

Paraconsistent Hybrid Theories*

Michael Fink

Institut für Informationssysteme
Technische Universität Wien, Austria
fink@kr.tuwien.ac.at

Abstract

We consider the problem of reasoning from inconsistent hybrid theories, i.e., combinations of a structural part given by a classical first order theory (e.g., an ontology) and a rules part as a set of declarative logic program rules (under answer-set semantics). Paraconsistent reasoning is achieved by defining an appropriate semantics, so-called paraconsistent semi-equilibrium model semantics for such hybrid theories. Appropriateness of the semantics is established with respect to desirable properties attesting design objectives, such as to generalize the underlying semantics in case of consistency, as well as to generalize existing paraconsistent semantics for the individual parts. A complexity analysis of corresponding reasoning tasks complements these results.

Introduction

Paraconsistent Reasoning is an important means to tolerate inconsistencies in knowledge representation (see e.g., Hunter 1998; Bertossi, Hunter, and Schaub 2005; Arieli, Avron, and Zamansky 2011a). In providing non-trivial semantics to contradictory pieces of knowledge, it not only allows to analyze reasons for inconsistency: reasoning systems also stay operable in case of contradictions, in the sense that they still can provide reasonable answers to queries. Therefore, developing paraconsistent semantics and studying their logical and computational properties for prominent knowledge-representation and reasoning (KRR) formalisms received considerable attention. Recent developments include, for instance, paraconsistent semantics for *Description Logics* (DLs) (Ma, Hitzler, and Lin 2008; 2007), as well as for nonmonotonic formalisms like *Answer-Set Programming* (ASP) (Eiter, Fink, and Moura 2010; Odintsov and Pearce 2005; Alcântara, Damásio, and Pereira 2004; Sakama and Inoue 1995).

Another line of research pursued actively in KRR is the combination of conceptual knowledge bases with rule-based formalisms, in order to establish powerful reasoning reasoning systems utilizing semantic domain information (of-

ten from the Web) beyond classifying data. In particular, combinations of DLs with rules under declarative nonmonotonic semantics have been studied intensively, both under so-called tight semantic couplings (de Bruijn et al. 2007; Rosati 2006; Motik and Rosati 2010; Kifer 2005; Grosz et al. 2003), as well as loose couplings (Eiter et al. 2008; Heymans et al. 2010; Lukasiewicz 2007). Such combinations have applications in various application domains, e.g., steel product management systems and quality control systems for automobile design are two applications addressed within the EU FP7 project ONTORULE¹.

In this paper, we focus on tight combinations under ASP semantics, tackling the issue of developing a paraconsistent semantics for this setting, based on its prevailing logical underpinning by the so-called logic of Here-and-There (see, e.g., Pearce and Valverde 2008). In general, combining different pieces of knowledge is more prone to cause contradiction than inconsistency occurring in the individual representations. For instance, consider the following example.

Example 1 Consider a hybrid theory $(\mathcal{T}, \mathcal{P})$ consisting of a structural part $\mathcal{T}$ representing information about cars and persons designing and engineering them:

$$\begin{aligned} \forall x \text{ car}(x) &\rightarrow \text{product}(x) \\ \forall x \text{ engineer}(x) &\rightarrow \text{person}(x) \\ \forall x \forall y \text{ designs}(x, y) &\rightarrow \text{engineer}(x) \\ \text{person}(p) \quad \text{car}(c) & \end{aligned}$$

The rules part $\mathcal{P}$ takes into account information on car assembly at a certain plant (here) of the company under consideration, making use of the structural predicates:

$$\begin{aligned} \forall x (\text{product}(x) \wedge \neg \text{built}(x, \text{here})) &\rightarrow \sim \text{car}(x) \\ \forall x \forall y (\text{person}(x) \wedge \text{product}(y) \wedge \neg \text{assembles}(x, y)) &\rightarrow \\ \text{designs}(x, y) & \\ \forall x \forall y (\text{engineer}(x) \wedge \text{car}(y)) &\rightarrow \text{assembles}(x, y) \end{aligned}$$

Intuitively, the first rule expresses that all cars built by the company are built here. This is done by applying the closed world assumption to the information on products built here, i.e., products not known to be built here are not a car. Moreover, the local information for the plant includes that persons not known to assemble a product are designers of the product (second rule), and that engineers assemble cars

*This work was supported by the EU FP7 project OnToRule ICT-2009-231875, the Austrian Science Fund (FWF) grant P20840, and by the Vienna Science and Technology Fund (WWTF) through project ICT08-020.

Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹www.ontorule-project.eu

(third rule). Also note the use of two types of negation: ‘ $\neg$ ’ for default negation and ‘ $\sim$ ’ for strong negation.

While both parts taken individually are consistent pieces of knowledge (the rules part also if the concrete instance data is taken into account, i.e., that c is a car, and p a person), their combination is both inconsistent (wrt. c being a car or not) and incoherent (wrt. p assembling c , or not).

The above example illustrates important aspects to be taken into account when analyzing inconsistency of hybrid theories. First of all, there may be different reasons for inconsistency, which is reflected in the distinction between inconsistency and incoherence (terms otherwise mainly used synonymously). We use ‘inconsistency’ in case a theory does not allow for a consistent model because this would require an atom and its strong negation to be true. E.g., this is the case for $car(c)$ in our example, due to the fact that cars are products (first formula of $\mathcal{T}$), and since the first rule of $\mathcal{P}$ applies for c . In the presence of nonmonotonic rules, a hybrid theory may also lack a (consistent) model due to cyclic dependencies through default negation. In this case, we say that the theory (or a program) is incoherent. For instance, in our example $assembles(p, c)$ depends on $\neg assembles(p, c)$ via $engineer(p)$ and $designes(p, c)$ (cf. the second and third rule of $\mathcal{P}$, and the second formula of $\mathcal{T}$).

Since both parts in the example are consistent when considered separately, it also shows that methods of consistency maintenance on the individual parts are of no avail. Similarly, solving inconsistency and incoherence separately does not solve the problem in many cases, due to logical dependencies. Furthermore, approximative semantics—usually geared towards efficiency in computation—intuitively treat all atoms involved in inconsistency as undefined, making them less attractive for inconsistency analysis.

Despite the fact that dealing with inconsistency in such cases is an important problem for applications, the problem has been unexplored to the best of our knowledge (with the single exception of recent work on MKNF knowledge bases (Huang, Li, and Hitzler 2011), i.e. a tight coupling in a different syntactic and semantic setting, however;² see the final section for further discussion).

Our work addresses this open problem, aiming at a paraconsistent semantics for tight couplings under ASP semantics. Thereby, one objective is to retain answer sets in case of consistency. Moreover, in case of inconsistency or incoherence, the semantics should generalize paraconsistent (Alcântara, Damásio, and Pereira 2004; Sakama and Inoue 1995) and paraconsistent (Eiter, Fink, and Moura 2010) ASP semantics on the rules part, as well as the paraconsistent semantics underlying (Ma, Hitzler, and Lin 2007; 2008) on the structural part.

Our respective contributions are summarized as follows:

- We consider a general setting of hybrid knowledge bases, so-called *hybrid theories*, and define *paraconsistent semi-equilibrium model (pseq-) semantics* on semantic structures reflecting the reasons of inconsistency (i.e., without

²Although originally based on the stable model semantics, by resorting to the logic of minimal knowledge and negation as failure (MKNF) it is based on a nonmonotonic modal logic.

intermingling the concerns of paraconsistency and para-coherence).

- We study semantic properties, and corresponding formal results also establish our objectives: that the semantics generalizes (1) existing paraconsistent semantics on the individual parts in case of inconsistency, and (2) the underlying ASP semantics. By retaining benign properties of paraconsistent ASP semantics, pseq-semantics is more accurate for inconsistency analysis than potential alternatives tailored to efficient query answering, such as generalizing well-founded semantics.
- We report complexity results for predominant paraconsistent reasoning tasks including consistency (model existence), as well as brave and cautious reasoning. The study encompasses combined as well as data complexity, relying on usual assumptions that guarantee decidability.

In summary, our results assure that pseq-semantics is a faithful extension and combination of recent work on paraconsistent reasoning in DLs with paraconsistent and paraconsistent ASP. It thus provides a paraconsistent semantics for hybrid theories with intuitive computational properties.

Preliminaries

Throughout the paper, let us consider function-free first order languages $\mathcal{L} = \langle C, P \rangle$ over a set C of *constant* symbols, and a set P of *predicate* symbols. The notions of well-formed $\mathcal{L}$ -formulas, atomic $\mathcal{L}$ -formulas, $\mathcal{L}$ -sentences and $\mathcal{L}$ -theories are as usual (including the symbol ‘ $\perp$ ’ for falsity). We will sometimes drop the prefix $\mathcal{L}$ if this is unambiguous, i.e., when the language is clear from the context or arbitrary.

