) *NMELP 1996*. LNCS, vol. 1216, pp. 57–70. Springer, Heidelberg (1997). <https://doi.org/10.1007/BFb0023801>
27. Pearce, D.: Equilibrium logic. *Ann. Math. Artif. Intell.* **47**(1–2), 3–41 (2006)
28. Przymusiński, T.C.: On the relationship between logic programming and nonmonotonic reasoning. In: Shrobe, H.E., Mitchell, T.M., Smith, R.G. (eds.) *Proceedings of the 7th National Conference on Artificial Intelligence*, St. Paul, MN, USA, 21–26 August 1988, pp. 444–448. AAAI Press/The MIT Press (1988). <http://www.aaai.org/Library/AAAI/1988/aaai88-078.php>
29. Shen, Y., Eiter, T.: Evaluating epistemic negation in answer set programming. *Artif. Intell.* **237**, 115–135 (2016). <https://doi.org/10.1016/j.artint.2016.04.004>
30. Shen, Y., Eiter, T.: Evaluating epistemic negation in answer set programming (extended abstract). In: Sierra, C. (ed.) *Proceedings of the 26th International Joint Conference on Artificial Intelligence, IJCAI 2017*, Melbourne, Australia, 19–25 August 2017, pp. 5060–5064. *ijcai.org* (2017). <https://doi.org/10.24963/ijcai.2017/722>

31. Son, T.C., Le, T., Kahl, P.T., Leclerc, A.P.: On computing world views of epistemic logic programs. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, 19–25 August 2017, pp. 1269–1275 (2017). <https://doi.org/10.24963/ijcai.2017/176>
32. Stalnaker, R.: What is a nonmonotonic consequence relation? *Fundam. Inform.* **21**(1/2), 7–21 (1994)
33. Su, E.I.: Extensions of equilibrium logic by modal concepts. (Extensions de la logique d'équilibre par des concepts modaux). Ph.D. thesis, Institut de Recherche en Informatique de Toulouse, France (2015). <https://tel.archives-ouvertes.fr/tel-01636791>
34. Truszczyński, M.: Revisiting epistemic specifications. In: Balduccini, M., Son, T.C. (eds.) *Logic Programming, Knowledge Representation, and Nonmonotonic Reasoning*. LNCS (LNAI), vol. 6565, pp. 315–333. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20832-4_20
35. Wang, K., Zhang, Y.: Nested epistemic logic programs. In: Baral, C., Greco, G., Leone, N., Terracina, G. (eds.) *LPNMR 2005*. LNCS (LNAI), vol. 3662, pp. 279–290. Springer, Heidelberg (2005). https://doi.org/10.1007/11546207_22
36. Watson, R.: A splitting set theorem for epistemic specifications. *CoRR* cs.AI/0003038 (2000). <http://arxiv.org/abs/cs.AI/0003038>
37. Zhang, Y.: Updating epistemic logic programs. *J. Log. Comput.* **19**(2), 405–423 (2009). <https://doi.org/10.1093/logcom/exn100>
38. Zhang, Y., Zhang, Y.: Epistemic specifications and conformant planning. In: Barták, R., McCluskey, T.L., Pontelli, E. (eds.) *Proceedings of the 2017 Workshop on Knowledge-Based Techniques for Problem Solving and Reasoning (KnowProS 2017)* (2017)
39. Zhang, Z.: Introspecting preferences in answer set programming. In: Palù, A.D., Tarau, P., Saeedloei, N., Fodor, P. (eds.) *Technical Communications of the 34th International Conference on Logic Programming, ICLP 2018*, 14–17 July 2018, Oxford, United Kingdom. OASICS, vol. 64, pp. 3:1–3:13. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2018). <https://doi.org/10.4230/OASICS.ICLP.2018.3>