

A Semantical Framework for Hybrid Knowledge Bases

Jos de Bruijn

Free University of Bozen-Bolzano, Bozen-Bolzano, Italy

David Pearce

Universidad Politécnica de Madrid, Spain

Axel Polleres

DERI Galway, National University of Ireland, Galway, Ireland

Agustín Valverde

Universidad de Málaga, Málaga, Spain

Abstract. In the ongoing discussion about combining rules and Ontologies on the Semantic Web a recurring issue is how to combine first-order classical logic with nonmonotonic rule languages. Whereas several modular approaches to define a combined semantics for such hybrid knowledge bases focus mainly on decidability issues, we tackle the matter from a more general point of view. In this paper we show how Quantified Equilibrium Logic (QEL) can function as a unified framework which embraces classical logic as well as disjunctive logic programs under the (open) answer set semantics. In the proposed variant of QEL we relax the unique names assumption, which was present in earlier versions of QEL. Moreover, we show that this framework elegantly captures the existing modular approaches for hybrid knowledge bases in a unified way.

Keywords: Hybrid Knowledge Bases, Ontologies, Nonmonotonic Rules, Semantic Web, Logic Programming, Quantified Equilibrium Logic, Answer Set Programming

1. Introduction

In the current discussions on the Semantic Web architecture a recurring issue is how to combine a first-order classical theory formalising an ontology with a (possibly non-

Received Oct 30, 2007

Revised Feb 25, 2010

Accepted Mar 18, 2010

monotonic) rule base. Nonmonotonic rule languages have received considerable attention and achieved maturity over the last few years especially due to the success of Answer Set Programming (ASP), a nonmonotonic, purely declarative logic programming and knowledge representation paradigm with many useful features such as aggregates, weak constraints and priorities, supported by efficient implementations (for an overview see (Baral, 2002)).

As a logical foundation for the answer set semantics and a tool for logical analysis in ASP, the system of Equilibrium Logic was presented in (Pearce, 1997) and further developed in subsequent works (see (Pearce, 2006) for an overview and references). The aim of this paper is to show how Equilibrium Logic can be used as a logical foundation for the combination of ASP and Ontologies.

In the quest to provide a formal underpinning for a nonmonotonic rules layer for the Semantic Web which can coexist in a semantically well-defined manner with the Ontology layer, various proposals for combining classical first-order logic with different variants of ASP have been presented in the literature.¹ We distinguish three kinds of approaches: At one end of the spectrum there are approaches which provide an entailment-based query interface to the Ontology in the bodies of ASP rules, resulting in a loose integration (e.g. (Eiter, Lukasiewicz, Schindlauer and Tompits, 2004; Eiter, Ianni, Schindlauer and Tompits, 2005)). At the other end there are approaches which use a unifying nonmonotonic formalism to embed both the Ontology and the rule base (e.g. (de Bruijn, Eiter, Polleres and Tompits, 2007; Motik and Rosati, 2007)), resulting in a tight coupling. Hybrid approaches (e.g. (Rosati, 2005a; Rosati, 2005b; Rosati, 2006; Heymans, Predoiu, Feier, de Bruijn and van Nieuwenborgh, 2006)) fall between these extremes. Common to hybrid approaches is the definition of a modular semantics based on classical first-order models, on the one hand, and stable models – often, more generally, referred to as answer sets² – on the other hand. Additionally, they require several syntactical restrictions on the use of classical predicates within rules, typically driven by considerations upon retaining decidability of reasoning tasks such as knowledge base satisfiability and predicate subsumption. With further restrictions of the classical part to decidable Description Logics (DLs), these semantics support straightforward implementation using existing DL reasoners and ASP engines, in a modular fashion. In this paper, we focus on such hybrid approaches, but from a more general point of view.

Example 1. Consider a hybrid knowledge base consisting of a classical theory \mathcal{T} :

$$\begin{aligned} \forall x. PERSON(x) &\rightarrow (AGENT(x) \wedge (\exists y. HAS-MOTHER(x, y))) \\ \forall x. (\exists y. HAS-MOTHER(x, y)) &\rightarrow ANIMAL(x) \end{aligned}$$

which says that every *PERSON* is an *AGENT* and has some (unknown) mother, and everyone who has a mother is an *ANIMAL*, and a nonmonotonic logic program \mathcal{P} :

$$\begin{aligned} PERSON(x) &\leftarrow AGENT(x), \neg machine(x) \\ AGENT(DaveBowman) & \end{aligned}$$

which says that *AGENTS* are by default *PERSONS*, unless *known* to be *machines*, and *DaveBowman* is an *AGENT*.

Using such a hybrid knowledge base consisting of \mathcal{T} and \mathcal{P} , we intuitively would conclude that $PERSON(DaveBowman)$ holds since he is not known to be a *machine*,

¹ Most of these approaches focus on the Description Logics fragments of first-order logic underlying the Web Ontology Language OWL.

² “answer sets” denote the extension of stable models, which originally have only been defined for normal logic programs to more general logic programs such as disjunctive programs.

and furthermore we would conclude that *DaveBowman* has some (unknown) mother, and thus $ANIMAL(DaveBowman)$. \diamond

We see two important shortcomings in current hybrid approaches:

(1) Current approaches to hybrid knowledge bases differ not only in terms of syntactic restrictions, motivated by decidability considerations, but also in the way they deal with more fundamental issues which arise when classical logic meets ASP, such as the domain closure and unique names assumptions.³ In particular, current proposals implicitly deal with these issues by either restricting the allowed models of the classical theory, or by using variants of the traditional answer set semantics which cater for open domains and non-unique names. So far, little effort has been spent in comparing the approaches from a more general perspective. In this paper we aim to provide a generic semantic framework for hybrid knowledge bases that neither restricts models (e.g. to unique names) nor imposes syntactical restrictions driven by decidability concerns. (2) The semantics of current hybrid knowledge bases is defined in a modular fashion. This has the important advantage that algorithms for reasoning with this combination can be based on existing algorithms for DL and ASP satisfiability. A single underlying logic for hybrid knowledge bases which, for example, allows to capture notions of equivalence between combined knowledge bases in a standard way, is lacking though.

Our main contribution with this paper is twofold. First, we survey and compare different (extensions of the) answer set semantics, as well as the existing approaches to hybrid knowledge bases, all of which define nonmonotonic models in a modular fashion. Second, we propose to use Quantified Equilibrium Logic (QEL) as a unified logical foundation for hybrid knowledge bases: As it turns out, the equilibrium models of the combined knowledge base coincide exactly with the modular nonmonotonic models for all approaches we are aware of (Rosati, 2005a; Rosati, 2005b; Rosati, 2006; Heymans et al., 2006).

The remainder of this paper is structured as follows: Section 2 recalls some basics of classical first-order logic. Section 3 reformulates different variants of the answer set semantics introduced in the literature using a common notation and points out correspondences and discrepancies between these variants. Next, definitions of hybrid knowledge bases from the literature are compared and generalised in Section 4. QEL and its relation to the different variants of ASP are clarified in Section 5. Section 6 describes an embedding of hybrid knowledge bases into QEL and establishes the correspondence between equilibrium models and nonmonotonic models of hybrid KBs. We discuss some immediate implications of our results in Section 7. In Section 8 we show how for finite knowledge bases an equivalent semantical characterisation can be given via a second-order operator NM. This behaves analogously to the operator SM used by Ferraris, Lee and Lifschitz (Ferraris, Lee and Lifschitz, 2007) to define the stable models of a first-order sentence, except that its minimisation condition applies only to the non-classical predicates. In Section 9 we discuss an application of the previous results: we propose a definition of strong equivalence for knowledge bases sharing a common structural language and show how this notion can be captured by deduction in the (monotonic) logic of here-and-there. These two Sections (9 and 8) particularly contain mostly new material which has not yet been presented in the conference version (de Bruijn, Pearce, Polleres and Valverde, 2007) of this article. We conclude with a discussion of further related approaches and an outlook to future work in Section 10.

³ See (de Bruijn, Eiter, Polleres and Tompits, 2006) for a more in-depth discussion of these issues.

2. First-Order Logic (FOL)

A *function-free first-order language* $\mathcal{L} = \langle C, P \rangle$ with equality consists of disjoint sets of constant and predicate symbols C and P . Moreover, we assume a fixed countably infinite set of variables, the symbols ‘ \rightarrow ’, ‘ \vee ’, ‘ \wedge ’, ‘ \neg ’, ‘ \exists ’, ‘ \forall ’, and auxiliary parentheses ‘(,)’. Each predicate symbol $p \in P$ has an assigned arity $ar(p)$. Atoms and formulas are constructed as usual. Closed formulas, or *sentences*, are those where each variable is bound by some quantifier. A *theory* \mathcal{T} is a set of sentences. Variable-free atoms, formulas, or theories are also called *ground*. If D is a non-empty set, we denote by $At_D(C, P)$ the set of ground atoms constructible from $\mathcal{L}' = \langle C \cup D, P \rangle$.

Given a first-order language \mathcal{L} , an \mathcal{L} -structure consists of a pair $\mathcal{I} = \langle U, I \rangle$, where the *universe* $U = (D, \sigma)$ (sometimes called *pre-interpretation*) consists of a non-empty domain D and a function $\sigma: C \cup D \rightarrow D$ which assigns a domain value to each constant such that $\sigma(d) = d$ for every $d \in D$. For tuples we write $\sigma(\vec{t}) = (\sigma(d_1), \dots, \sigma(d_n))$. We call $d \in D$ an *unnamed individual* if there is no $c \in C$ such that $\sigma(c) = d$. The function I assigns a relation $p^I \subseteq D^n$ to each n -ary predicate symbol $p \in P$ and is called the \mathcal{L} -*interpretation over* D . The designated binary predicate symbol eq , occasionally written ‘=’ in infix notation, is assumed to be associated with the fixed interpretation function $eq^I = \{(d, d) : d \in D\}$. If \mathcal{I} is an \mathcal{L}' -structure we denote by $\mathcal{I}|_{\mathcal{L}}$ the restriction of \mathcal{I} to a sublanguage $\mathcal{L} \subseteq \mathcal{L}'$.

An \mathcal{L} -structure $\mathcal{I} = \langle U, I \rangle$ *satisfies* an atom $p(d_1, \dots, d_n)$ of $At_D(C, P)$, written $\mathcal{I} \models p(d_1, \dots, d_n)$, iff $(\sigma(d_1), \dots, \sigma(d_n)) \in p^I$. This is extended as usual to sentences and theories. \mathcal{I} is a *model* of an atom (sentence, theory, respectively) φ , written $\mathcal{I} \models \varphi$, if it satisfies φ . A theory \mathcal{T} *entails* a sentence φ , written $\mathcal{T} \models \varphi$, if every model of \mathcal{T} is also a model of φ . A theory is *consistent* if it has a model.

In the context of logic programs, the following assumptions often play a role: We say that the *parameter names assumption (PNA)* applies in case $\sigma|_C$ is surjective, i.e., there are no unnamed individuals in D ; the *unique names assumption (UNA)* applies in case $\sigma|_C$ is injective; in case both the PNA and UNA apply, the *standard names assumption (SNA)* applies, i.e. $\sigma|_C$ is a bijection. In the following, we will speak about PNA-, UNA-, or SNA-structures, (or PNA-, UNA-, or SNA-models, respectively), depending on σ .

An \mathcal{L} -interpretation I over D can be seen as a subset of $At_D(C, P)$. So, we can define a subset relation for \mathcal{L} -structures $\mathcal{I}_1 = \langle (D, \sigma_1), I_1 \rangle$ and $\mathcal{I}_2 = \langle (D, \sigma_2), I_2 \rangle$ over the same domain by setting $\mathcal{I}_1 \subseteq \mathcal{I}_2$ if $I_1 \subseteq I_2$.⁴ Whenever we speak about subset minimality of models/structures in the following, we thus mean minimality among all models/structures over the same domain.

3. Answer Set Semantics

In this paper we assume non-ground disjunctive logic programs with negation allowed in rule heads and bodies, interpreted under the answer set semantics (Lifschitz and Woo, 1992).⁵ A program \mathcal{P} consists of a set of rules of the form

$$a_1 \vee a_2 \vee \dots \vee a_k \vee \neg a_{k+1} \vee \dots \vee \neg a_l \leftarrow b_1, \dots, b_m, \neg b_{m+1}, \dots, \neg b_n \quad (1)$$

⁴ Note that this is not the substructure or submodel relation in classical model theory, which holds between a structure and its restriction to a subdomain.