Given a non-empty set D of domain objects, by $At(D, P)$ we denote the set of ground atomic sentences of $\mathcal{L}' = \langle D, P \rangle$. Any subset of $At(D, P)$ is called an $\mathcal{L}$ -interpretation over D . We consider *classical $\mathcal{L}$ -structures* as tuples $\mathcal{M} = \langle (D, \sigma), I \rangle$, where I is an $\mathcal{L}$ -interpretation over D and σ is an *assignment*, i.e., a mapping $\sigma: C \cup D \rightarrow D$ such that $\sigma(d) = d$ for all $d \in D$. If $D = C$ and $\sigma = id$, then $\mathcal{M}$ is called a *Herbrand structure*. Let $\sigma|_C$ be the restriction of σ to constants from C , then the *parameter names assumption (PNA)* applies if $\sigma|_C$ is surjective, i.e., there are no unnamed individuals in D ; the *unique names assumption (UNA)* applies if $\sigma|_C$ is injective; and if both PNA and UNA apply, then the *standard names assumption (SNA)* applies, i.e. $\sigma|_C$ is a bijection.

Four-valued Logic. A *four-valued classical $\mathcal{L}$ -structure* is a tuple $\mathcal{M} = \langle (D, \sigma), I_t, I_f \rangle$, where both substructures, $\langle (D, \sigma), I_t \rangle$ and $\langle (D, \sigma), I_f \rangle$, are classical $\mathcal{L}$ -structures. Intuitively, these two substructures serve the purpose of decoupling the evaluation of truth (t) from the evaluation of falsity (f)—hence the subscripts—paving the way for assigning a designated truth value ($\ddot{t}$) to contradictory formulas, and thus their paraconsistent treatment. More formally, a four-valued satisfaction relation $\mathcal{M} \models_4 \varphi$, for sentences φ over $At(C \cup D, P)$ is defined recursively as follows. Let $k \in \{t, f\}$, and consider arbitrary constants t_i from $C \cup D$,

respectively any predicate p from P , then

$$\begin{aligned}\mathcal{M}, \mathbf{t} \models t_1 = t_2 & \quad \text{iff} \quad \sigma(t_1) = \sigma(t_2); \\ \mathcal{M}, \mathbf{f} \models t_1 = t_2 & \quad \text{iff} \quad \sigma(t_1) \neq \sigma(t_2); \\ \mathcal{M}, k \models p(t_1, \dots, t_n) & \quad \text{iff} \quad p(\sigma(t_1), \dots, \sigma(t_n)) \in I_k.\end{aligned}$$

This definition for ground atomic sentences is extended recursively by the following, where given k and an expression of the form x/y , x applies if $k = \mathbf{t}$, and y otherwise:

- $\mathcal{M}, \mathbf{t} \not\models \perp$ and $\mathcal{M}, \mathbf{f} \models \perp$ for all $\mathcal{L}$ -structures $\mathcal{M}$;
- $\mathcal{M}, k \models \sim \varphi$ iff $\mathcal{M}, \bar{k} \models \varphi$ ($\bar{k} = \mathbf{f}$ if $k = \mathbf{t}$, $\mathbf{t}$ otherwise);
- $\mathcal{M}, k \models \varphi \wedge \psi$ iff $\mathcal{M}, k \models \varphi$ and/or³ $\mathcal{M}, k \models \psi$;
- $\mathcal{M}, k \models \varphi \vee \psi$ iff $\mathcal{M}, k \models \varphi$ or/and $\mathcal{M}, k \models \psi$;
- $\mathcal{M}, \mathbf{t} \models \varphi \rightarrow \psi$ iff $\mathcal{M}, \mathbf{t} \models \varphi$ implies $\mathcal{M}, \mathbf{t} \models \psi$;
- $\mathcal{M}, \mathbf{f} \models \varphi \rightarrow \psi$ iff $\mathcal{M}, \mathbf{t} \models \varphi$ and $\mathcal{M}, \mathbf{f} \models \psi$;
- $\mathcal{M}, k \models \forall x \varphi(x)$ iff $\mathcal{M}, k \models \varphi(d)$ for all/some $d \in D$;
- $\mathcal{M}, k \models \exists x \varphi(x)$ iff $\mathcal{M}, k \models \varphi(d)$ for some/all $d \in D$.

Note that ‘ $\sim$ ’ denotes negation, and the inclusion of a symbol for falsity ($\perp$), as well as an implication connective ($\rightarrow$) representing so-called *internal implication*. Truth and other implication connectives such as *material implication* and *strong implication* are definable (e.g., by $\sim \perp$, $\sim \varphi \vee \psi$, and $(\varphi \rightarrow \psi) \wedge (\sim \varphi \rightarrow \sim \psi)$, respectively).

We say that $\mathcal{M}$ is a *four-valued model* of a sentence φ , in symbols $\mathcal{M} \models_4 \varphi$, iff $\mathcal{M}, \mathbf{t} \models \varphi$. The *valuation* of φ wrt. a four-valued structure $\mathcal{M}$, denoted as $v(\mathcal{M}, \varphi)$, assigns one of four truth values true $\mathbf{t}$, false $\mathbf{f}$, contradictory $\ddot{\top}$, or undefined $\ddot{\perp}$ to φ as follows:

- $v(\mathcal{M}, \varphi) = \mathbf{t}$ iff $\mathcal{M}, \mathbf{t} \models \varphi$ and $\mathcal{M}, \mathbf{f} \not\models \varphi$;
- $v(\mathcal{M}, \varphi) = \mathbf{f}$ iff $\mathcal{M}, \mathbf{t} \not\models \varphi$ and $\mathcal{M}, \mathbf{f} \models \varphi$;
- $v(\mathcal{M}, \varphi) = \ddot{\top}$ iff $\mathcal{M}, \mathbf{t} \models \varphi$ and $\mathcal{M}, \mathbf{f} \models \varphi$;
- $v(\mathcal{M}, \varphi) = \ddot{\perp}$ iff $\mathcal{M}, \mathbf{t} \not\models \varphi$ and $\mathcal{M}, \mathbf{f} \not\models \varphi$.

Note that by the above definition of model, $\mathbf{t}$ and $\ddot{\top}$ are designated truth values.

TODO: Add intuitive explanation.

As usual, a sentence φ is valid if it is true in all four-valued $\mathcal{L}$ -structures, denoted by $\models_4 \varphi$. A sentence φ is a four-valued consequence of a set of sentences Γ , denoted $\Gamma \models_4 \varphi$, if every four-valued model of Γ is a four-valued model of φ .

Quantified Here-and-There and Equilibrium Logic. As a logical basis for the nonmonotonic answer-set semantics of logic programs and hybrid knowledge bases, we build on *Quantified Equilibrium Logic (QEL)*, following Lifschitz, Pearce, and Valverde (2007), and Pearce and Valverde (2008). QEL is the nonmonotonic extension of a monotonic intuitionistic base logic called *Quantified Logic of Here-and-There (QHT)*.

Again we restrict our attention to the function-free languages introduced at the beginning of this section, including

³Recall that ‘and’ applies if $k = \mathbf{t}$, while ‘or’ applies if $k = \mathbf{f}$ (reflecting the dual nature of evaluating falsity).

a single (intuitionistic) negation denoted by ‘ $\neg$ ’. A *here-and-there $\mathcal{L}$ -structure* (with static domain), or *QHT $\mathcal{L}$ -structure*, is a tuple $\mathcal{M} = \langle (D, \sigma), I_h, I_t \rangle$, where again $\langle (D, \sigma), I_h \rangle$ and $\langle (D, \sigma), I_t \rangle$ are classical $\mathcal{L}$ -structures, however such that $I_h \subseteq I_t$.

Here-and-there $\mathcal{L}$ -structures $\mathcal{M}$ are similar to four-valued $\mathcal{L}$ -structures in having two parts, or components, that are interpretations—indexed and identified with h and t . In the case of the logic of here-and-there however, they 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 $h \leq t$. Every world $w \in \{h, t\}$ verifies a set of atoms I_w over the expanded language with domain D . The structures are termed static, because the same domain serves both worlds.

Again, the associated semantic consequence relation is defined recursively, taking into account the two components, i.e., worlds, as well as the fact that whatever is verified at the here world h remains true at the there world t . More precisely, let $w \in \{h, t\}$, and consider arbitrary constants t_i from $C \cup D$, respectively any predicate p from P , then

$$\begin{aligned}\mathcal{M}, w \models t_1 = t_2 & \quad \text{iff} \quad \sigma(t_1) = \sigma(t_2); \\ \mathcal{M}, w \models p(t_1, \dots, t_n) & \quad \text{iff} \quad p(\sigma(t_1), \dots, \sigma(t_n)) \in I_w.\end{aligned}$$

The recursive extension of this definition to sentences φ is given by:

- $\mathcal{M}, w \not\models \perp$ for all here-and-there $\mathcal{L}$ -structures $\mathcal{M}$;
- $\mathcal{M}, w \models \varphi \wedge \psi$ 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 $\mathcal{M}, w \models \psi$;
- $\mathcal{M}, w \models \varphi \rightarrow \psi$ iff $\mathcal{M}, w' \models \varphi$ implies $\mathcal{M}, w' \models \psi$ for all $w' \geq w$;
- $\mathcal{M}, w \models \forall x \varphi(x)$ iff $\mathcal{M}, w' \models \varphi(d)$ for all $d \in D$ and $w' \geq w$;
- $\mathcal{M}, w \models \exists x \varphi(x)$ iff $\mathcal{M}, w \models \varphi(d)$ for some $d \in D$.

Negation $\neg \varphi$ is considered an abbreviation for $\varphi \rightarrow \perp$, and therefore omitted in the above definition. It is easily verified (and well-known) that $\mathcal{M}, w \models \neg \varphi$ iff $\mathcal{M}, t \not\models \varphi$.

The QHT $\mathcal{L}$ -models of a sentence φ are those QHT $\mathcal{L}$ -structures $\mathcal{M}$, that satisfy φ at both worlds, in symbols $\mathcal{M} \models_{\text{QHT}} \varphi$ iff $\mathcal{M}, w \models \varphi$ for every $w \in \{h, t\}$. A sentence φ is a consequence of a set of sentences Γ , denoted $\Gamma \models_{\text{QHT}} \varphi$, if every QHT $\mathcal{L}$ -model of Γ is a model of φ ; it is valid iff every QHT $\mathcal{L}$ -structure is a model.

The logic thus defined is called *Quantified Here-and-There Logic with static domains and decidable equality*, and has been denoted by $\text{SQHT}^=$ in Lifschitz, Pearce, and Valverde (2007). We simply refer to it as QHT, though. For a complete axiomatisation of QHT, based on the axioms and rules of first-order intuitionistic logic (van Dalen 1983), cf., e.g., Lifschitz, Pearce, and Valverde (2007).

The logic of Here-and-There serves as a basis to characterize stable model semantics, or answer-set semantics, of logic programs. Such a characterization builds on a selection of QHT $\mathcal{L}$ -models through an additional minimization criterion yielding so-called equilibrium models as follows.

Definition 1 Given a $\mathcal{L}$ -theory Γ , a total QHT $\mathcal{L}$ -structure $\mathcal{M} = \langle (D, \sigma), T, T \rangle$ is called an equilibrium model of Γ iff $\mathcal{M} \models \Gamma$ and $\mathcal{M}' \not\models \Gamma$, for all $\mathcal{M}' = \langle (D, \sigma), H, T \rangle$, such that $H \subset T$.

Note that the above definition does not rely on any of the domain assumptions, PNA, UNA, or SNA introduced before. It thus represents stable models, or answer-set semantics in its most general form, also called generalized open answer-set semantics, as it has been developed and studied more recently. On the other hand, traditionally answer-set semantics for a logic program has been defined in terms of Herbrand models of its grounding. This semantics is obtained by restricting to equilibrium models among Herbrand structures.

Hybrid Knowledge Bases. We study hybrid knowledge bases as defined in de Bruijn et al. (2007), where it has also been shown that the corresponding semantic treatment in terms of QHT equilibrium models appropriately generalizes several previous approaches (Rosati 2005b; 2006; Heymans et al. 2008).

More formally, we consider a (single) theory over a function-free language $\mathcal{L}$, which however is composed of two parts. In particular, a *hybrid knowledge base* $\mathcal{K} = (\mathcal{T}, \mathcal{P})$ over a language $\mathcal{L} = \langle C, P_{\mathcal{T}} \cup P_{\mathcal{P}} \rangle$, where $P_{\mathcal{T}} \cap P_{\mathcal{P}} = \emptyset$, consists of a classical first-order theory $\mathcal{T}$ (also called the *structural* part of $\mathcal{K}$) and a program $\mathcal{P}$ (also called *rules* part of $\mathcal{K}$). Thereby, the structural part is over the language $\mathcal{L}_{\mathcal{T}} = \langle C, P_{\mathcal{T}} \rangle$, whereas the program part is over the language $\mathcal{L}$. Note that thus both parts share a single set of constants, and that the predicate names allowed in $\mathcal{P}$ are a superset of the predicate names in $\mathcal{L}_{\mathcal{T}}$. Additionally, we use $\mathcal{L}_{\mathcal{P}} = \langle C, P_{\mathcal{P}} \rangle$ to refer to the language built from predicate names that are allowed in $\mathcal{P}$ only.

The intuition of considering such a composition of two parts and a respective split of the language, is to interpret the predicates in $\mathcal{L}_{\mathcal{T}}$, hence the structural part, classically, whereas the predicates in $\mathcal{L}_{\mathcal{P}}$ are interpreted nonmonotonically. More specifically, let $\mathcal{K} = (\mathcal{T}, \mathcal{P})$ be a hybrid knowledge base, then $\mathcal{T} \cup \mathcal{P} \cup st(\mathcal{T})$ is called the *stable closure* of $\mathcal{K}$, where $st(\mathcal{T}) = \{\forall x(p(x) \vee \neg p(x)) \mid p \in \mathcal{L}_{\mathcal{T}}\}$. The following property of $st(\mathcal{T})$ guarantees that the stable closure behaves as intended: for all $\varphi \in \mathcal{T}$ it holds that $st(\mathcal{T}) \models_{\text{QHT}} \neg\neg\varphi \rightarrow \varphi$. Thus, the structural part behaves classically (taking into account that $\varphi \rightarrow \neg\neg\varphi$ is valid in QHT, the former yields excluded middle). Consequently, the semantics of a hybrid knowledge base is given by the equilibrium models of its stable closure. For further details we refer to de Bruijn et al. (2007).

Definition 2 Given a hybrid knowledge base $\mathcal{K} = (\mathcal{T}, \mathcal{P})$ over $\mathcal{L}$, a total QHT $\mathcal{L}$ -structure $\mathcal{M} = \langle (D, \sigma), T, T \rangle$ is called an equilibrium model of $\mathcal{K}$ iff it is an equilibrium model of the stable closure of $\mathcal{K}$, ie. of $\mathcal{T} \cup \mathcal{P} \cup st(\mathcal{T})$.

Static Quantified Paraconsistent Nelson Logic

In order to deal with hybrid knowledge bases paraconsistently, we have to slightly weaken the stable closure, which

otherwise trivializes in case of contradiction, as, for instance, it is the case in Example 1. Intuitively speaking, we aim at separating the semantic treatment (and hence the closure) of positive structural atoms from that of negative structural atoms. For this purpose, we use a second type of negation: in addition to default or intuitionistic negation ($\neg$), strong negation ($\sim$) is allowed. Hence, we consider such a syntactically enriched language, and rather than focusing on hybrid knowledge bases a priori, we start providing the necessary foundations for theories in general.

The corresponding logical basis is a paraconsistent version of Nelson logic, i.e., with explicit strong negation, and an additional (definable) intuitionistic negation. A respective nine-valued logic, called N_9 has been defined, axiomatized, and semantically characterized in the propositional case by means of so-called Routley models by Odintsov and Pearce (2005), based on ideas developed by Routley for constructive logics with strong negation, including paraconsistent versions. In particular, confer Routley (1974) for one of the first paraconsistent extensions of Nelson logic (in the propositional case), while Almkudad and Nelson (1984) provide an early first-order treatment.

For our purpose, we require a first-order version of N_9 that we have developed resorting to the ideas applied for QHT, specifically the restriction to static domains for all worlds under consideration. This seems to be the most natural choice for an underlying logic if one aims at extending ASP semantics, given that the logic of Here-and-There, respectively Nelson logic, correspondingly serve its logical underpinning. The resulting logic is called QN_9 and, for space reasons, we refer to Fink (2012) and Fink et al. (2011) for a more detailed account of the logic (including a sound and complete axiomatization, as well as semantic properties including normal forms). Subsequently, we rather present an alternative semantic characterization of QN_9 equivalent to Routley model semantics, which more suitably reflects our intentions and objectives when used as a basis for defining a paraconsistent semantics for hybrid theories.