⁵ By \neg we mean negation as failure and not classical, or strong negation, which is also sometimes considered in ASP.

where a_i ($i \in \{1, \dots, l\}$) and b_j ($j \in \{1, \dots, n\}$) are atoms, called head (body, respectively) atoms of the rule, in a function-free first-order language $\mathcal{L} = \langle C, P \rangle$ without equality. By $C_{\mathcal{P}} \subseteq C$ we denote the set of constants which appear in \mathcal{P} . A rule with $k = l$ and $m = n$ is called *positive*. Rules where each variable appears in b_1, \dots, b_m are called *safe*. A program is *positive (safe)* if all its rules are positive (safe).

For the purposes of this paper, we give a slightly generalised definition of the common notion of the *grounding* of a program: The *grounding* $gr_U(\mathcal{P})$ of \mathcal{P} wrt. a universe $U = (D, \sigma)$ denotes the set of all rules obtained as follows: For $r \in \mathcal{P}$, replace (i) each constant c appearing in r with $\sigma(c)$ and (ii) each variable with some element in D . Observe that thus $gr_U(\mathcal{P})$ is a ground program over the atoms in $At_D(C, P)$.

For a ground program \mathcal{P} and first-order structure \mathcal{I} the *reduct* $\mathcal{P}^{\mathcal{I}}$ consists of rules

$$a_1 \vee a_2 \vee \dots \vee a_k \leftarrow b_1, \dots, b_m$$

obtained from all rules of the form (1) in \mathcal{P} for which it holds that $\mathcal{I} \models a_i$ for all $k < i \leq l$ and $\mathcal{I} \not\models b_j$ for all $m < j \leq n$.

Answer set semantics is usually defined in terms of *Herbrand structures* over $\mathcal{L} = \langle C, P \rangle$. Herbrand structures have a fixed universe, the *Herbrand universe* $\mathcal{H} = (C, id)$, where id is the identity function. For a Herbrand structure $\mathcal{I} = \langle \mathcal{H}, I \rangle$, I can be viewed as a subset of the *Herbrand base*, \mathcal{B} , which consists of the ground atoms of \mathcal{L} . Note that by definition of \mathcal{H} , Herbrand structures are SNA-structures. A Herbrand structure \mathcal{I} is an *answer set* (Lifschitz and Woo, 1992) of \mathcal{P} if \mathcal{I} is subset minimal among the structures satisfying $gr_{\mathcal{H}}(\mathcal{P})^{\mathcal{I}}$. Two variations of this semantics, the open (Heymans, Nieuwenborgh and Vermeir, 2005) and generalised open answer set (Heymans et al., 2006) semantics, consider open domains, thereby relaxing the PNA. An *extended Herbrand structure* is a first-order structure based on a universe $U = (D, id)$, where $D \supseteq C$.

Definition 1. A first-order \mathcal{L} -structure $\mathcal{I} = \langle U, I \rangle$ is called a *generalised open answer set* of \mathcal{P} if \mathcal{I} is subset minimal among the structures satisfying all rules in $gr_U(\mathcal{P})^{\mathcal{I}}$. If, additionally, \mathcal{I} is an extended Herbrand structure, then \mathcal{I} is an *open answer set* of \mathcal{P} .

In the open answer set semantics the UNA applies. We have the following correspondence with the answer set semantics. First, as a straightforward consequence from the definitions, we can observe:

Proposition 1. If \mathcal{M} is an answer set of \mathcal{P} then \mathcal{M} is also an open answer set of \mathcal{P} .

The converse does not hold in general:

Example 2. Consider $\mathcal{P} = \{p(a); ok \leftarrow \neg p(x); \leftarrow \neg ok\}$ over $\mathcal{L} = \langle \{a\}, \{p, ok\} \rangle$. We leave it as an exercise to the reader to show that \mathcal{P} is inconsistent under the answer set semantics, but $\mathcal{M} = \langle (\{a, c_1\}, id), \{p(a), ok\} \rangle$ is an open answer set of \mathcal{P} . \diamond

Open answer set programs allow the use of the equality predicate ‘=’ in the body of rules. However, since this definition of open answer sets adheres to the UNA, one could argue that equality in open answer set programming is purely syntactical. Positive equality predicates in rule bodies can thus be eliminated by simple preprocessing, applying unification. This is not the case for negative occurrences of equality, but, since the interpretation of equality is fixed, these can be eliminated during grounding.

An alternative approach to relax the UNA has been presented by Rosati in (Rosati, 2005b): Instead of grounding with respect to U , programs are grounded with respect to the Herbrand universe $\mathcal{H} = (C, id)$, and minimality of the models of $gr_{\mathcal{H}}(\mathcal{P})^{\mathcal{I}}$ wrt. U is redefined: $\mathcal{I} \upharpoonright_{\mathcal{H}} = \{p(\sigma(c_1), \dots, \sigma(c_n)) : p(c_1, \dots, c_n) \in \mathcal{B}, \mathcal{I} \models p(c_1, \dots, c_n)\}$,

i.e., $\mathcal{I}|_{\mathcal{H}}$ is the restriction of \mathcal{I} to ground atoms of \mathcal{B} . Given \mathcal{L} -structures $\mathcal{I}_1 = (U_1, I_1)$ and $\mathcal{I}_2 = (U_2, I_2)$,⁶ the relation $\mathcal{I}_1 \subseteq_{\mathcal{H}} \mathcal{I}_2$ holds if $\mathcal{I}_1|_{\mathcal{H}} \subseteq \mathcal{I}_2|_{\mathcal{H}}$.

Definition 2. An \mathcal{L} -structure \mathcal{I} is called a *generalised answer set* of \mathcal{P} if \mathcal{I} is $\subseteq_{\mathcal{H}}$ -minimal among the structures satisfying all rules in $gr_{\mathcal{H}}(\mathcal{P})^{\mathcal{I}}$.

The following Lemma (implicit in (Heymans, 2006)) establishes that, for safe programs, all atoms of $At_D(C, P)$ satisfied in an open answer set of a safe program are ground atoms over $C_{\mathcal{P}}$:

Lemma 2. Let \mathcal{P} be a safe program over $\mathcal{L} = \langle C, P \rangle$ with $\mathcal{M} = \langle U, I \rangle$ a (generalised) open answer set over universe $U = (D, \sigma)$. Then, for any atom from $At_D(C, P)$ such that $\mathcal{M} \models p(d_1, \dots, d_n)$, there exist $c_i \in C_{\mathcal{P}}$ such that $\sigma(c_i) = d_i$ for each $1 \leq i \leq n$.

Proof: First, we observe that any atom $\mathcal{M} \models p(d_1, \dots, d_n)$ must be derivable from a sequence of rules $(r_0; \dots; r_l)$ in $gr_U(\mathcal{P})^{\mathcal{M}}$. We prove the lemma by induction over the length l of this sequence. $l = 0$: Assume $\mathcal{M} \models p(d_1, \dots, d_n)$, then r_0 must be (by safety) a ground fact in \mathcal{P} such that $p(\sigma(c_1), \dots, \sigma(c_n)) = p(d_1, \dots, d_n)$ and $c_1, \dots, c_n \in C_{\mathcal{P}}$. As for the induction step, let $p(d_1, \dots, d_n)$ be inferred by application of rule $r_l \in gr_U(\mathcal{P})^{\mathcal{M}}$. By safety, again each d_j either stems from a constant $c_j \in C_{\mathcal{P}}$ such that $\sigma(c_j) = d_j$ which appears in some true head atom of r_l or d_j also appears in a positive body atom $q(\dots, d_j, \dots)$ of r_l such that $\mathcal{M} \models q(\dots, d_j, \dots)$, derivable by $(r_0; \dots; r_{l-1})$, which, by the induction hypothesis, proves the existence of a $c_j \in C_{\mathcal{P}}$ with $\sigma(c_j) = d_j$. \square

From this Lemma, the following correspondence follows directly. Note that the answer sets and open answer sets of *safe* programs coincide as a direct consequence of Lemma 2:

Proposition 3. \mathcal{M} is an answer set of a *safe* program \mathcal{P} if and only if \mathcal{M} is an open answer set of \mathcal{P} .

Similarly, on unsafe programs, generalised answer sets and generalised open answer sets do not necessarily coincide, as demonstrated by Example 2. However, the following correspondence follows straightforwardly from Lemma 2:

Proposition 4. Given a safe program \mathcal{P} , \mathcal{M} is a generalised open answer set of \mathcal{P} if and only if \mathcal{M} is a generalised answer set of \mathcal{P} .

Proof:

(\Rightarrow) Assume \mathcal{M} is a generalised open answer set of \mathcal{P} . By Lemma 2 we know that rules in $gr_U(\mathcal{P})^{\mathcal{M}}$ involving unnamed individuals do not contribute to answer sets, since their body is always false. It follows that $\mathcal{M} = \mathcal{M}|_{\mathcal{H}}$ which in turn is a $\subseteq_{\mathcal{H}}$ -minimal model of $gr_{\mathcal{H}}(\mathcal{P})^{\mathcal{M}}$. This follows from the observation that each rule in $gr_{\mathcal{H}}(\mathcal{P})^{\mathcal{M}}$ and its corresponding rules in $gr_U(\mathcal{P})^{\mathcal{M}}$ are satisfied under the same models.

(\Leftarrow) Analogously. \square

By similar arguments, generalised answer sets and generalised open answer sets coincide in case the parameter name assumption applies:

Proposition 5. Let \mathcal{M} be a PNA-structure. Then \mathcal{M} is a generalised answer set of \mathcal{P} if and only if \mathcal{M} is a generalised open answer of \mathcal{P} .

⁶ Not necessarily over the same domain.

If the SNA applies, consistency with respect to all semantics introduced so far boils down to consistency under the original definition of answer sets:

Proposition 6. A program \mathcal{P} has an answer set if and only if \mathcal{P} has a generalised open answer set under the SNA.

Answer sets under SNA may differ from the original answer sets since also non-Herbrand structures are allowed. Further, we observe that there are programs which have generalised (open) answer sets but do not have (open) answer sets, even for safe programs, as shown by the following simple example:

Example 3. Consider $\mathcal{P} = \{p(a); \leftarrow \neg p(b)\}$ over $\mathcal{L} = \langle \{a, b\}, \{p\} \rangle$. \mathcal{P} is ground, thus obviously safe. However, although \mathcal{P} has a *generalised* (open) answer set – the reader may verify this by, for instance, considering the one-element universe $U = (\{d\}, \sigma)$, where $\sigma(a) = \sigma(b) = d$ – it is inconsistent under the open answer set semantics, i.e. the program does not have any open (non-generalised) answer set. \diamond

4. Hybrid Knowledge Bases

We now turn to the concept of hybrid knowledge bases, which combine classical theories with the various notions of answer sets. We define a notion of hybrid knowledge bases which generalizes definitions in the literature (Rosati, 2005a; Rosati, 2005b; Rosati, 2006; Heymans et al., 2006). We then compare and discuss the differences between the various definitions. It turns out that the differences are mainly concerned with the notion of answer sets, and syntactical restrictions, but do not change the general semantics. This will allow us to base our embedding into Quantified Equilibrium Logic on a unified definition.

A *hybrid knowledge base* $\mathcal{K} = (\mathcal{T}, \mathcal{P})$ over the function-free language $\mathcal{L} = \langle C, P_{\mathcal{T}} \cup P_{\mathcal{P}} \rangle$ consists of a classical first-order theory \mathcal{T} (also called the *structural* part of \mathcal{K}) over the language $\mathcal{L}_{\mathcal{T}} = \langle C, P_{\mathcal{T}} \rangle$ and a program \mathcal{P} (also called *rules* part of \mathcal{K}) over the language \mathcal{L} , where $P_{\mathcal{T}} \cap P_{\mathcal{P}} = \emptyset$, i.e. \mathcal{T} and \mathcal{P} share a single set of constants, and the predicate symbols allowed to be used in \mathcal{P} are a superset of the predicate symbols in $\mathcal{L}_{\mathcal{T}}$. Intuitively, the predicates in $\mathcal{L}_{\mathcal{T}}$ are interpreted classically, whereas the predicates in $\mathcal{L}_{\mathcal{P}}$ are interpreted nonmonotonically under the (generalised open) answer set semantics. With $\mathcal{L}_{\mathcal{P}} = \langle C, P_{\mathcal{P}} \rangle$ we denote the restricted language of \mathcal{P} to only the distinct predicates $P_{\mathcal{P}}$ which are not supposed to occur in \mathcal{T} .