Definition 3 Given language $\mathcal{L}$, a nine-valued $\mathcal{L}$ -structure (with static domain) is a five-tuple $\mathcal{M} = \langle (D, \sigma), I_{h,t}, I_{h,f}, I_{t,t}, I_{t,f} \rangle$, where $\langle (D, \sigma), I_{w,k} \rangle$ are classical $\mathcal{L}$ -structures, such that $I_{h,k} \subseteq I_{t,k}$, for all $w \in \{h, t\}$ and $k \in \{t, f\}$.

Generalizing from four-valued logic, we define for a world w from $\{h, t\}$ and $k \in \{t, f\}$:

$$\begin{aligned} \mathcal{M}, w, t &\models t_1 = t_2 \quad \text{iff} \quad \sigma(t_1) = \sigma(t_2); \\ \mathcal{M}, w, f &\models t_1 = t_2 \quad \text{iff} \quad \sigma(t_1) \neq \sigma(t_2); \\ \mathcal{M}, w, k &\models p(t_1, \dots, t_n) \quad \text{iff} \quad p(\sigma(t_1), \dots, \sigma(t_n)) \in I_{w,k}. \end{aligned}$$

and recursively extend this definition (for expressions of the form x/y , x applies if $k = t$, and y otherwise):

- $\mathcal{M}, w, t \not\models \perp$ and $\mathcal{M}, w, f \models \perp$ for all $\mathcal{L}$ -structures $\mathcal{M}$;
- $\mathcal{M}, w, k \models \sim\varphi$ iff $\mathcal{M}, w, \bar{k} \models \varphi$ ($\bar{k} = f$ if $k = t$, and t otherwise);
- $\mathcal{M}, w, k \models \varphi \wedge \psi$ iff $\mathcal{M}, w, k \models \varphi$ and/or $\mathcal{M}, w, k \models \psi$;
- $\mathcal{M}, w, k \models \varphi \vee \psi$ iff $\mathcal{M}, w, k \models \varphi$ or/and $\mathcal{M}, w, k \models \psi$;

- $\mathcal{M}, w, \mathbf{t} \models \varphi \rightarrow \psi$ iff $\mathcal{M}, w', \mathbf{t} \models \varphi$ implies $\mathcal{M}, w', \mathbf{t} \models \psi$ for all $w' \geq w$;
- $\mathcal{M}, w, \mathbf{f} \models \varphi \rightarrow \psi$ iff $\mathcal{M}, w, \mathbf{t} \models \varphi$ and $\mathcal{M}, w, \mathbf{f} \models \psi$;
- $\mathcal{M}, w, k \models \forall x \varphi(x)$ iff $\mathcal{M}, w', k \models \varphi(d)$ for all/some $d \in D$ and $w' \geq w$;
- $\mathcal{M}, w, k \models \exists x \varphi(x)$ iff $\mathcal{M}, w, k \models \varphi(d)$ for some/all $d \in D$.

Definition 4 Let $\mathcal{M} = \langle (D, \sigma), I_{h,t}, I_{h,f}, I_{t,t}, I_{t,f} \rangle$ be a nine-valued $\mathcal{L}$ -structure, and let α be a $\mathcal{L}$ -sentence. Then, $\mathcal{M}$ is a nine-valued $\mathcal{L}$ -model of α , in symbols $\mathcal{M} \models_{\text{QN}_9} \alpha$ if and only-if $\mathcal{M}, w, \mathbf{t} \models \alpha$ for all $w \in \{h, t\}$.

The notion of semantic consequence is extended to theories as usual; the same holds for the definition of validity.

Example 2 Consider a theory composed of $\mathcal{T} \cup \mathcal{P}$ as in Example 1. Then, e.g., the following are QN_9 -models of $\mathcal{T} \cup \mathcal{P}$:

$$\begin{aligned} \mathcal{M}_1 &= \langle U, I, I_1, I \cup \{\text{assembles}(p, c)\}, I_1 \rangle \\ \mathcal{M}_2 &= \langle U, I, I_2, I \cup \{\text{assembles}(p, c)\}, I_2 \rangle \\ \mathcal{M}_3 &= \langle U, I, \emptyset, I \cup \{\text{assembles}(p, c), \text{built}(c, \text{here})\}, \emptyset \rangle. \end{aligned}$$

where U is the Herbrand domain over $\{c, p, \text{here}\}$, and $I = \{\text{product}(c), \text{person}(p), \text{car}(c), \text{engineer}(p), \text{designs}(p, c)\}$, $I_1 = \{\text{car}(c), \text{product}(c)\}$, $I_2 = \{\text{car}(c)\}$.

Intuitively, the semantics combines the idea underlying the logic of Here-and-There, namely to restrict to frames composed of two worlds with a fixed reachability relation (reflexive and there from here, but not vice versa), with the idea of decoupling the evaluation of truth from the evaluation of falsity underlying (paraconsistent) four-valued logic. The latter is a variant of Routley's basic idea: to use separate, so-called 'starred' worlds, to validate strong negation. Hence, there is a simple bijection between Routley $\mathcal{L}$ -structures and nine-valued $\mathcal{L}$ -structures, by considering the complements (wrt. $\text{At}(D, P)$) of $I_{h,f}$ and $I_{t,f}$. Given this bijection and a semantic consequence relation for Routley $\mathcal{L}$ -structures, it can be shown that both semantics capture the same logic. Again we refer to Fink (2012) and Fink et al. (2011) for corresponding formal results.

Therefore both, Routley models and nine-valued models, characterize QN_9 . For the remainder of the paper we consider it more appropriate to work with nine-valued models. However, we remark that by their close relation, all subsequent results easily carry over to Routley semantics.

Properties of nine-valued models. The valuation of a sentence φ wrt. a nine-valued structure $\mathcal{M}$, denoted as $v(\mathcal{M}, \varphi)$, assigns one of nine truth values true $\mathbf{t}$, false $\mathbf{f}$, contradictory $\ddot{\mathbf{t}}$, undefined $\ddot{\mathbf{f}}$, believed true $\mathbf{Kt}$, believed false $\mathbf{Kf}$, believed contradictory $\mathbf{K}\ddot{\mathbf{t}}$, believed contradictory with coherent truth $\mathbf{K}\ddot{\mathbf{t}}\mathbf{t}$, or believed contradictory with coherent falsity $\mathbf{K}\ddot{\mathbf{t}}\mathbf{f}$ to φ as follows:

- $v(\mathcal{M}, \varphi) = \mathbf{t}$ iff $\mathcal{M}, w, \mathbf{t} \models \varphi$ and $\mathcal{M}, w, \mathbf{f} \not\models \varphi$;
- $v(\mathcal{M}, \varphi) = \mathbf{f}$ iff $\mathcal{M}, w, \mathbf{t} \not\models \varphi$ and $\mathcal{M}, w, \mathbf{f} \models \varphi$;
- $v(\mathcal{M}, \varphi) = \ddot{\mathbf{t}}$ iff $\mathcal{M}, w, \mathbf{t} \models \varphi$ and $\mathcal{M}, w, \mathbf{f} \models \varphi$;
- $v(\mathcal{M}, \varphi) = \ddot{\mathbf{f}}$ iff $\mathcal{M}, w, \mathbf{t} \not\models \varphi$ and $\mathcal{M}, w, \mathbf{f} \not\models \varphi$;




Figure 1: The bilattice $\mathcal{NINE}$.

- $v(\mathcal{M}, \varphi) = \mathbf{Kt}$ iff $\mathcal{M}, h, \mathbf{t} \not\models \varphi$, $\mathcal{M}, t, \mathbf{t} \models \varphi$, and $\mathcal{M}, w, \mathbf{f} \not\models \varphi$;
- $v(\mathcal{M}, \varphi) = \mathbf{Kf}$ iff $\mathcal{M}, w, \mathbf{t} \not\models \varphi$, $\mathcal{M}, h, \mathbf{f} \not\models \varphi$, and $\mathcal{M}, t, \mathbf{f} \models \varphi$;
- $v(\mathcal{M}, \varphi) = \mathbf{K}\ddot{\mathbf{t}}$ iff $\mathcal{M}, h, \mathbf{t} \not\models \varphi$, $\mathcal{M}, t, \mathbf{t} \models \varphi$, $\mathcal{M}, h, \mathbf{f} \not\models \varphi$, and $\mathcal{M}, t, \mathbf{f} \models \varphi$;
- $v(\mathcal{M}, \varphi) = \mathbf{K}\ddot{\mathbf{t}}\mathbf{t}$ iff $\mathcal{M}, w, \mathbf{t} \models \varphi$, $\mathcal{M}, h, \mathbf{f} \not\models \varphi$, and $\mathcal{M}, t, \mathbf{f} \models \varphi$;
- $v(\mathcal{M}, \varphi) = \mathbf{K}\ddot{\mathbf{t}}\mathbf{f}$ iff $\mathcal{M}, h, \mathbf{t} \not\models \varphi$, $\mathcal{M}, t, \mathbf{t} \models \varphi$, and $\mathcal{M}, w, \mathbf{f} \models \varphi$.