We do not consider the alternative classical semantics defined in (Rosati, 2005a; Rosati, 2005b; Rosati, 2006), as these are straightforward.

We define the *projection* of a ground program \mathcal{P} with respect to an \mathcal{L} -structure $\mathcal{I} = \langle U, I \rangle$, denoted $\Pi(\mathcal{P}, \mathcal{I})$, as follows: for each rule $r \in \mathcal{P}$, r^{Π} is defined as:

1. $r^{\Pi} = \emptyset$ if there is a literal over $At_D(C, P_{\mathcal{T}})$ in the head of r of form $p(\vec{t})$ such that $p(\sigma(\vec{t})) \in I$ or of form $\neg p(\vec{t})$ with $p(\sigma(\vec{t})) \notin I$;
2. $r^{\Pi} = \emptyset$ if there is a literal over $At_D(C, P_{\mathcal{T}})$ in the body of r of form $p(\vec{t})$ such that $p(\sigma(\vec{t})) \notin I$ or of form $\neg p(\vec{t})$ such that $p(\sigma(\vec{t})) \in I$;
3. otherwise r^{Π} is the singleton set resulting from r by deleting all occurrences of literals from $\mathcal{L}_{\mathcal{T}}$,

and $\Pi(\mathcal{P}, \mathcal{I}) = \bigcup \{r^{\Pi} : r \in \mathcal{P}\}$. Intuitively, the projection “evaluates” all classical literals in \mathcal{P} with respect to \mathcal{I} .

Definition 3. Let $\mathcal{K} = (\mathcal{T}, \mathcal{P})$ be a hybrid knowledge base over the language $\mathcal{L} =$

$\langle C, P_{\mathcal{T}} \cup P_{\mathcal{P}} \rangle$. An NM-model $\mathcal{M} = \langle U, I \rangle$ (with $U = (D, \sigma)$) of \mathcal{K} is a first-order \mathcal{L} -structure such that $\mathcal{M}|_{\mathcal{L}_{\mathcal{T}}}$ is a model of \mathcal{T} and $\mathcal{M}|_{\mathcal{L}_{\mathcal{P}}}$ is a generalised open answer set of $\Pi(\text{gr}_U(\mathcal{P}), \mathcal{M})$.

Analogous to first-order models, we speak about PNA-, UNA-, and SNA-NM-models.

Example 4. Consider the hybrid knowledge base $\mathcal{K} = (\mathcal{T}, \mathcal{P})$, with \mathcal{T} and \mathcal{P} as in Example 1, with the capitalised predicates being predicates in $P_{\mathcal{T}}$. Now consider the interpretation $\mathcal{I} = \langle U, I \rangle$ (with $U = (D, \sigma)$) with $D = \{\text{DaveBowman}, k\}$, σ the identity function, and $I = \{\text{AGENT}(\text{DaveBowman}), \text{HAS-MOTHER}(\text{DaveBowman}, k), \text{ANIMAL}(\text{DaveBowman}), \text{machine}(\text{DaveBowman})\}$. Clearly, $\mathcal{I}|_{\mathcal{L}_{\mathcal{T}}}$ is a model of \mathcal{T} . The projection $\Pi(\text{gr}_U(\mathcal{P}), \mathcal{I})$ is

$$\leftarrow \neg \text{machine}(\text{DaveBowman}),$$

which does not have a stable model, and thus \mathcal{I} is not an NM-model of \mathcal{K} . In fact, the logic program \mathcal{P} ensures that an interpretation cannot be an NM-model of \mathcal{K} if there is an *AGENT* which is neither a *PERSON* nor known (by conclusions from \mathcal{P}) to be a *machine*. It is easy to verify that, for any NM-model of \mathcal{K} , the atoms $\text{AGENT}(\text{DaveBowman})$, $\text{PERSON}(\text{DaveBowman})$, and $\text{ANIMAL}(\text{DaveBowman})$ must be true, and are thus entailed by \mathcal{K} . The latter cannot be derived from neither \mathcal{T} nor \mathcal{P} individually. \diamond

4.1. r-hybrid KBs

We now proceed to compare our definition of NM-models with the various definitions in the literature. The first kind of hybrid knowledge base we consider was introduced by Rosati in (Rosati, 2005a) (and extended in (Rosati, 2006) under the name $\mathcal{DL}+\text{log}$), and was labeled *r-hybrid* knowledge base. Syntactically, *r-hybrid* KBs do not allow negated atoms in rule heads, i.e. for rules of the form (1) $l = k$, and do not allow atoms from $\mathcal{L}_{\mathcal{T}}$ to occur negatively in the rule body.⁷ Moreover, in (Rosati, 2005a), Rosati deploys a restriction which is stronger than standard safety: each variable must appear in at least one positive body atom with a predicate from $\mathcal{L}_{\mathcal{P}}$. We call this condition $\mathcal{L}_{\mathcal{P}}$ -safe in the remainder. In (Rosati, 2006) this condition is relaxed to *weak $\mathcal{L}_{\mathcal{P}}$ -safety*: there is no special safety restriction for variables which occur only in body atoms from $P_{\mathcal{T}}$.

Semantically, Rosati assumes (an infinite number of) standard names, i.e. C is countably infinite, and normal answer sets, in his version of NM-models:

Definition 4. Let $\mathcal{K} = (\mathcal{T}, \mathcal{P})$ be an r-hybrid knowledge base, over the language $\mathcal{L} = \langle C, P_{\mathcal{T}} \cup P_{\mathcal{P}} \rangle$, where C is countably infinite, and \mathcal{P} is a (weak) $\mathcal{L}_{\mathcal{P}}$ -safe program. An *r-NM-model* $\mathcal{M} = \langle U, I \rangle$ of \mathcal{K} is a first-order \mathcal{L} -SNA-structure such that $\mathcal{M}|_{\mathcal{L}_{\mathcal{T}}}$ is a model of \mathcal{T} and $\mathcal{M}|_{\mathcal{L}_{\mathcal{P}}}$ is an answer set of $\Pi(\text{gr}_U(\mathcal{P}), \mathcal{M})$.

In view of the (weak) $\mathcal{L}_{\mathcal{P}}$ -safety condition, we observe that r-NM-model existence coincides with SNA-NM-model existence on r-hybrid knowledge bases, by Lemma 2 and Proposition 6.

The syntactic restrictions in r-hybrid knowledge bases guarantee decidability of the satisfiability problem in case satisfiability (in case of $\mathcal{L}_{\mathcal{P}}$ -safety) or conjunctive query

⁷ Note that by projection, negation of predicates from $P_{\mathcal{T}}$ is treated classically, whereas negation of predicates from $P_{\mathcal{P}}$ is treated nonmonotonically. This might be considered unintuitive and therefore a reason why Rosati disallows structural predicates to occur negated. The negative occurrence of classical predicates in the body is equivalent to the positive occurrence of the predicate in the head.

	SNA	variables	disjunctive rule heads	negated $\mathcal{L}_{\mathcal{T}}$ atoms
r-hybrid	yes	$\mathcal{L}_{\mathcal{P}}$ -safe	pos. only	no
r^+ -hybrid	no	$\mathcal{L}_{\mathcal{P}}$ -safe	pos. only	no
r_w -hybrid	yes	weak $\mathcal{L}_{\mathcal{P}}$ -safe	pos. only	no
g-hybrid	no	guarded	neg. allowed*	yes

* g-hybrid allows negation in the head but at most one positive head atom

Table 1. Different variants of hybrid KBs

containment (in case of weak $\mathcal{L}_{\mathcal{P}}$ -safety) in \mathcal{T} is decidable. Rosati (Rosati, 2005a; Rosati, 2006) presents sound and complete algorithms for both cases.

4.2. r^+ -hybrid KBs

In (Rosati, 2005b), Rosati relaxes the UNA for what we will call here r^+ -hybrid knowledge bases. In this variant the $\mathcal{L}_{\mathcal{P}}$ -safety restriction is kept but generalised answer sets under arbitrary interpretations are considered:

Definition 5. Let $\mathcal{K} = (\mathcal{T}, \mathcal{P})$ be an r^+ -hybrid knowledge base consisting of a theory \mathcal{T} and an $\mathcal{L}_{\mathcal{P}}$ -safe program \mathcal{P} . An r^+ -NM-model, $\mathcal{M} = \langle U, I \rangle$ of \mathcal{K} is a first-order \mathcal{L} -structure such that $\mathcal{M}|_{\mathcal{L}_{\mathcal{T}}}$ is a model of \mathcal{T} and $\mathcal{M}|_{\mathcal{L}_{\mathcal{P}}}$ is a generalised answer set of $\Pi(\text{gr}_U(\mathcal{P}), \mathcal{M})$.

$\mathcal{L}_{\mathcal{P}}$ -safety guarantees safety of $\Pi(\text{gr}_U(\mathcal{P}), \mathcal{M})$. Thus, by Proposition 3, we can conclude that r^+ -NM-models coincide with NM-models on r -hybrid knowledge bases. The relaxation of the UNA does not affect decidability

4.3. g-hybrid KBs

G-hybrid knowledge bases (Heymans et al., 2006) allow a different form of rules in the program. In order to regain decidability, rules are not required to be safe, but they are required to be *guarded* (hence the ‘g’ in g-hybrid): All variables in a rule are required to occur in a single positive body atom, the *guard*, with the exception that unsafe choice rules of the form

$$p(c_1, \dots, c_n) \vee \neg p(c_1, \dots, c_n) \leftarrow$$

are allowed. Moreover, disjunction in rule heads is limited to at most one positive atom, i.e. for rules of the form (1) we have that $k \leq 1$, but an arbitrary number of negated head atoms is allowed. Another significant difference is that, as opposed to the approaches based on r -hybrid KBs, negative structural predicates are allowed in the rules part within g-hybrid knowledge bases (see also Footnote 7). The definition of NM-models in (Heymans et al., 2006) coincides precisely with our Definition 3.

Table 4.3 summarises the different versions of hybrid knowledge bases introduced in the literature.

5. Quantified Equilibrium Logic (QEL)

Equilibrium logic for propositional theories and logic programs was presented in (Pearce, 1997) as a foundation for answer set semantics, and extended to the first-order case in (Pearce and Valverde, 2005), as well as, in slightly more general, modified form, in (Pearce and Valverde, 2006). For a survey of the main properties of equilibrium logic, see (Pearce, 2006). Usually in quantified equilibrium logic we consider a full first-order language allowing function symbols and we include a second, strong negation operator as occurs in several ASP dialects. For the present purpose of drawing comparisons with approaches to hybrid knowledge bases, it will suffice to consider the function-free language with a single negation symbol, ‘ \neg ’. In particular, we shall work with a quantified version of the logic HT of *here-and-there*. In other respects we follow the treatment of (Pearce and Valverde, 2006).

5.1. General Structures for Quantified Here-and-There Logic

As before, we consider a function-free first order language $\mathcal{L} = \langle C, P \rangle$ built over a set of *constant* symbols, C , and a set of *predicate* symbols, P . The sets of \mathcal{L} -formulas, \mathcal{L} -sentences and atomic \mathcal{L} -sentences are defined in the usual way. Again, we only work with *sentences*, and, as in Section 2, by an \mathcal{L} -interpretation I over a set D we mean a subset I of $At_D(C, P)$. A *here-and-there* \mathcal{L} -structure with static domains, or **QHT^s(\mathcal{L})-structure**, is a tuple $\mathcal{M} = \langle (D, \sigma), I_h, I_t \rangle$ where $\langle (D, \sigma), I_h \rangle$ and $\langle (D, \sigma), I_t \rangle$ are \mathcal{L} -structures such that $I_h \subseteq I_t$.

We can think of \mathcal{M} as a structure similar to a first-order classical model, but having two parts, or components, h and t that correspond to two different points or “worlds”, ‘here’ and ‘there’, in the sense of Kripke semantics for intuitionistic logic (van Dalen, 1983), where the worlds are ordered by $h \leq t$. At each world $w \in \{h, t\}$ one verifies a set of atoms I_w in the expanded language for the domain D . We call the model static, since, in contrast to say intuitionistic logic, the same domain serves each of the worlds.⁸ Since $h \leq t$, whatever is verified at h remains true at t . The satisfaction relation for \mathcal{M} is defined so as to reflect the two different components, so we write $\mathcal{M}, w \models \varphi$ to denote that φ is true in \mathcal{M} with respect to the w component. Evidently we should require that an atomic sentence is true at w just in case it belongs to the w -interpretation. Formally, if $p(t_1, \dots, t_n) \in At_D$ then

$$\mathcal{M}, w \models p(t_1, \dots, t_n) \quad \text{iff} \quad p(\sigma(t_1), \dots, \sigma(t_n)) \in I_w. \quad (2)$$

Then \models is extended recursively as follows⁹:

- $\mathcal{M}, w \models \varphi \wedge \psi$ iff $\mathcal{M}, w \models \varphi$ and $\mathcal{M}, w \models \psi$.
- $\mathcal{M}, w \models \varphi \vee \psi$ iff $\mathcal{M}, w \models \varphi$ or $\mathcal{M}, w \models \psi$.
- $\mathcal{M}, t \models \varphi \rightarrow \psi$ iff $\mathcal{M}, t \not\models \varphi$ or $\mathcal{M}, t \models \psi$.
- $\mathcal{M}, h \models \varphi \rightarrow \psi$ iff $\mathcal{M}, t \models \varphi \rightarrow \psi$ and $\mathcal{M}, h \not\models \varphi$ or $\mathcal{M}, h \models \psi$.
- $\mathcal{M}, w \models \neg\varphi$ iff $\mathcal{M}, t \not\models \varphi$.
- $\mathcal{M}, t \models \forall x\varphi(x)$ iff $\mathcal{M}, t \models \varphi(d)$ for all $d \in D$.

⁸ Alternatively it is quite common to speak of a logic with *constant* domains. However this is slightly ambiguous since it might suggest that the domain is composed only of constants, which is not intended here.

⁹ The reader may easily check that the following correspond exactly to the usual Kripke semantics for intuitionistic logic given our assumptions about the two worlds h and t and the single domain D , see e.g. (van Dalen, 1983)

- $\mathcal{M}, h \models \forall x\varphi(x)$ iff $\mathcal{M}, t \models \forall x\varphi(x)$ and $\mathcal{M}, h \models \varphi(d)$ for all $d \in D$.
- $\mathcal{M}, w \models \exists x\varphi(x)$ iff $\mathcal{M}, w \models \varphi(d)$ for some $d \in D$.

Truth of a sentence in a model is defined as follows: $\mathcal{M} \models \varphi$ iff $\mathcal{M}, w \models \varphi$ for each $w \in \{h, t\}$. A sentence φ is valid if it is true in all models, denoted by $\models \varphi$. A sentence φ is a consequence of a set of sentences Γ , denoted $\Gamma \models \varphi$, if every model of Γ is a model of φ . In a model \mathcal{M} we often use the symbols H and T , possibly with subscripts, to denote the interpretations I_h and I_t respectively; so, an \mathcal{L} -structure may be written in the form $\langle U, H, T \rangle$, where $U = (D, \sigma)$.

The resulting logic is called *Quantified Here-and-There Logic with static domains*, denoted by \mathbf{QHT}^s . In terms of satisfiability and validity this logic is equivalent to the logic introduced before in (Pearce and Valverde, 2005). By $\mathbf{QHT}_{=}^s$ we denote the version of QEL with equality. The equality predicate in $\mathbf{QHT}_{=}^s$ is interpreted as the actual equality in both worlds, ie $\mathcal{M}, w \models t_1 = t_2$ iff $\sigma(t_1) = \sigma(t_2)$.

The logic $\mathbf{QHT}_{=}^s$ can be axiomatised as follows. Let $\mathbf{INT}^=$ denote first-order intuitionistic logic (van Dalen, 1983) with the usual axioms for equality:

$$\begin{aligned} & x = x, \\ & x = y \rightarrow (F(x) \rightarrow F(y)), \end{aligned}$$

for every formula $F(x)$ such that y is substitutable for x in $F(x)$. To this we add the axiom of Hosoi

$$\alpha \vee (\neg\beta \vee (\alpha \rightarrow \beta)),$$

which determines 2-element here-and-there models in the propositional case, and the axiom SQHT (static quantified here-and-there):

$$\exists x(F(x) \rightarrow \forall xF(x)).$$

Lastly we add the “decidable equality” axiom:

$$x = y \vee x \neq y.$$

For a completeness proof for $\mathbf{QHT}_{=}^s$, see (Lifschitz, Pearce and Valverde, 2007).

As usual in first order logic, satisfiability and validity are independent from the language. If $\mathcal{M} = \langle (D, \sigma), H, T \rangle$ is an $\mathbf{QHT}_{=}^s(\mathcal{L}')$ -structure and $\mathcal{L} \subset \mathcal{L}'$, we denote by $\mathcal{M}|_{\mathcal{L}}$ the restriction of \mathcal{M} to the sublanguage \mathcal{L} : $\mathcal{M}|_{\mathcal{L}} = \langle (D, \sigma|_{\mathcal{L}}), H|_{\mathcal{L}}, T|_{\mathcal{L}} \rangle$.

Proposition 7. Suppose that $\mathcal{L}' \supset \mathcal{L}$, Γ is a theory in \mathcal{L} and \mathcal{M} is an \mathcal{L}' -structure such $\mathcal{M} \models \Gamma$. Then $\mathcal{M}|_{\mathcal{L}}$ is a model of Γ in $\mathbf{QHT}_{=}^s(\mathcal{L})$.

Proposition 8. Suppose that $\mathcal{L}' \supset \mathcal{L}$ and $\varphi \in \mathcal{L}$. Then φ is valid (resp. satisfiable) in $\mathbf{QHT}_{=}^s(\mathcal{L})$ if and only if is valid (resp. satisfiable) in $\mathbf{QHT}_{=}^s(\mathcal{L}')$.

Analogous to the case of classical models we can define special kinds of \mathbf{QHT}^s (resp. $\mathbf{QHT}_{=}^s$) models. Let $\mathcal{M} = \langle (D, \sigma), H, T \rangle$ be an \mathcal{L} -structure that is a model of a universal theory T . Then, we call \mathcal{M} a PNA-, UNA-, or SNA-model if the restriction of σ to constants in \mathcal{C} is surjective, injective or bijective, respectively.

5.2. Equilibrium Models

As in the propositional case, quantified equilibrium logic is based on a suitable notion of minimal model.

Definition 6. Among $\mathbf{QHT}_{=}^s(\mathcal{L})$ -structures we define the order \leq as: $\langle (D, \sigma), H, T \rangle \leq$

$\langle\langle D', \sigma' \rangle, H', T' \rangle$ if $D = D'$, $\sigma = \sigma'$, $T = T'$ and $H \subseteq H'$. If the subset relation is strict, we write ' \triangleleft '.

Definition 7. Let Γ be a set of sentences and $\mathcal{M} = \langle\langle D, \sigma \rangle, H, T \rangle$ a model of Γ .

1. \mathcal{M} is said to be *total* if $H = T$.
2. \mathcal{M} is said to be an *equilibrium* model of Γ (for short, we say: “ \mathcal{M} is in equilibrium”) if it is minimal under \trianglelefteq among models of Γ , and it is total.

Notice that a total $\mathbf{QHT}_{=}^s$ model of a theory Γ is equivalent to a classical first order model of Γ .

Proposition 9. Let Γ be a theory in \mathcal{L} and \mathcal{M} an equilibrium model of Γ in $\mathbf{QHT}_{=}^s(\mathcal{L}')$ with $\mathcal{L}' \supset \mathcal{L}$. Then $\mathcal{M}|_{\mathcal{L}}$ is an equilibrium model of Γ in $\mathbf{QHT}_{=}^s(\mathcal{L})$.

5.3. Relation to Answer Sets

The above version of QEL is described in more detail in (Pearce and Valverde, 2006). If we assume all models are UNA-models, we obtain the version of QEL found in (Pearce and Valverde, 2005). There, the relation of QEL to (ordinary) answer sets for logic programs with variables was established (in (Pearce and Valverde, 2005, Corollary 7.7)). For the present version of QEL the correspondence can be described as follows.

Proposition 10 ((Pearce and Valverde, 2006)). Let Γ be a universal theory in $\mathcal{L} = \langle C, P \rangle$. Let $\langle U, T, T \rangle$ be a total $\mathbf{QHT}_{=}^s$ model of Γ . Then $\langle U, T, T \rangle$ is an equilibrium model of Γ iff $\langle T, T \rangle$ is a propositional equilibrium model of $gr_U(\Gamma)$.

By convention, when \mathcal{P} is a logic program with variables we consider the models and equilibrium models of its universal closure expressed as a set of logical formulas. So, from Proposition 10 we obtain:

Corollary 11. Let \mathcal{P} be a logic program. A total $\mathbf{QHT}_{=}^s$ model $\langle U, T, T \rangle$ of \mathcal{P} is an equilibrium model of \mathcal{P} iff it is a generalised open answer set of \mathcal{P} .

Proof: It is well-known that for propositional programs equilibrium models coincide with answer sets (Pearce, 1997). Using Proposition 10 and Definition 4 for generalised open answer sets, the result follows. \square

6. Relation between Hybrid KBs and QEL

In this section we show how equilibrium models for hybrid knowledge bases relate to the NM models defined earlier and we show that QEL captures the various approaches to the semantics of hybrid KBs in the literature (Rosati, 2005a; Rosati, 2005b; Rosati, 2006; Heymans et al., 2006).

Given a hybrid KB $\mathcal{K} = (\mathcal{T}, \mathcal{P})$ we call $\mathcal{T} \cup \mathcal{P} \cup st(\mathcal{T})$ the *stable closure* of \mathcal{K} , where $st(\mathcal{T}) = \{\forall x(p(x) \vee \neg p(x)) : p \in \mathcal{L}_{\mathcal{T}}\}$.¹⁰ From now on, unless otherwise clear from context, the symbol ' \models ' denotes the truth relation for $\mathbf{QHT}_{=}^s$. Given a ground program \mathcal{P} and an \mathcal{L} -structure $\mathcal{M} = \langle U, H, T \rangle$, the *projection* $\Pi(\mathcal{P}, \mathcal{M})$ is understood to be defined relative to the component T of \mathcal{M} .

¹⁰ Evidently \mathcal{T} becomes *stable* in \mathcal{K} in the sense that $\forall \varphi \in \mathcal{T}, st(\mathcal{T}) \models \neg\neg\varphi \rightarrow \varphi$. The terminology is drawn from intuitionistic logic and mathematics.

Lemma 12. Let $\mathcal{M} = \langle U, H, T \rangle$ be a $\mathbf{QHT}_{=}^s$ -model of $\mathcal{T} \cup st(\mathcal{T})$. Then $\mathcal{M} \models \mathcal{P}$ iff $\mathcal{M}|_{\mathcal{L}_{\mathcal{P}}} \models \Pi(gr_U(\mathcal{P}), \mathcal{M})$.

Proof: By the hypothesis $\mathcal{M} \models \{\forall x(p(x) \vee \neg p(x)) : p \in \mathcal{L}_{\mathcal{T}}\}$. It follows that $H|_{\mathcal{L}_{\mathcal{T}}} = T|_{\mathcal{L}_{\mathcal{T}}}$. Consider any $r \in \mathcal{P}$, such that $r^{\Pi} \neq \emptyset$. Then there are four cases to consider. (i) r has the form $\alpha \rightarrow \beta \vee p(t)$, $p(t) \in \mathcal{L}_{\mathcal{T}}$ and $p(\sigma(t)) \notin T$, so $\mathcal{M} \models \neg p(t)$. W.l.o.g. assume that $\alpha, \beta \in \mathcal{L}_{\mathcal{P}}$, so $r^{\Pi} = \alpha \rightarrow \beta$ and

$$\mathcal{M} \models r \Leftrightarrow \mathcal{M} \models r^{\Pi} \Leftrightarrow \mathcal{M}|_{\mathcal{L}_{\mathcal{P}}} \models r^{\Pi} \quad (3)$$

by the semantics for $\mathbf{QHT}_{=}^s$ and Theorem 7. (ii) r has the form $\alpha \rightarrow \beta \vee \neg p(t)$, where $p(\sigma(t)) \in T$; so $p(\sigma(t)) \in H$ and $\mathcal{M} \models p(t)$. Again it is easy to see that (3) holds. Case (iii): r has the form $\alpha \wedge p(t) \rightarrow \beta$ and $p(\sigma(t)) \in H, T$, so $\mathcal{M} \models p(t)$. Case (iv): r has the form $\alpha \wedge \neg p(t) \rightarrow \beta$ and $\mathcal{M} \models \neg p(t)$. Clearly for these two cases (3) holds as well. It follows that if $\mathcal{M} \models \mathcal{P}$ then $\mathcal{M}|_{\mathcal{L}_{\mathcal{P}}} \models \Pi(gr_U(\mathcal{P}), \mathcal{M})$.