Note that by the above definition of model, $\mathbf{t}$, $\ddot{\mathbf{t}}$, and $\mathbf{K}\ddot{\mathbf{t}}\mathbf{t}$ are designated truth values. The nine truth values form a bilattice $\mathcal{NINE}$ wrt. two partial orders called truth order (t) and knowledge order (k), as depicted in Figure 1: for illustration, the nodes are given as binary strings of length four and labeled with a corresponding truth value. This encodes, e.g., for a ground atom a its valuation by a QN_9 -structure where a is present (1) or not present (0) in $I_{h,t}$, $I_{h,f}$, $I_{t,t}$, and $I_{t,f}$, respectively.

Example 3 Given the models $\mathcal{M}_1$, $\mathcal{M}_2$, and $\mathcal{M}_3$ as in our previous Example 2 above, the following valuations are obtained: $v(\mathcal{M}_1, \text{car}(c)) = v(\mathcal{M}_1, \text{product}(c)) = \ddot{\mathbf{t}}$ (contradictory), while $v(\mathcal{M}_2, \text{product}(c)) = v(\mathcal{M}_3, \text{product}(c)) = \mathbf{t}$ and $v(\mathcal{M}_3, \text{car}(c)) = \mathbf{t}$. All three models assign $\mathbf{Kt}$ (believed true) to $\text{assembles}(p, c)$ and make $\text{designs}(p, c)$ true ($\mathbf{t}$).

A nine-valued structure $\mathcal{M} = \langle (D, \sigma), I_{h,t}, I_{h,f}, I_{t,t}, I_{t,f} \rangle$ is called *total* iff $I_{h,t} = I_{t,t}$ and $I_{h,f} = I_{t,f}$; it is *consistent* iff $I_{w,t} \cap I_{w,f} = \emptyset$ for $w \in \{h, t\}$. Intuitively, wrt. total nine-valued models the implication connective behaves classically. Consequently, total nine-valued models of a sentence (and thus also of theories) coincide with its four-valued models. This correspondence, for instance, is more obvious given the definition of four-valued semantics we consider (because it has also been used to study paraconsistency in Description Logics) and the alternative semantic characterization of QN_9 compared to Routley semantics.

Proposition 1 *Given a $\mathcal{L}$ -theory Γ , a universe $U = (D, \sigma)$, and two $\mathcal{L}$ -interpretations I_t, I_f over D , the following holds: $\langle U, I_t, I_f \rangle \models_4 \Gamma$ iff $\langle U, I_t, I_f, I_t, I_f \rangle \models_{\text{QN}_9} \Gamma$.*

We also remark, that by the above correspondence one can introduce four-valued logic by reference to QN_9 and corresponding structures. However, we opted to stick with a presentation that is closer to its treatment in work on paraconsistent semantics for Description Logics.

Moreover, total nine-valued models which also are consistent characterize classical models of the underlying theory. More specifically, the following property holds.

Proposition 2 *Let α be a $\mathcal{L}$ -sentence, and let $\mathcal{M} = \langle (D, \sigma), I_t, I_f, I_t, I_f \rangle$ be a total and consistent nine-valued $\mathcal{L}$ -structure, i.e., such that $I_t \cap I_f = \emptyset$, then $\mathcal{M}, w, t \models \sim \varphi$ implies $\mathcal{M}, w, t \not\models \varphi$. The converse holds for $I_f = \text{At}(D, P) \setminus I_t$.*

Note that (due to the totality of $\mathcal{M}$) this result intuitively relates strong negation with a classical interpretation (I_t) of the theory (where in addition to implication—and thus default negation—also ' $\sim$ ' is interpreted classically). Consequently, with such QN_9 models (i.e., total and consistent ones) we can associate classical models of the theory and vice versa, as follows.

Corollary 1 *Consider a $\mathcal{L}$ -theory Γ , a universe $U = (D, \sigma)$, and two $\mathcal{L}$ -interpretations I_t, I_f over D , such that $I_t \cap I_f = \emptyset$. Then, $\langle U, I_t, I_f, I_t, I_f \rangle \models_{\text{QN}_9} \Gamma$ implies $\langle U, I_t \rangle \models \Gamma$. The converse holds for $I_f = \text{At}(D, P) \setminus I_t$.*

Let us eventually remark that by restricting to consistent—but not necessarily total—structures, one obtains nine-valued models that correspond to models of a static quantified version of Nelson logic (with decidable equality), cf. QN_9^s in Pearce and Valverde (2005). The latter serves as a logical foundation for answer-set semantics when strong negation is explicitly present in the language. Accordingly, we finally define equilibrium models of a $\mathcal{L}$ -theory in terms of consistent nine-valued models as follows.

Given two pairs of sets, (X_1, Y_1) and (X_2, Y_2) , we say that (X_1, Y_1) is a subset of (X_2, Y_2) , in symbols $(X_1, Y_1) \subseteq (X_2, Y_2)$, iff $X_1 \subseteq X_2$ and $Y_1 \subseteq Y_2$; it is a strict subset iff at least one of the two inclusions is strict.

Definition 5 *Given a $\mathcal{L}$ -theory Γ , a total and consistent nine-valued $\mathcal{L}$ -structure $\mathcal{M} = \langle (D, \sigma), I_t, I_f, I_t, I_f \rangle$, i.e., such that $I_t \cap I_f = \emptyset$, is called an equilibrium model of Γ iff $\mathcal{M} \models \Gamma$ and $\mathcal{M}' \not\models \Gamma$, for all $\mathcal{M}' = \langle (D, \sigma), I'_t, I'_f, I_t, I_f \rangle$, such that $(I'_t, I'_f) \subset (I_t, I_f)$. The set of all equilibrium models of Γ is denoted by $\mathcal{EQ}(\Gamma)$.*

Paraconsistent Semi-Equilibrium Semantics

We are now ready to develop a paraconsistent semantics for hybrid knowledge bases with QN_9 as its underlying logic. The reason why QN_9 is particularly suited for this purpose, will become more clear when we investigate properties of the semantics defined. Intuitively speaking, it allows us to extend existing approaches quite straight forwardly to the hybrid setting, and our corresponding design objectives in terms of desirable properties can be achieved.

Paraconsistent Semi-Equilibrium Models

By restricting to nine-valued models that satisfy certain minimality criteria, it is possible to generalize corresponding equilibrium semantics quite directly to a paraconsistent version. Since our goal is to study paraconsistent hybrid theories, specifically combinations of nonmonotonic rules with classical theories, our interest is in a paraconsistent and quantified version of so-called *semi-equilibrium semantics* for logic programs (Eiter, Fink, and Moura 2010). The latter has been developed as a paracoherent semantics, i.e., to deal with incoherence due to weak negation, for (propositional) nonmonotonic logic programs without strong negation (hence not concerned with issues of paraconsistency, i.e., inconsistency wrt. strong negation). Generalizing to first-order theories under answer-set semantics (but still without strong negation) the main characterization in Eiter, Fink, and Moura (2010) provides a definition of semi-equilibrium models as follows. Suppose that Γ is a $\mathcal{L}$ -theory not involving ' $\sim$ ', and assume that $\mathcal{M} = \langle U, H, T \rangle$ is a QHT $\mathcal{L}$ -model of Γ . Then, $\mathcal{M}$ is called a *semi-equilibrium model* of Γ iff (a) $\langle U, H', T \rangle \not\models \Gamma$, for all $H' \subset H$, and (b) there is no QHT $\mathcal{L}$ -model $\langle U, H', T' \rangle$ of Γ that satisfies (a) and $T' \setminus H' \subset T \setminus H$. The set of all semi-equilibrium models of Γ is denoted by $\mathcal{SEQ}(\Gamma)$.

The main intuition of this semantics is to weaken equilibrium-model (thus answer-set) semantics, by giving up on the equilibrium condition, allowing for a 'gap' between H and T , i.e. for certain atoms to be unfounded. However, atoms not in the gap still need to be founded (by Condition (a)) and the gap should be globally (subset-)minimal (cf. Condition (b)). Thus, one aims at staying as close as possible to equilibrium semantics.

The following definition generalizes this notion to arbitrary first-order theories with strong negation that is dealt with paraconsistently (generalizing so-called PAS semantics (Alcântara, Damásio, and Pereira 2004) in this respect).