To check the converse condition we need now only examine the cases where $r^{\Pi} = \emptyset$. Suppose this arises because $p(\sigma(t)) \in H, T$, so $\mathcal{M} \models p(t)$. Now, if $p(t)$ is in the head of r , clearly $\mathcal{M} \models r$. Similarly if $\neg p(t)$ is in the body of r , by the semantics $\mathcal{M} \models r$. The cases where $p(\sigma(t)) \notin T$ are analogous and left to the reader. Consequently if $\mathcal{M}|_{\mathcal{L}_{\mathcal{P}}} \models \Pi(gr_U(\mathcal{P}), \mathcal{M})$, then $\mathcal{M} \models \mathcal{P}$. \square

We now state the relation between equilibrium models and NM-models.

Theorem 13. Let $\mathcal{K} = (\mathcal{T}, \mathcal{P})$ be a hybrid KB. $\mathcal{M} = \langle U, T, T \rangle$ is an equilibrium model of the stable closure of \mathcal{K} if and only if $\langle U, T \rangle$ is an NM-model of \mathcal{K} .

Proof: Assume the hypothesis and suppose that \mathcal{M} is in equilibrium. Since \mathcal{T} contains only predicates from $\mathcal{L}_{\mathcal{T}}$ and $\mathcal{M} \models \mathcal{T} \cup st(\mathcal{T})$, evidently

$$\mathcal{M}|_{\mathcal{L}_{\mathcal{T}}} \models \mathcal{T} \cup st(\mathcal{T}) \quad (4)$$

and so in particular $(U, \mathcal{M}|_{\mathcal{L}_{\mathcal{T}}})$ is a model of \mathcal{T} . By Lemma 12,

$$\mathcal{M} \models \mathcal{P} \Leftrightarrow \mathcal{M}|_{\mathcal{L}_{\mathcal{P}}} \models \Pi(gr_U(\mathcal{P}), \mathcal{M}). \quad (5)$$

We claim (i) that $\mathcal{M}|_{\mathcal{L}_{\mathcal{P}}}$ is an equilibrium model of $\Pi(gr_U(\mathcal{P}), \mathcal{M})$. If not, there is a model $\mathcal{M}' = \langle H', T' \rangle$ with $H' \subset T' = T|_{\mathcal{L}_{\mathcal{P}}}$ and $\mathcal{M}' \models \Pi(gr_U(\mathcal{P}), \mathcal{M})$. Lift (U, \mathcal{M}') to a (first order) \mathcal{L} -structure \mathcal{N} by interpreting each $p \in \mathcal{L}_{\mathcal{T}}$ according to \mathcal{M} . So $\mathcal{N}|_{\mathcal{L}_{\mathcal{T}}} = \mathcal{M}|_{\mathcal{L}_{\mathcal{T}}}$ and by (4) clearly $\mathcal{N} \models \mathcal{T} \cup st(\mathcal{T})$. Moreover, by Lemma 12 $\mathcal{N} \models \mathcal{P}$ and by assumption $\mathcal{N} \triangleleft \mathcal{M}$, contradicting the assumption that \mathcal{M} is an equilibrium model of $\mathcal{T} \cup st(\mathcal{T}) \cup \mathcal{P}$. This establishes (i). Lastly, we note that since $\langle T|_{\mathcal{L}_{\mathcal{P}}}, T|_{\mathcal{L}_{\mathcal{P}}} \rangle$ is an equilibrium model of $\Pi(gr_U(\mathcal{P}), \mathcal{M})$, $\mathcal{M}|_{\mathcal{L}_{\mathcal{P}}}$ is a generalised open answer set of $\Pi(gr_U(\mathcal{P}), \mathcal{M})$ by Corollary 11, so that $\mathcal{M} = \langle U, T, T \rangle$ is an NM-model of \mathcal{K} .

For the converse direction, assume the hypothesis but suppose that \mathcal{M} is not in equilibrium. Then there is a model $\mathcal{M}' = \langle U, H, T \rangle$ of $\mathcal{T} \cup st(\mathcal{T}) \cup \mathcal{P}$, with $H \subset T$. Since $\mathcal{M}' \models \mathcal{P}$ we can apply Lemma 12 to conclude that $\mathcal{M}'|_{\mathcal{L}_{\mathcal{P}}} \models \Pi(gr_U(\mathcal{P}), \mathcal{M}')$. But clearly

$$\Pi(gr_U(\mathcal{P}), \mathcal{M}') = \Pi(gr_U(\mathcal{P}), \mathcal{M}).$$

However, since evidently $\mathcal{M}'|_{\mathcal{L}_{\mathcal{T}}} = \mathcal{M}|_{\mathcal{L}_{\mathcal{T}}}$, thus $\mathcal{M}'|_{\mathcal{L}_{\mathcal{P}}} \triangleleft \mathcal{M}|_{\mathcal{L}_{\mathcal{P}}}$, so this shows that $\mathcal{M}|_{\mathcal{L}_{\mathcal{P}}}$ is not an equilibrium model of $\Pi(gr_U(\mathcal{P}), \mathcal{M})$ and therefore $T|_{\mathcal{L}_{\mathcal{P}}}$ is not an answer set of $\Pi(gr_U(\mathcal{P}), \mathcal{M})$ and \mathcal{M} is not an NM-model of \mathcal{K} . \square

This establishes the main theorem relating to the various special types of hybrid KBs discussed earlier.

Theorem 14 (Main Theorem). (i) Let $\mathcal{K} = (\mathcal{T}, \mathcal{P})$ be a g-hybrid (resp. an r^+ -hybrid)

knowledge base. $\mathcal{M} = \langle U, T, T \rangle$ is an equilibrium model of the stable closure of \mathcal{K} if and only if $\langle U, T \rangle$ is an NM-model (resp. r^+ -NM-model) of \mathcal{K} .

(ii) Let $\mathcal{K} = (\mathcal{T}, \mathcal{P})$ be an r -hybrid knowledge base. Let $\mathcal{M} = \langle U, T, T \rangle$ be an Herbrand model of the stable closure of \mathcal{K} . Then \mathcal{M} is in equilibrium in the sense of (Pearce and Valverde, 2005) if and only if $\langle U, T \rangle$ is an r -NM-model of \mathcal{K} .

Example 5. Consider again the hybrid knowledge base $\mathcal{K} = (\mathcal{T}, \mathcal{P})$, with \mathcal{T} and \mathcal{P} as in Example 1. The stable closure of \mathcal{K} , $st(\mathcal{K}) = \mathcal{T} \cup st(\mathcal{T}) \cup \mathcal{P}$ is

$$\begin{aligned} & \forall x. PERSON(x) \rightarrow (AGENT(x) \wedge (\exists y. HAS-MOTHER(x, y))) \\ & \forall x. (\exists y. HAS-MOTHER(x, y)) \rightarrow ANIMAL(x) \\ & \forall x. PERSON(x) \vee \neg PERSON(x) \\ & \forall x. AGENT(x) \vee \neg AGENT(x) \\ & \forall x. ANIMAL(x) \vee \neg ANIMAL(x) \\ & \forall x, y. HAS-MOTHER(x, y) \vee \neg HAS-MOTHER(x, y) \\ & \forall x. AGENT(x) \wedge \neg machine(x) \rightarrow PERSON(x) \\ & AGENT(DaveBowman) \end{aligned}$$

Consider the total HT-model $\mathcal{M}_{HT} = \langle U, I, I \rangle$ of $st(\mathcal{K})$, with U, I as in Example 4. \mathcal{M}_{HT} is not an equilibrium model of $st(\mathcal{K})$, since \mathcal{M}_{HT} is not minimal among all models: $\langle U, I', I \rangle$, with $I' = I \setminus \{machine(DaveBowman)\}$, is a model of $st(\mathcal{K})$. Furthermore, it is easy to verify that $\langle U, I', I' \rangle$ is not a model of $st(\mathcal{K})$.

Now, consider the total HT-model $\mathcal{M}'_{HT} = \langle U, M, M \rangle$, with U as before, and

$$M = \{AGENT(DaveBowman), PERSON(DaveBowman), ANIMAL(DaveBowman), HAS-NAME(DaveBowman, k)\}.$$

\mathcal{M}'_{HT} is an equilibrium model of $st(\mathcal{K})$. Indeed, consider any $M' \subset M$. It is easy to verify that $\langle U, M', M \rangle$ is not a model of $st(\mathcal{K})$. \diamond

7. Discussion

We have seen that quantified equilibrium logic captures three of the main approaches to integrating classical, first-order or DL knowledge bases with nonmonotonic rules under the answer set semantics, in a modular, hybrid approach. However, QEL has a quite distinct flavor from those of r -hybrid, r^+ -hybrid and g -hybrid KBs. Each of these hybrid approaches has a semantics composed of two different components: a classical model on the one hand and an answer set on the other. Integration is achieved by the fact that the classical model serves as a pre-processing tool for the rule base. The style of QEL is different. There is one semantics and one kind of model that covers both types of knowledge. There is no need for any pre-processing of the rule base. In this sense, the integration is more far-reaching. The only distinction we make is that for that part of the knowledge base considered to be classical and monotonic we add a stability condition to obtain the intended interpretation.

There are other features of the approach using QEL that are worth highlighting. First, it is based on a simple minimal model semantics in a known non-classical logic, actually a quantified version of Gödel's 3-valued logic. No reducts are involved and, consequently, the equilibrium construction applies directly to arbitrary first-order theories. The rule part \mathcal{P} of a knowledge base might therefore comprise, say, a nested logic program, where the heads and bodies of rules may be arbitrary boolean formulas, or perhaps rules permitting nestings of the implication connective. While answer sets have recently been defined for such general formulas, more work would be needed to provide

integration in a hybrid KB setting.¹¹ Evidently QEL in the general case is undecidable, so for extensions of the rule language syntax for practical applications one may wish to study restrictions analogous to safety or guardedness. Second, the logic $\mathbf{QHT}_{=}^s$ can be applied to characterise properties such as the strong equivalence of programs and theories (Lifschitz et al., 2007; Pearce and Valverde, 2006). While strong equivalence and related concepts have been much studied recently in ASP, their characterisation in the case of hybrid KBs remains uncharted territory. The fact that QEL provides a single semantics for hybrid KBs means that a simple concept of strong equivalence is applicable to such KBs and characterisable using the underlying logic, $\mathbf{QHT}_{=}^s$. In Section 9 below we describe how $\mathbf{QHT}_{=}^s$ can be applied in this context.

8. Hybrid KBs and the SM operator

Recently, Ferraris, Lee and Lifschitz (Ferraris et al., 2007) have presented a new definition of stable models. It is applicable to sentences or finitely axiomatisable theories in first-order logic. The definition is syntactical and involves an operator SM that resembles parallel circumscription. The stable models of a sentence F are the structures that satisfy a certain second-order sentence, $\text{SM}[F]$. This new definition of stable model agrees with equilibrium logic in the sense that the models of $\text{SM}[F]$ from (Ferraris et al., 2007) are essentially the equilibrium models of F as defined in this article.

We shall now show that by slightly modifying the SM operator we can also capture the NM semantics of hybrid knowledge bases. First, we need to introduce some notation, essentially following (Lifschitz et al., 2007).