Definition 6 (pseq-models of a theory) *Let Γ be a $\mathcal{L}$ -theory and let $\mathcal{M} = \langle U, H_t, H_f, T_t, T_f \rangle$ be a nine-valued $\mathcal{L}$ -model of Γ . Then, $\mathcal{M}$ is a paraconsistent semi-equilibrium model (pseq-model) of Γ iff*

- (i) $\langle U, H'_t, H'_f, T_t, T_f \rangle \not\models \Gamma$, for all $(H'_t, H'_f) \subset (H_t, H_f)$, and
- (ii) *there is no nine-valued $\mathcal{L}$ -model $\langle U, H'_t, H'_f, T'_t, T'_f \rangle$ of Γ satisfying (i) and $(T'_t \setminus H'_t, T'_f \setminus H'_f) \subset (T_t \setminus H_t, T_f \setminus H_f)$.*

The set of all pseq-models of Γ is denoted by $\text{pSEQ}(\Gamma)$.

For our purpose this semantic definition yields the basis to deal with hybrid theories in a paracoherent and paraconsistent way. That this can essentially be achieved by considering Routley structures in place of here-and-there interpretations has already been put forward in Eiter, Fink, and Moura (2010), and we leave a more detailed discussion of its relation to PAS and other semantics (in the propositional case), as well as the issue of alternative characterizations in terms of program transformations, to an extended version thereof.

Paraconsistent Hybrid Theories

Inspired by hybrid knowledge bases that combine classical theories with nonmonotonic rules, we first slightly generalize this idea to so-called hybrid theories. Further on, we will obtain and study a paraconsistent semantics for hybrid theories that comprises hybrid knowledge bases as special (restricted) settings.

Definition 7 (hybrid theory) A hybrid $\mathcal{L}$ -theory is a theory over a language $\mathcal{L} = \langle C, P \rangle$, where a subset $P_{\mathcal{T}}$ of the predicates P is designated as classical (also called structural) predicates, i.e., $\mathcal{L} = \langle C, P_{\mathcal{T}} \cup P_{\mathcal{P}} \rangle$ such that $P_{\mathcal{T}} \cap P_{\mathcal{P}} = \emptyset$.

In analogy to hybrid knowledge bases, the restriction of the language to structural predicates $\mathcal{L}_{\mathcal{T}} = \langle C, P_{\mathcal{T}} \rangle$ is called the structural language of $\mathcal{L}$. Moreover, a hybrid theory Γ can be regarded as being composed of a ‘classical’ (four-valued) first-order theory $\mathcal{T}$, the *structural part* of Γ , i.e., sentences entirely over $\mathcal{L}_{\mathcal{T}}$, and a set of $\mathcal{L}$ -sentences $\mathcal{P} = \Gamma \setminus \mathcal{T}$. We keep referring to $\mathcal{P}$ as the *rules part* of Γ , however, without imposing any syntactic restrictions to rules of a particular form. In the following, we thus denote hybrid $\mathcal{L}$ -theories Γ as pairs $(\mathcal{T}, \mathcal{P})$ in order to avoid confusion with (ordinary) $\mathcal{L}$ -theories. Moreover, given a set of atomic $\mathcal{L}$ -sentences S , we write $S|_{\mathcal{L}_{\mathcal{T}}}$, respectively $S|_{\mathcal{L}_{\mathcal{P}}}$, to denote its restriction to atomic sentences over $P_{\mathcal{T}}$ or $P_{\mathcal{P}}$, respectively.

Example 4 Our running example, i.e., $(\mathcal{T}, \mathcal{P})$ of Example 1, is a hybrid theory over $\mathcal{L} = \langle C, P_{\mathcal{T}} \cup P_{\mathcal{P}} \rangle$, where $\{c, p, \text{here}\} \subseteq C$, $P_{\mathcal{T}} = \{\text{car}, \text{designs}, \text{engineer}, \text{person}, \text{product}\}$, and $P_{\mathcal{P}} = \{\text{assembles}, \text{built}\}$.

Obviously, just declaring certain predicates structural does not make them behave classically without a corresponding semantic treatment. For hybrid knowledge bases, the stable closure of the structural part yields the desired classical interpretations. However, as we already pointed out before, the stable closure prohibits a paraconsistent semantics on the basis of QHT: if the structural part is inconsistent, then the stable closure has no QHT model at all. In our envisaged paraconsistent semantics on the basis of QN_9 , we aim at dealing with such situations by four-valued interpretations. Thus, we require a paraconsistent version of the stable closure that, intuitively speaking, forces structural predicates to behave ‘classically’ in a four-valued sense.

Definition 8 (pstable closure) Given a hybrid $\mathcal{L}$ -theory $(\mathcal{T}, \mathcal{P})$, its paraconsistent stable-closure (*pstable closure*) is $\mathcal{T} \cup \text{pst}(\mathcal{T}) \cup \mathcal{P}$, where $\text{pst}(\mathcal{T})$ is the set of sentences: $\{\forall x((p(x) \vee \neg p(x)) \wedge (\sim p(x) \vee \neg \sim p(x))) \mid p \in \mathcal{L}_{\mathcal{T}}\}$.

Note that the sentences in $\text{pst}(\mathcal{T})$ are structural since they involve structural predicates only. The occurrence of intuitionistic negation (or internal implication since, e.g., $\neg p(x)$) is just an abbreviation for $p(x) \rightarrow \perp$ may tempt one to regard them as non-classical. However, it is precisely the effect of the pstable closure that these connectives are interpreted ‘classically’ (in a paraconsistent four-valued sense) over the structural language (as opposed to their ‘non-classical’ nine-valued interpretation in QN_9). This is captured more formally by the following lemma.

Lemma 1 Let $(\mathcal{T}, \mathcal{P})$ be a hybrid $\mathcal{L}$ -theory, then $\text{pst}(\mathcal{T}) \models_{\text{QN}_9} \neg \neg \varphi \rightarrow \varphi$ holds for every $\mathcal{L}_{\mathcal{T}}$ -sentence φ .

Eventually, by means of the pstable closure we are in the position to define a paracoherent and paraconsistent semantics for hybrid theories.

Definition 9 (pseq-models of a hybrid theory) The *pseq-models* of a hybrid theory $(\mathcal{T}, \mathcal{P})$ are the *pseq-models* of its pstable closure, in symbols $\text{pSEQ}((\mathcal{T}, \mathcal{P})) = \text{pSEQ}(\mathcal{T} \cup \text{pst}(\mathcal{T}) \cup \mathcal{P})$.

That the semantics thus defined provably captures our intuitions and behaves as intended is the subject of the next section. Before turning to these more formal issues, let us briefly reconsider our running example.

Example 5 Given the hybrid theory $(\mathcal{T}, \mathcal{P})$ of Example 1, $\text{pst}(\mathcal{T})$ adds respective formulas for predicates *car*, *designs*, *engineer*, *person*, and *product*. Thus, they are essentially interpreted under four-valued logic. Applying pseq-semantics, observe that $\mathcal{M}_1$ and $\mathcal{M}_2$ are in $\text{pSEQ}(\mathcal{T} \cup \text{pst}(\mathcal{T}) \cup \mathcal{P})$, while $\mathcal{M}_3$ is a consistent model but not a pseq-model of the pstable closure. Indeed, $\mathcal{M}_1$ yields a smaller gap between the *here* worlds and the *there* worlds (see Condition (ii) in Definition 6).

Also note, how $\mathcal{M}_1$ and $\mathcal{M}_2$ cope with the inconsistency wrt. *c* being a car, and the incoherence concerning its assembly by *p*, by assigning $\ddagger$ (contradictory) to $\text{car}(c)$ and $\mathbf{Kt}$ (believed true) to $\text{assembles}(p, c)$ (cf. Example 3). Moreover, they allow for $\text{designs}(p, c)$ as a non-trivial conclusion that is $\mathbf{t}$ (true).

Properties

As for desirable properties, one would expect a particular behavior for two special cases of hybrid theories, namely when the entire theory is either a logic program, or a classical theory, i.e., when either $\mathcal{T} = \emptyset$, or $\mathcal{P} = \emptyset$, respectively. Recall, that in the former case, our goal was to generalize paracoherent ASP semantics, while in the latter case we aimed at generalizing (paraconsistent) four-valued logic. The following results verify these objectives, starting with the case $\mathcal{P} = \emptyset$, where the corresponding result is a consequence of the previous Lemma 1.