If p and q are predicate constants of the same arity then $p = q$ stands for the formula

$$\forall \mathbf{x}(p(\mathbf{x}) \leftrightarrow q(\mathbf{x})),$$

and $p \leq q$ stands for

$$\forall \mathbf{x}(p(\mathbf{x}) \rightarrow q(\mathbf{x})),$$

where \mathbf{x} is a tuple of distinct object variables. If \mathbf{p} and \mathbf{q} are tuples p_1, \dots, p_n and q_1, \dots, q_n of predicate constants then $\mathbf{p} = \mathbf{q}$ stands for the conjunction

$$p_1 = q_1 \wedge \dots \wedge p_n = q_n,$$

and $\mathbf{p} \leq \mathbf{q}$ for

$$p_1 \leq q_1 \wedge \dots \wedge p_n \leq q_n.$$

Finally, $\mathbf{p} < \mathbf{q}$ is an abbreviation for $\mathbf{p} \leq \mathbf{q} \wedge \neg(\mathbf{p} = \mathbf{q})$. The operator $\text{NM}|_{\mathcal{P}}$ defines second-order formulas and the previous notation can be also applied to tuples of predicate variables.

$$\text{NM}|_{\mathcal{P}}[F] = F \wedge \neg \exists \mathbf{u}((\mathbf{u} < \mathbf{p}) \wedge F^*(\mathbf{u})),$$

where \mathbf{p} is the list of all predicate constants $p_1, \dots, p_n \notin \mathcal{L}_{\mathcal{T}}$ occurring in F , \mathbf{u} is a list of n distinct predicate variables u_1, \dots, u_n . The $\text{NM}|_{\mathcal{P}}$ operator works just like in the SM operator from (Ferraris et al., 2007) except that the substitution of predicates p_i is restricted to those not in $\mathcal{L}_{\mathcal{T}}$. Notice that in the definition of $\text{NM}|_{\mathcal{P}}[F]$ the second conjunct specifies the minimality condition on interpretations while the third conjunct involves a translation ‘*’ that provides a reduction of the non-classical here-and-there logic to classical logic. This translation is recursively defined as follows:

¹¹ For a recent extension of answer sets to first-order formulas, see (Ferraris et al., 2007), which is explained in more detail in Section 8.

- $p_i(t_1, \dots, t_m)^* = u_i(t_1, \dots, t_m)$ if $p_i \notin \mathcal{L}_{\mathcal{T}}$;
- $p_i(t_1, \dots, t_m)^* = p_i(t_1, \dots, t_m)$ if $p_i \in \mathcal{L}_{\mathcal{T}}$;
- $(t_1 = t_2)^* = (t_1 = t_2)$;
- $\perp^* = \perp$;
- $(F \odot G)^* = F^* \odot G^*$, where $\odot \in \{\wedge, \vee\}$;
- $(F \rightarrow G)^* = (F^* \rightarrow G^*) \wedge (F \rightarrow G)$;
- $(QxF)^* = QxF^*$, where $Q \in \{\forall, \exists\}$.

(There is no clause for negation here, because $\neg F$ is treated as shorthand for $F \rightarrow \perp$.)

Theorem 15. $\mathcal{M} = \langle U, T \rangle$ is a NM-model of $\mathcal{K} = (\mathcal{T}, \mathcal{P})$ if and only if it satisfies \mathcal{T} and $\text{NM}|_{\mathcal{P}}[\mathcal{P}]$.

We assume here that both \mathcal{T} and \mathcal{P} are finite, so that the operator $\text{NM}|_{\mathcal{P}}$ is well-defined.
Proof:

(\Rightarrow) If $\langle U, T \rangle$, $U = (D, \sigma)$, is a NM-model of $\mathcal{K} = (\mathcal{T}, \mathcal{P})$, then $\langle U, T \rangle \models \mathcal{T}$, and $\langle U, T \rangle \models \mathcal{P}$, and $\langle U, T, T \rangle$ is an equilibrium model of $\mathcal{T} \cup \text{st}(\mathcal{T}) \cup \mathcal{P}$. So we only need to prove that $\langle U, T \rangle \models \neg \exists \mathbf{u}((\mathbf{u} < \mathbf{p}) \wedge P^*(\mathbf{u}))$. For the contradiction, let us assume that

$$\langle U, T \rangle \models \exists \mathbf{u}((\mathbf{u} < \mathbf{p}) \wedge P^*(\mathbf{u}))$$

This means that:

Fact 1: For every $p_i \notin \mathcal{L}_{\mathcal{T}}$, there exists $\bar{p}_i \subset D^n$ such that $(\mathbf{u} < \mathbf{p}) \wedge P^*(\mathbf{u})$ is valid in the structure $\langle U, T \rangle$ where u_i is interpreted as \bar{p}_i .

If we consider the set

$$H = \{p_i(d_1, \dots, d_k) : (d_1, \dots, d_k) \in \bar{p}_i\} \cup \{p_i(d_1, \dots, d_k) : p_i \in \mathcal{L}_{\mathcal{T}}, p_i(d_1, \dots, d_k) \in T\},$$

and u_i is interpreted as \bar{p}_i , then $\mathbf{u} < \mathbf{p}$ is valid in $\langle U, T \rangle$ iff $H \subset T$ and $P^*(\mathbf{u})$ is valid in $\langle U, T \rangle$ iff $\langle U, H \rangle \models P^*(\mathbf{p})$; that is, the *Fact 1* is equivalent to:

$$H \subset T \quad \text{and} \quad \langle U, H \rangle \models P^*(\mathbf{p}) \quad (6)$$

Since $T \setminus H$ does not include predicate symbols of $\mathcal{L}_{\mathcal{T}}$, $\langle U, H, T \rangle \models \mathcal{T} \cup \text{st}(\mathcal{T})$. So, to finish the proof, we need to prove the following for every formula φ :

Fact 2: $\langle U, H, T \rangle, h \models \varphi$ if and only if $\langle U, H \rangle \models \varphi^*(\mathbf{p})$.

As a consequence of *Fact 2*, we have that $\langle U, H, T \rangle \models \mathcal{P}$ and thus $\langle U, H, T \rangle$ is a model of the stable closure of \mathcal{K} , which contradicts that $\langle U, T, T \rangle$ is in equilibrium.

Fact 2 is proved by induction on φ :

(i) If $\varphi = p_i(d_1, \dots, d_k)$, then $\varphi^*(\mathbf{p}) = \varphi$:

$$\langle U, H, T \rangle, h \models \varphi \Leftrightarrow p_i(d_1, \dots, d_k) \in H \Leftrightarrow \langle U, H \rangle \models \varphi^*(\mathbf{p})$$

(ii) Let $\varphi = \psi_1 \wedge \psi_2$ and assume that, for $i = 1, 2$,

$$\langle U, H, T \rangle, h \models \psi_i \quad \text{iff} \quad \langle U, H \rangle \models \psi_i^*(\mathbf{p}). \quad (7)$$

* For $\varphi = \psi_1 \wedge \psi_2$ under assumption (7):

$$\begin{aligned} \langle U, H, T \rangle, h \models \psi_1 \wedge \psi_2 &\Leftrightarrow \langle U, H, T \rangle, h \models \psi_1 \text{ and } \langle U, H, T \rangle, h \models \psi_2 \\ &\Leftrightarrow \langle U, H \rangle \models \psi_1^*(\mathbf{p}) \text{ and } \langle U, H \rangle \models \psi_2^*(\mathbf{p}) \\ &\Leftrightarrow \langle U, H \rangle \models (\psi_1 \wedge \psi_2)^* \end{aligned}$$

* Similarly, for $\varphi = \psi_1 \rightarrow \psi_2$ under assumption (7):

$$\begin{aligned} & \langle U, H, T \rangle, h \models \psi_1 \rightarrow \psi_2 \Leftrightarrow \\ & \Leftrightarrow \langle U, H, T \rangle, t \models \psi_1 \rightarrow \psi_2 \text{ and} \\ & \qquad \text{either } \langle U, H, T \rangle, h \not\models \psi_1 \text{ or } \langle U, H, T \rangle, h \models \psi_2 \\ & \Leftrightarrow \langle U, T \rangle \models \psi_1 \rightarrow \psi_2 \text{ and either } \langle U, H \rangle \not\models \psi_1^*(\mathbf{p}) \text{ or } \langle U, H \rangle \models \psi_2^*(\mathbf{p}) \\ & \Leftrightarrow \langle U, H \rangle \models (\psi_1 \rightarrow \psi_2)^* \end{aligned}$$

(\Leftarrow) If $\langle U, T \rangle, U = (D, \sigma)$, satisfies \mathcal{T} and $\text{NM}|_{\mathcal{P}}[\mathcal{P}]$, then trivially $\langle U, T, T \rangle$ is a here-and-there model of the closure of \mathcal{K} ; we only need to prove that this model is in equilibrium. By contradiction, let us assume that $\langle U, H, T \rangle$ is a here-and-there model of the closure of \mathcal{K} with $H \subset T$. For every $p_i \notin \mathcal{L}_{\mathcal{T}}$, we define

$$\bar{p}_i = \{(d_i, \dots, d_k) : p_i(d_i, \dots, d_k) \in H\}$$

Fact 3: $(\mathbf{u} < \mathbf{p}) \wedge \mathcal{P}^*(\mathbf{u})$ is valid in the structure $\langle U, T \rangle$ if the variables u_i are interpreted as \bar{p}_i .

As a consequence of *Fact 3*, we have that $\exists \mathbf{u}((\mathbf{u} < \mathbf{p}) \wedge \mathcal{P}^*(\mathbf{u}))$ is satisfied by $\langle U, T \rangle$ which contradicts that $\text{NM}|_{\mathcal{P}}[\mathcal{P}]$ is satisfied by the structure.

As in the previous item, *Fact 3* is equivalent to

$$H \subset T \quad \text{and} \quad \langle U, H \rangle \models \mathcal{P}^*(\mathbf{p})$$

The first condition, $H \subset T$, is trivial by definition and the second one is a consequence of *Fact 2*.

□

9. The strong equivalence of knowledge bases

Let us see how the previous results, notably Theorem 13, can be applied to characterise a concept of strong equivalence between hybrid knowledge bases. It is important to know when different reconstructions of a given body of knowledge or state of affairs are equivalent and lead to essentially the same problem solutions. In the case of knowledge reconstructed in classical logic, ordinary logical equivalence can serve as a suitable concept when applied to theories formulated in the same vocabulary. In the case where nonmonotonic rules are present, however, the situation changes: two sets of rules may have the same answer sets yet behave very differently once they are embedded in some larger context. Thus for hybrid knowledge bases one may also like to know that equivalence is robust or modular. A robust notion of equivalence for logic programs will require that programs behave similarly when extended by any further programs. This leads to the following concept of *strong* equivalence: programs Π_1 and Π_2 are strongly equivalent if and only if for any set of rules Σ , $\Pi_1 \cup \Sigma$ and $\Pi_2 \cup \Sigma$ have the same answer sets. This concept of strong equivalence for logic programs in ASP was introduced and studied in (Lifschitz, Pearce and Valverde, 2001) and has given rise to a substantial body of further work looking at different characterisations, new variations and applications of the idea, as well as the development of systems to test for strong equivalence. Strong equivalence has also been defined and studied for logic programs with variables and first-order nonmonotonic theories under the stable model or equilibrium logic semantics (Lin, 2002; Eiter, Fink, Tompits and Woltran, 2005; Lifschitz et al., 2007; Pearce and Valverde, 2006). In equilibrium logic we say that two (first-order) theories Π_1 and Π_2 are strongly equivalent if and only if for any theory Σ , $\Pi_1 \cup \Sigma$ and $\Pi_2 \cup \Sigma$ have the

same equilibrium models (Lifschitz et al., 2007; Pearce and Valverde, 2006). Under this definition we have:

Theorem 16 ((Lifschitz et al., 2007; Pearce and Valverde, 2006)). Two (first-order) theories Π_1 and Π_2 are strongly equivalent if and only if they are equivalent in $\mathbf{QHT}_{=}^s$.

Different proofs of Theorem 16 are given in (Lifschitz et al., 2007) and (Pearce and Valverde, 2006). For present purposes, the proof contained in (Pearce and Valverde, 2006) is more useful. It shows that if theories are not strongly equivalent, the set of formulas Σ such that $\Pi_1 \cup \Sigma$ and $\Pi_2 \cup \Sigma$ do not have the same equilibrium models can be chosen to have the form of implications ($A \rightarrow B$) where A and B are atomic. So if we are interested in the case where Π_1 and Π_2 are sets of rules, Σ can also be regarded as a set of rules. We shall make use of this property below.

In the case of hybrid knowledge bases $\mathcal{K} = (\mathcal{T}, \mathcal{P})$, various kinds of equivalence can be specified, according to whether one or other or both of the components \mathcal{T} and \mathcal{P} are allowed to vary. The following form is rather general.

Definition 8. Let $\mathcal{K}_1 = (\mathcal{T}_1, \mathcal{P}_1)$ and $\mathcal{K}_2 = (\mathcal{T}_2, \mathcal{P}_2)$, be two hybrid KBs sharing the same structural language, ie. $\mathcal{L}_{\mathcal{T}_1} = \mathcal{L}_{\mathcal{T}_2}$. \mathcal{K}_1 and \mathcal{K}_2 are said to be *strongly equivalent* if for any theory \mathcal{T} and set of rules \mathcal{P} , $(\mathcal{T}_1 \cup \mathcal{T}, \mathcal{P}_1 \cup \mathcal{P})$ and $(\mathcal{T}_2 \cup \mathcal{T}, \mathcal{P}_2 \cup \mathcal{P})$ have the same NM-models.

Until further notice, let us suppose that $\mathcal{K}_1 = (\mathcal{T}_1, \mathcal{P}_1)$ and $\mathcal{K}_2 = (\mathcal{T}_2, \mathcal{P}_2)$ are hybrid KBs sharing a common structural language \mathcal{L} .

Proposition 17. \mathcal{K}_1 and \mathcal{K}_2 are strongly equivalent if and only if $\mathcal{T}_1 \cup st(\mathcal{T}_1) \cup \mathcal{P}_1$ and $\mathcal{T}_2 \cup st(\mathcal{T}_2) \cup \mathcal{P}_2$ are logically equivalent in $\mathbf{QHT}_{=}^s$.

Proof: Let $\mathcal{K}_1 = (\mathcal{T}_1, \mathcal{P}_1)$ and $\mathcal{K}_2 = (\mathcal{T}_2, \mathcal{P}_2)$ be hybrid KBs such that $\mathcal{L}_{\mathcal{T}_1} = \mathcal{L}_{\mathcal{T}_2} = \mathcal{L}$. Suppose $\mathcal{T}_1 \cup st(\mathcal{T}_1) \cup \mathcal{P}_1$ and $\mathcal{T}_2 \cup st(\mathcal{T}_2) \cup \mathcal{P}_2$ are logically equivalent in $\mathbf{QHT}_{=}^s$. Clearly $(\mathcal{T}_1 \cup \mathcal{T} \cup st(\mathcal{T}_1 \cup \mathcal{T}) \cup \mathcal{P}_1 \cup \mathcal{P})$ and $(\mathcal{T}_2 \cup \mathcal{T} \cup st(\mathcal{T}_2 \cup \mathcal{T}) \cup \mathcal{P}_2 \cup \mathcal{P})$ have the same $\mathbf{QHT}_{=}^s$ -models and hence the same equilibrium models. Strong equivalence of \mathcal{K}_1 and \mathcal{K}_2 follows by Theorem 13.

For the ‘only-if’ direction, suppose that $\mathcal{T}_1 \cup st(\mathcal{T}_1) \cup \mathcal{P}_1$ and $\mathcal{T}_2 \cup st(\mathcal{T}_2) \cup \mathcal{P}_2$ are not logically equivalent in $\mathbf{QHT}_{=}^s$, so there is an $\mathbf{QHT}_{=}^s$ -model of one of these theories that is not an $\mathbf{QHT}_{=}^s$ of the other. Applying the proof of Theorem 16 given in (Pearce and Valverde, 2006) we can infer that there is a set \mathcal{P} of rules of a simple sort such that the equilibrium models of $\mathcal{T}_1 \cup st(\mathcal{T}_1) \cup \mathcal{P}_1 \cup \mathcal{P}$ and $\mathcal{T}_2 \cup st(\mathcal{T}_2) \cup \mathcal{P}_2 \cup \mathcal{P}$ do not coincide. Hence by Theorem 13 \mathcal{K}_1 and \mathcal{K}_2 are not strongly equivalent. \square

Notice that from the proof of Proposition 17 it follows that the non-strong equivalence of two hybrid knowledge bases can always be made manifest by choosing extensions having a simple form, obtained by adding simple rules to the rule base.

We mention some conditions to test for strong equivalence and non-equivalence.

Corollary 18. (a) \mathcal{K}_1 and \mathcal{K}_2 are strongly equivalent if \mathcal{T}_1 and \mathcal{T}_2 are classically equivalent and \mathcal{P}_1 and \mathcal{P}_2 are equivalent in $\mathbf{QHT}_{=}^s$.
 (b) \mathcal{K}_1 and \mathcal{K}_2 are not strongly equivalent if $\mathcal{T}_1 \cup \mathcal{P}_1$ and $\mathcal{T}_2 \cup \mathcal{P}_2$ are not equivalent in classical logic.

Proof:

(a) Assume the hypothesis. Since $\mathcal{K}_1 = (\mathcal{T}_1, \mathcal{P}_1)$ and $\mathcal{K}_2 = (\mathcal{T}_2, \mathcal{P}_2)$ share a common structural language \mathcal{L} , it follows that $st(\mathcal{T}_1) = st(\mathcal{T}_2) = \mathcal{S}$, say. Since \mathcal{T}_1 and \mathcal{T}_2 are classically equivalent, $\mathcal{T}_1 \cup \mathcal{S}$ and $\mathcal{T}_2 \cup \mathcal{S}$ have the same (total) $\mathbf{QHT}_{=}^s$ -models and so for any \mathcal{T} also $\mathcal{T}_1 \cup \mathcal{T} \cup \mathcal{S} \cup st(\mathcal{T})$ and $\mathcal{T}_2 \cup \mathcal{T} \cup \mathcal{S} \cup st(\mathcal{T})$ have the same (total)

$\mathbf{QHT}_{=}^s$ -models. Since \mathcal{P}_1 and \mathcal{P}_2 are equivalent in $\mathbf{QHT}_{=}^s$ it follows also that for any \mathcal{P} , $(\mathcal{T}_1 \cup \mathcal{T} \cup \mathcal{S} \cup st(\mathcal{T}) \cup \mathcal{P}_1 \cup \mathcal{P})$ and $(\mathcal{T}_2 \cup \mathcal{T} \cup \mathcal{S} \cup st(\mathcal{T}) \cup \mathcal{P}_2 \cup \mathcal{P})$ have the same $\mathbf{QHT}_{=}^s$ -models and hence the same equilibrium models. The conclusion follows by Theorem 13.

(b) Suppose that $\mathcal{T}_1 \cup \mathcal{P}_1$ and $\mathcal{T}_2 \cup \mathcal{P}_2$ are not equivalent in classical logic. Assume again that $st(\mathcal{T}_1) = st(\mathcal{T}_2) = \mathcal{S}$, say. Then clearly $\mathcal{T}_1 \cup \mathcal{S} \cup \mathcal{P}_1$ and $\mathcal{T}_2 \cup \mathcal{S} \cup \mathcal{P}_2$ are not classically equivalent and hence they cannot be $\mathbf{QHT}_{=}^s$ -equivalent. Applying the second part of the proof of Proposition 17 completes the argument. \square

Special cases of strong equivalence arise when hybrid KBs are based on the same classical theory, say, or share the same rule base. That is, $(\mathcal{T}, \mathcal{P}_1)$ and $(\mathcal{T}, \mathcal{P}_2)$ are strongly equivalent if \mathcal{P}_1 and \mathcal{P}_2 are $\mathbf{QHT}_{=}^s$ -equivalent.¹² Analogously:

$$(\mathcal{T}_1, \mathcal{P}) \text{ and } (\mathcal{T}_2, \mathcal{P}) \text{ are strongly equivalent if } \mathcal{T}_1 \text{ and } \mathcal{T}_2 \text{ are classically equivalent.} \quad (8)$$

Let us briefly comment on a restriction that we imposed on strong equivalence, namely that the KBs in question share a common structural language. Intuitively the reason for this is that the structural language $\mathcal{L}_{\mathcal{T}}$ associated with a hybrid knowledge base $\mathcal{K} = (\mathcal{T}, \mathcal{P})$ is part of its identity or ‘meaning’. Precisely the predicates in $\mathcal{L}_{\mathcal{T}}$ are the ones treated classically. In fact, another KB, $\mathcal{K}' = (\mathcal{T}', \mathcal{P})$, where \mathcal{T}' is completely equivalent to \mathcal{T} in classical logic, may have a different semantics if $\mathcal{L}_{\mathcal{T}'}$ is different from $\mathcal{L}_{\mathcal{T}}$. To see this, let us consider a simple example in propositional logic. Let $\mathcal{K}_1 = (\mathcal{T}_1, \mathcal{P}_1)$ and $\mathcal{K}_2 = (\mathcal{T}_2, \mathcal{P}_2)$, be two hybrid KBs where $\mathcal{P}_1 = \mathcal{P}_2 = \{(p \rightarrow q)\}$, $\mathcal{T}_1 = \{(r \wedge (r \vee p))\}$, $\mathcal{T}_2 = \{r\}$. Clearly, \mathcal{T}_1 and \mathcal{T}_2 are classically and even $\mathbf{QHT}_{=}^s$ -equivalent. However, \mathcal{K}_1 and \mathcal{K}_2 are not even in a weak sense semantically equivalent. $st(\mathcal{T}_1) = \{r \vee \neg r; p \vee \neg p\}$, while $st(\mathcal{T}_2) = \{r \vee \neg r\}$. It is easy to check that $\mathcal{T}_1 \cup st(\mathcal{T}_1) \cup \mathcal{P}_1$ and $\mathcal{T}_2 \cup st(\mathcal{T}_2) \cup \mathcal{P}_2$ have different $\mathbf{QHT}_{=}^s$ -models, different equilibrium models and (hence) \mathcal{K}_1 and \mathcal{K}_2 have different NM-models. So we see that without the assumption of a common structural language, the natural properties expressed in Corollary 18 (a) and (8) would no longer hold.

It is interesting to note here that meaning-preserving relations among ontologies have recently become a topic of interest in the description logic community where logical concepts such as that of conservative extension are currently being studied and applied (Ghilardi, Lutz and Wolter, 2006). A unified, logical approach to hybrid KBs such as that developed here should lend itself well to the application of such concepts.

10. Related Work and Conclusions

We have provided a general notion of hybrid knowledge base, combining first-order theories with nonmonotonic rules, with the aim of comparing and contrasting some of the different variants of hybrid KBs found in the literature (Rosati, 2005a; Rosati, 2005b; Rosati, 2006; Heymans et al., 2006). We presented a version of quantified equilibrium logic, QEL, without the unique names assumption, as a unified logical foundation for hybrid knowledge bases. We showed how for a hybrid knowledge base \mathcal{K} there is a natural correspondence between the nonmonotonic models of \mathcal{K} and the equilibrium models of what we call the *stable closure* of \mathcal{K} . This yields a way to capture in QEL the semantics of the g-hybrid KBs of Heymans et al. (Heymans et al., 2006) and the r-hybrid KBs

¹² In (de Bruijn, Pearce, Polleres and Valverde, 2007) it was incorrectly stated in Proposition 7 that this condition was both necessary and sufficient for strong equivalence, instead of merely sufficient.

of Rosati (Rosati, 2005b), where the latter is defined without the UNA but for safe programs. Similarly, the version of QEL with UNA captures the semantics of r-hybrid KBs as defined in (Rosati, 2005a; Rosati, 2006). It is important to note that the aim of this paper was not that of providing new kinds of safety conditions or decidability results; these issues are ably dealt with in the literature reviewed here. Rather our objective has been to show how classical and nonmonotonic theories might be unified under a single semantical model. In part, as (Heymans et al., 2006) show with their reduction of DL knowledge bases to open answer set programs, this can also be achieved (at some cost of translation) in other approaches. What distinguishes QEL is the fact that it is based on a standard, nonclassical logic, $\mathbf{QHT}_{\underline{=}}^s$, which can therefore provide a unified logical foundation for such extensions of (open) ASP. To illustrate the usefulness of our framework we showed how the logic $\mathbf{QHT}_{\underline{=}}^s$ also captures a natural concept of strong equivalence between hybrid knowledge bases.

There are several other approaches to combining languages for Ontologies with nonmonotonic rules which can be divided into two main streams (de Bruijn et al., 2006): approaches which define integration of rules and ontologies (a) by entailment, ie. querying classical knowledge bases through special predicates the rules body, and (b) on the basis of single models, ie. defining a common notion of combined model.

The most prominent of the former kind of approaches are dl-programs (Eiter et al., 2004) and their generalization, HEX-programs (Eiter, Ianni, Schindlauer and Tompits, 2005). Although these approaches both are based on Answer Set programming like our approach, the orthogonal view of integration by entailment can probably not be captured by a simple embedding in QEL. Another such approach which allows querying classical KBs from a nonmonotonic rules language is based on Defeasible Logic (Wang, Billington, Blee and Antoniou, 2004).