Theorem 1 Let $\mathcal{L} = \langle C, P_{\mathcal{T}} \cup P_{\mathcal{P}} \rangle$, and consider a hybrid $\mathcal{L}$ -theory $(\mathcal{T}, \emptyset)$. Then, $\langle (D, \sigma), I_{h,t}, I_{h,f}, I_{t,t}, I_{t,f} \rangle \in \text{pSEQ}((\mathcal{T}, \emptyset))$ iff $I_{h,t} = I_{t,t}$, $I_{h,f} = I_{t,f}$, $I_{t,t}|_{\mathcal{L}_{\mathcal{P}}} = I_{t,f}|_{\mathcal{L}_{\mathcal{P}}} = \emptyset$, and $\langle (D, \sigma), I_{t,t}, I_{t,f} \rangle \models_4 \mathcal{T}$.

In order to verify our second objective, we need to resort to paracoherent ASP semantics. Due to the fact that in Eiter, Fink, and Moura (2010) only propositional programs without strong negation have been considered, we consider here the same restriction for strong negation but general first-order theories. More precisely, we resort to semi-equilibrium models for $\mathcal{L}$ -theories as introduced in the beginning of this section. We remark, however, that the subsequent result also generalizes to programs (and general first-order theories) with strong negation, given a corresponding extension of paracoherent ASP semantics as sketched in Eiter, Fink, and Moura (2010).

Theorem 2 Let $\mathcal{L} = \langle C, P_{\mathcal{T}} \cup P_{\mathcal{P}} \rangle$, such that $P_{\mathcal{T}} = \emptyset$, and consider a hybrid $\mathcal{L}$ -theory $(\emptyset, \mathcal{P})$, where $\mathcal{P}$ does not involve strong negation. Then, $\langle (D, \sigma), I_{h,t}, I_{h,f}, I_{t,t}, I_{t,f} \rangle \in p\mathcal{SEQ}((\emptyset, \mathcal{P}))$ iff $I_{h,f} = \emptyset$, $I_{t,f} = \emptyset$, and $\langle (D, \sigma), I_{h,t}, I_{t,t} \rangle \in \mathcal{SEQ}(\mathcal{P})$.

Note that rather than $\mathcal{T} = \emptyset$, the stronger requirement that $P_{\mathcal{T}} = \emptyset$ is necessary for the one-to-one correspondence of the above theorem. If the structural language is non-empty, then even if $\mathcal{T} = \emptyset$ and the program part does not involve predicates from $\mathcal{L}_{\mathcal{T}}$, the relationship becomes many to one: structural predicates can be interpreted arbitrarily. More formally, the following result holds.

Proposition 3 Let $\mathcal{L} = \langle C, P_{\mathcal{T}} \cup P_{\mathcal{P}} \rangle$, and consider a hybrid $\mathcal{L}$ -theory $(\emptyset, \mathcal{P})$, where $\mathcal{P}$ does neither involve strong negation nor predicates over $P_{\mathcal{T}}$. Then,

- (i) $\langle (D, \sigma), I_{h,t} \cup I_1, I_2, I_{t,t} \cup I_1, I_2 \rangle \in p\mathcal{SEQ}((\emptyset, \mathcal{P}))$ if $I_1 \subseteq \text{At}(D, P_{\mathcal{T}})$, $I_2 \subseteq \text{At}(D, P_{\mathcal{T}})$, and $\langle (D, \sigma), I_{h,t}, I_{t,t} \rangle \in \mathcal{SEQ}(\mathcal{P})$;
- (ii) $\langle (D, \sigma), I_{h,t}|_{\mathcal{L}_{\mathcal{P}}}, I_{t,t}|_{\mathcal{L}_{\mathcal{P}}} \rangle \in \mathcal{SEQ}(\mathcal{P})$ if $\langle (D, \sigma), I_{h,t}, I_{h,f}, I_{t,t}, I_{t,f} \rangle \in p\mathcal{SEQ}((\emptyset, \mathcal{P}))$.

Resorting to paracoherent ASP semantics for the rules part, when defining a paraconsistent semantics for hybrid theories, was motivated by benign properties of semi-equilibrium model semantics. In particular, *classical coherence*, *answer-set congruence*, and *answer-set coverage* are desirable properties (Eiter, Fink, and Moura 2010), also in the context of hybrid theories. Therefore, we next study whether our generalization to hybrid theories by means of pseq-models retained these properties.

Let us start by restating classical coherence, answer-set congruence, and answer-set coverage in the setting of hybrid theories. Consider a hybrid theory $(\mathcal{T}, \mathcal{P})$, let Γ denote its pstable closure, and recall that Γ can also be treated classically (since default negation has been introduced as an abbreviation for a particular implication, which can also be interpreted classically). Then, the properties of interest are:

classical coherence If Γ has a classical model, then it has a pseq-model;

coverage Every equilibrium model of Γ corresponds to a pseq-model;

congruence If an equilibrium model exists for Γ , then all pseq-models correspond to equilibrium models.

The next result establishes classical coherence for paraconsistent semi-equilibrium semantics.

Proposition 4 Let $\mathcal{L} = \langle C, P_{\mathcal{T}} \cup P_{\mathcal{P}} \rangle$, consider a hybrid $\mathcal{L}$ -theory $(\mathcal{T}, \mathcal{P})$, and let Γ be its pstable closure $\mathcal{T} \cup \text{pst}(\mathcal{T}) \cup \mathcal{P}$. Then, $p\mathcal{SEQ}((\mathcal{T}, \mathcal{P})) \neq \emptyset$ if $\langle U, I \rangle \models \Gamma$, for a universe $U = (D, \sigma)$ and some $\mathcal{L}$ -interpretation I .

Also, coverage holds for pseq-models.

Proposition 5 Let $\mathcal{L} = \langle C, P_{\mathcal{T}} \cup P_{\mathcal{P}} \rangle$, consider a hybrid $\mathcal{L}$ -theory $(\mathcal{T}, \mathcal{P})$, and let Γ be its pstable closure $\mathcal{T} \cup \text{pst}(\mathcal{T}) \cup \mathcal{P}$. Then, $\mathcal{M} \in \mathcal{EQ}(\Gamma)$ implies $\mathcal{M} \in p\mathcal{SEQ}((\mathcal{T}, \mathcal{P}))$.

However, the converse, i.e., congruence, does not hold (although all pseq-models are total if $\mathcal{EQ}(\Gamma) \neq \emptyset$). Intuitively, this is a consequence of adopting the rather weak

four-valued semantics for the paraconsistent treatment of the structural part of hybrid theories: pseq-models need not be consistent, whereas equilibrium models are consistent by definition. Of course, restricting to consistent pseq-models would degrade the ability to deal with hybrid theories paraconsistently. But in analogy to stronger paraconsistent semantics for classical theories, for instance like in Priest's Logic, one may think of imposing further minimization criteria that intuitively aim at giving preference to consistent interpretations. This motivates the following definition.

Definition 10 (preferred pseq-models of a theory) Let Γ be a $\mathcal{L}$ -theory and let $\mathcal{M} = \langle U, H_t, H_f, T_t, T_f \rangle$ be a pseq-model of Γ . Then, $\mathcal{M}$ is a preferred paraconsistent semi-equilibrium model (preferred pseq-model) of Γ iff

- (iii) there is no pseq-model $\langle U, H'_t, H'_f, T'_t, T'_f \rangle$ of Γ , such that $T'_t \cap T'_f \subset T_t \cap T_f$.

As before, preferred pseq-models of a hybrid theory are defined as the preferred pseq-models of its pstable closure.

Example 6 Reconsider our running example, i.e., $(\mathcal{T}, \mathcal{P})$ as given in Example 1 and the pseq-models $\mathcal{M}_1$ and $\mathcal{M}_2$ of $\Gamma = \mathcal{T} \cup \text{pst}(\mathcal{T}) \cup \mathcal{P}$. Observe that $\mathcal{M}_2$ is preferred over $\mathcal{M}_1$ in terms of Condition (iii) in Definition 10 above, i.e., in terms of minimizing the contradictory interpretation of structural ground atoms. The intuitive explanation is that, contrary to *car(c)*, there is no reason for interpreting also *product(c)* as contradictory. Actually, $\mathcal{M}_2$ is a preferred pseq-model of $(\mathcal{T}, \mathcal{P})$.

As it can be easily verified, classical coherence and coverage still hold on preferred pseq-models, which in addition also satisfy congruence:

Proposition 6 Let $\mathcal{L} = \langle C, P_{\mathcal{T}} \cup P_{\mathcal{P}} \rangle$, consider a hybrid $\mathcal{L}$ -theory $(\mathcal{T}, \mathcal{P})$, and let Γ be its pstable closure $\mathcal{T} \cup \text{pst}(\mathcal{T}) \cup \mathcal{P}$. Then, $\mathcal{EQ}(\Gamma) \neq \emptyset$ implies $\mathcal{M} \in \mathcal{EQ}(\Gamma)$, for every preferred pseq-model $\mathcal{M}$ in $p\mathcal{SEQ}((\mathcal{T}, \mathcal{P}))$.