As for the second stream, variants of Autoepistemic Logic (de Bruijn, Eiter, Polleres and Tompits, 2007), and the logic of minimal knowledge and negation as failure (MKNF) (Motik and Rosati, 2007) have been recently proposed in the literature. Similar to our approach, both these approaches embed a combined knowledge base in a unifying logic. However, both purchase use modal logics fact syntactically and semantically extend first-order logics. Thus, in these approaches, embedding of the classical part of the theory is trivial, whereas the nonmonotonic rules part needs to be rewritten in terms of modal formulas. Our approach is orthogonal, as we use a non-classical logic where the nonmonotonic rules are trivially embedded, but the stable closure guarantees classical behavior of certain predicates. In addition, the fact that we include the stable closure ensures that the predicates from the classical parts of the theory behave classically, also when used in rules with negation. In contrast, in both modal approaches occurrences of classical predicates are not interpreted classically, as illustrated in the following example.

Example 6. Consider the theory $\mathcal{T} = \{A(a)\}$ and the program $\mathcal{P} = \{r \leftarrow \neg A(b)\}$. We have that there exists an NM-model \mathcal{M} of $(\mathcal{T}, \mathcal{P})$ such that $\mathcal{M} \models A(a), A(b)$ and $\mathcal{M} \not\models A(a), A(b)$, and so r is not entailed. Consider now the embedding τ_{HP} of logic programs into autoepistemic logic (de Bruijn, Eiter, Polleres and Tompits, 2007). We have $\tau_{HP}(\mathcal{P}) = \{\neg \mathbf{L}A(b) \rightarrow r\}$. In autoepistemic logic, $\mathbf{L}A(b)$ is true iff $A(b)$ is included in a *stable expansion* T , which is essentially the set of all entailed formulas, assuming T . We have that $A(b)$ is not entailed from $\mathcal{T} \cup \tau_{HP}(\mathcal{P})$ under any stable expansion, and so $\mathbf{L}A(b)$ is false, and thus r is necessarily true in every model. We thus have that r is a consequence of $\mathcal{T} \cup \tau_{HP}(\mathcal{P})$.

Similar for the hybrid MKNF knowledge bases by (Motik and Rosati, 2007).

As shown by (de Bruijn, Eiter, and Tompits, 2008), adding *classical interpretation ax-*

ioms – essentially a modal version of the stable closure axioms – to the theory $\mathcal{T} \cup \tau_{HP}(\mathcal{P})$ allows one to capture the hybrid knowledge base semantics we considered in this paper.

In future work we hope to consider further aspects of applying QEL to the domain of hybrid knowledge systems. Extending the language with functions symbols and with strong negation is a routine task, since QEL includes these items already. We also plan to consider in the future how QEL can be used to define a catalogue of logical relations between hybrid KBs. Last, but not least, let us mention that in this paper we exclusively dealt with hybrid combinations of classical theories with logic programs under variants of the stable-model semantics. Recently, also hybrid rule combinations based on the well-founded semantics have been proposed by Drabent et al. (Drabent, Henriks-son and Maluszynski, 2007) or Knorr et al. (Knorr, Alferes and Hitzler, 2008), defining an analogous, modular semantics like hybrid knowledge bases considered here. In this context, we plan to investigate whether a first-order version of Partial Equilibrium Logic (Cabalar, Odintsov, Pearce and Valverde, 2006), which has been recently shown to capture the well-founded semantics in the propositional case, can similarly work as a foundation for hybrid rule combinations à la Drabent et al. (Drabent et al., 2007).

We believe that on the long run the general problem addressed by the theoretical foundations layed in this paper could potentially provide essential insights for realistic applications of Ontologies and Semantic Web technologies in general, since for most of these applications current classical ontology languages provide too limited expresivity and the addition of non-monotonic rules is key to overcome these limitations. As an example let us mention the “clash” between the open world assumption in Ontologies and the nonmonotonicity/closed world nature of typical Semantic Web query languages such as SPARQL, which contains nonmonotonic constructs and in fact can be translated to rules with nonmonotonic negation (Polleres, 2007). While some initial works exist in the direction of using SPARQL on top of OWL (Jing, Jeong and Baik, 2009), the foundations and exact semantic treatment of corner cases is still an open problem.¹³ As another example, let us mention mappings between modular ontologies as for instance investigated by (Ensan and Du, 2009); non-monotonic rules could provide a powerful tool to describe mappings between ontologies.

Acknowledgements. Part of the results in this paper are contained, in preliminary form, in the proceedings of the 1st International Conference on Web Reasoning and Rule Systems (RR2007), and in the informal proceedings of the RuleML-06 Workshop on Ontology and Rule Integration. The authors thank the anonymous reviewers of those preliminary versions of the article for their helpful comments. This research has been partially supported by the Spanish MEC (now MCI) under the projects TIC-2003-9001, TIN2006-15455-CO3, and CSD2007-00022, also by the project URJC-CM-2006-CET-0300 and by the European Commission under the projects Knowledge Web (IST-2004-507482) and inContext (IST-034718), as well as by Science Foundation Ireland under Grant No. SFI/08/CE/I1380 (Lion-2).

References

- Baral, C. (2002), *Knowledge Representation, Reasoning and Declarative Problem Solving*, Cambridge University Press.
 Cabalar, P., Odintsov, S. P., Pearce, D. and Valverde, A. (2006), Analysing and extending well-founded and

¹³ One of the authors of the present paper is in fact chairing the W3C SPARQL working group in which at the time of writing of this paper this topic has been being discussed actively.

- partial stable semantics using partial equilibrium logic, in 'Proceedings of the 22nd International Conference on Logic Programming (ICLP 2006)', Vol. 4079 of *Lecture Notes in Computer Science*, Springer, Seattle, WA, USA, pp. 346–360.
- de Bruijn, J., Eiter, T., Polleres, A. and Tompits, H. (2006), On representational issues about combinations of classical theories with nonmonotonic rules, in 'Proceedings of the First International Conference on Knowledge Science, Engineering and Management (KSEM'06)', number 4092 in 'Lecture Notes in Computer Science', Springer-Verlag, Guilin, China.
- de Bruijn, J., Eiter, T., Polleres, A. and Tompits, H. (2007), Embedding non-ground logic programs into autoepistemic logic for knowledge-base combination, in 'Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)', AAAI, Hyderabad, India, pp. 304–309.
- de Bruijn, J., Pearce, D., Polleres, A. and Valverde, A. (2007), Quantified equilibrium logic and hybrid rules, in 'First International Conference on Web Reasoning and Rule Systems (RR2007)', Vol. 4524 of *Lecture Notes in Computer Science*, Springer, Innsbruck, Austria, pp. 58–72.
- de Bruijn, J., Eiter, T., and Tompits, H. (2008), Embedding approaches to combining rules and ontologies into autoepistemic logic, in 'Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR2008)', AAAI, Sydney, Australia, pp. 485–495.
- Drabent, W., Henriksson, J. and Maluszynski, J. (2007), Hybrid reasoning with rules and constraints under well-founded semantics, in 'First International Conference on Web Reasoning and Rule Systems (RR2007)', Vol. 4524 of *Lecture Notes in Computer Science*, Springer, Innsbruck, Austria, pp. 348–357.
- Eiter, T., Fink, M., Tompits, H. and Woltran, S. (2005), Strong and uniform equivalence in answer-set programming: Characterizations and complexity results for the non-ground case, in 'Proceedings of the Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference', pp. 695–700.
- Eiter, T., Ianni, G., Schindlauer, R. and Tompits, H. (2005), A uniform integration of higher-order reasoning and external evaluations in answer-set programming, in 'IJCAI 2005', pp. 90–96.
- Eiter, T., Lukasiewicz, T., Schindlauer, R. and Tompits, H. (2004), Combining answer set programming with description logics for the semantic web, in 'Proceedings of the Ninth International Conference on Principles of Knowledge Representation and Reasoning (KR'04)'.
- Ensan, F. and Du, W. (2009), A knowledge encapsulation approach to ontology modularization. *Knowledge and Information Systems Online First*, to appear.
- Ferraris, P., Lee, J. and Lifschitz, V. (2007), A new perspective on stable models, in 'Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)', AAAI, Hyderabad, India, pp. 372–379.
- Ghilardi, S., Lutz, C. and Wolter, F. (2006), Did I damage my ontology: A Case for Conservative Extensions of Description Logics, in 'Proceedings of the Tenth International Conference on Principles of Knowledge Representation and Reasoning (KR'06)', pp. 187–197.
- Heymans, S. (2006), Decidable Open Answer Set Programming, PhD thesis, Theoretical Computer Science Lab (TINF), Department of Computer Science, Vrije Universiteit Brussel, Brussels, Belgium.
- Heymans, S., Nieuwenborgh, D. V. and Vermeir, D. (2005), Guarded Open Answer Set Programming, in '8th International Conference on Logic Programming and Non Monotonic Reasoning (LPNMR 2005)', volume 3662 in 'LNAI', Springer, pp. 92–104.
- Heymans, S., Predoiu, L., Feier, C., de Bruijn, J. and van Nieuwenborgh, D. (2006), G-hybrid knowledge bases, in 'Workshop on Applications of Logic Programming in the Semantic Web and Semantic Web Services (ALPSWS 2006)'.
- Jing, Y., Jeong, D. and Baik, D.-K. (2009), SPARQL graph pattern rewriting for OWL-DL inference queries. *Knowledge and Information Systems* **20**, pp. 243–262.
- Knorr, M., Alferes, J. and Hitzler, P. (2008), A coherent well-founded model for hybrid mknf knowledge bases, in '18th European Conference on Artificial Intelligence (ECAI2008)', volume 178 in 'Frontiers in Artificial Intelligence and Applications', IOS Press, pp. 99–103.
- Lifschitz, V., Pearce, D. and Valverde, A. (2001), 'Strongly equivalent logic programs', *ACM Transactions on Computational Logic* **2**(4), 526–541.
- Lifschitz, V., Pearce, D. and Valverde, A. (2007), A characterization of strong equivalence for logic programs with variables, in '9th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR)', Vol. 4483 of *Lecture Notes in Computer Science*, Springer, Tempe, AZ, USA, pp. 188–200.
- Lifschitz, V. and Woo, T. (1992), Answer sets in general nonmonotonic reasoning (preliminary report), in B. Nebel, C. Rich and W. Swartout, eds, 'KR'92. Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference', Morgan Kaufmann, San Mateo, California, pp. 603–614.
- Lin, F. (2002), Reducing strong equivalence of logic programs to entailment in classical propositional logic, in 'Proceedings of the Eighth International Conference on Principles of Knowledge Representation and Reasoning (KR'02)', pp. 170–176.

- Motik, B. and Rosati, R. (2007), A faithful integration of description logics with logic programming, in 'Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)', AAAI, Hyderabad, India, pp. 477–482.
- Pearce, D. (1997), A new logical characterization of stable models and answer sets, in 'Proceedings of NMELP 96', Vol. 1216 of *Lecture Notes in Computer Science*, Springer, pp. 57–70.
- Pearce, D. (2006), 'Equilibrium logic', *Annals of Mathematics and Artificial Intelligence* **47** 3–41.
- Pearce, D. and Valverde, A. (2005), 'A first-order nonmonotonic extension of constructive logic', *Studia Logica* **80**, 321–246.
- Pearce, D. and Valverde, A. (2006), Quantified equilibrium logic, Technical report, Universidad Rey Juan Carlos. in press.
- Polleres, A. (2007), From SPARQL to Rules (and back), in 'WWW 2007', pp. 787–796.
- Rosati, R. (2005a), 'On the decidability and complexity of integrating ontologies and rules', *Journal of Web Semantics* **3**(1), 61–73.
- Rosati, R. (2005b), Semantic and computational advantages of the safe integration of ontologies and rules, in 'Proceedings of the Third International Workshop on Principles and Practice of Semantic Web Reasoning (PPSWR 2005)', Vol. 3703 of *Lecture Notes in Computer Science*, Springer, pp. 50–64.
- Rosati, R. (2006), $\mathcal{DL}+log$: Tight integration of description logics and disjunctive datalog, in 'Proceedings of the Tenth International Conference on Principles of Knowledge Representation and Reasoning (KR'06)', pp. 68–78.
- van Dalen, D. (1983), *Logic and Structure*, Springer.
- Wang, K., Billington, D., Blee, J. and Antoniou, G. (2004), Combining description logic and defeasible logic for the semantic web, in 'Proceedings of the Third International Workshop Rules and Rule Markup Languages for the Semantic Web (RuleML 2004)', Vol. 3323 of *Lecture Notes in Computer Science*, Springer, pp. 170–181.