Paraconsistent Hybrid Knowledge Bases. These are special cases of hybrid theories, where the rules part is syntactically restricted to rules of the usual form, i.e., universally quantified implications, where the antecedent is a conjunction of literals or default negated literals, and the consequent is a disjunction of literals. Sometimes also the structural part of a hybrid knowledge base is implicitly given. In particular, when it is composed of a Description Logic knowledge base, i.e., by a DL TBox and ABox, of a Description Logic that can be translated into first-order logic.

Independently of whether the structural part is explicitly represented or not, a paraconsistent treatment may in practice require some preprocessing (translation) of the knowledge bases, both the structural part as well as the rule base. The domain expert should thus be aware, that the outcome of these translations, respectively otherwise semantically equivalent variants of representing knowledge bases, may yield semantic differences in the paraconsistent setting. For the rules part, corresponding considerations may concern the use of strong negation instead of default negation (if the original setting was one without strong negation)

which, however, is a well-understood knowledge representation aspect in answer-set programming (Gelfond and Lifschitz 1991). For the structural part, the transition from classical to four-valued semantics allows for different forms of implication: in addition to internal implication, material and strong implication are definable giving rise to different translations. We refer the reader to Ma, Hitzler, and Lin (2007; 2008), where this knowledge representation aspect is considered for DL knowledge bases under paraconsistent four-valued semantics and thus equally applies to pseq-semantics (likewise for aspects concerning the usage of $\perp$ and omniconsistency, i.e., guaranteed four-valued model existence).

Computational Complexity

We next consider reasoning tasks on hybrid theories under paraconsistent semi-equilibrium semantics and analyze their computational complexity for practically relevant (decidable) fragments of the language. In general, the problems are undecidable, which is an obvious consequence of well-known undecidability results, e.g., for first-order logic. We therefore restrict our attention to practically relevant fragments of hybrid theories that allow for decidable reasoning.

Assumptions. Throughout this section, we assume that $\mathcal{P}$ consists of universally quantified sentences of the form

$$l_1 \wedge \dots \wedge l_m \wedge \neg l_{m+1} \dots \neg l_n \rightarrow h_1 \vee \dots \vee h_l$$

where $l_1, \dots, l_b$ and $h_1, \dots, h_l$ are literals (atoms or strongly negated atoms) of $\mathcal{L}_{\mathcal{P}}$ called body literals and head literals, respectively. Furthermore, we say that $\{l_1, \dots, l_m\}$ is the positive body of a rule, $\{l_{m+1}, \dots, l_n\}$ is its negative body, and head literals are required to be equality-free.

In addition, we require rules to satisfy a syntactic safety restriction that is obtained by adopting *weak DL-safeness* (Rosati 2008) to our setting of hybrid theories. The rules part $\mathcal{P}$ is called *weakly $\mathcal{T}$ -safe* iff every rule $r \in \mathcal{P}$ of the above form satisfies the following condition: for every variable x in r , x either appears in a positive body literal over $P_{\mathcal{P}}$, or x only occurs in positive literals over $P_{\mathcal{T}}$. In the latter case, we say that x is a $\mathcal{T}$ -variable in r .

Intuitively, these restrictions guarantee that pseq-models of a weakly $\mathcal{T}$ -safe hybrid theory $(\mathcal{T}, \mathcal{P})$ can be considered as being composed of two parts. Thereby, one part, call it the $\mathcal{P}$ -part $\mathcal{M}_{\mathcal{P}}$ of a model $\mathcal{M}$, is always finite. It is given by interpretations of $\mathcal{L}'_{\mathcal{P}} = \langle D', P'_{\mathcal{T}} \cup P_{\mathcal{P}} \rangle$, where $P'_{\mathcal{T}}$ is the restriction of $P_{\mathcal{T}}$ to predicates occurring in $\mathcal{P}$, over a finite domain D' whose size only depends on $\mathcal{P}$. Moreover, the rules part of any pseq-model is domain expansion safe, in the sense that the interpretation of predicates from $P_{\mathcal{P}}$ does not change when considering (possibly infinite) domains $D \supseteq D'$. Also, Conditions (i) and (ii) of Definition 6 only effect $\mathcal{M}_{\mathcal{P}}$ and can be established by grounding $\mathcal{P}$ over D' .

Nevertheless, the structural part of a weakly $\mathcal{T}$ -safe hybrid theory may still cause associated reasoning problems to be undecidable (even for literals over $P_{\mathcal{P}}$), intuitively because the corresponding interpretation of predicates from $P'_{\mathcal{T}}$ in the rules part of a pseq-model must be compliant with (or extendable to) a four-valued model of $\mathcal{T}$. There-

Problem	normal $\mathcal{P}$	disjunctive $\mathcal{P}$
CONSISTENCY	NEXP/NP	NEXP/NP
MODELCHECK	coNEXP/coNP	coNEXP ^{NP} / Π_2^P
BRAVECONS	NEXP ^{NP} / Σ_2^P	NEXP ^{Σ_2^P} / Σ_3^P
CAUTIOUSCONS	coNEXP ^{NP} / Π_2^P	coNEXP ^{Σ_2^P} / Π_3^P

Table 1: Combined/Data complexity of reasoning tasks for hybrid theories under pseq-semantics.

fore, we restrict our attention to structural parts $\mathcal{T}$, such that the arity of predicates in $P_{\mathcal{T}}$ is bounded by a constant and four-valued existential entailment is decidable. More specifically, assume that $\mathcal{T} \models_4 \varphi_{Bool}$ is decidable for any ground sentence (boolean combination of ground literals) φ_{Bool} . We consider the case where the (combined) complexity of the corresponding decision problem is in EXP, and its data complexity is in P, respectively. Note that these assumptions seem to be reasonable as far as combinations of rules with DL knowledge bases are concerned. To wit, we refer to results in Ma, Hitzler, and Lin (2008) that essentially establish the above for various relevant Description Logics.

Problems and Results. In particular we are interested in the following problems of paraconsistent reasoning:

CONSISTENCY deciding whether a hybrid theory $(\mathcal{T}, \mathcal{P})$ has a pseq-model (i.e., whether $pSEQ((\mathcal{T}, \mathcal{P})) \neq \emptyset$);

MODELCHECK recognizing the (finite) $\mathcal{P}$ -part of pseq-models, given $\mathcal{M}_{\mathcal{P}}$ and $(\mathcal{T}, \mathcal{P})$;

BRAVECONS brave reasoning; and

CAUTIOUSCONS cautious reasoning from the pseq-models of a hybrid theory $(\mathcal{T}, \mathcal{P})$.

The task in Problem MODELCHECK is to decide whether a given finite nine-valued structure $\mathcal{M}_{\mathcal{P}}$ can be extended to a pseq-model $E(\mathcal{M}_{\mathcal{P}})$ of $(\mathcal{T}, \mathcal{P})$, i.e., whether $E(\mathcal{M}_{\mathcal{P}}) \in pSEQ((\mathcal{T}, \mathcal{P}))$. Concerning brave and cautious reasoning, more specifically, we consider the problem of deciding whether a ground literal l , i.e., a ground atom a or a strongly negated ground atom $\sim a$, is a brave (respectively, cautious) consequence of a hybrid theory $(\mathcal{T}, \mathcal{P})$ with a particular truth value $e \in \{\mathbf{t}, \mathbf{f}, \mathbf{\ddot{t}}, \mathbf{\ddot{f}}, \mathbf{Kt}, \mathbf{Kf}, \mathbf{K\ddot{t}}, \mathbf{K\ddot{f}}, \mathbf{K\ddot{t}t}, \mathbf{K\ddot{f}f}\}$. We use $(\mathcal{T}, \mathcal{P}) \models_b^e l$, respectively $(\mathcal{T}, \mathcal{P}) \models_c^e l$, to denote these problems, and say that l is a brave (resp., cautious) consequence of $(\mathcal{T}, \mathcal{P})$ iff $v(\mathcal{M}, l) = e$ for some (resp., every) $\mathcal{M} \in pSEQ((\mathcal{T}, \mathcal{P}))$.

Given the above assumptions, our results concerning combined complexity and data complexity of the reasoning tasks considered are summarized in Table 1. All entries in the table refer to corresponding completeness results. Unsurprisingly, the results for combined complexity exhibit an exponential blow-up compared to corresponding results for semi-equilibrium models of propositional logic programs. Intuitively, this is explained by the fact that the main source of complexity stems from establishing Conditions (i) and (ii) of Definition 6, which is confined to the rules part of $(\mathcal{T}, \mathcal{P})$.

In doing so, since we deal with non-ground programs, exponentially larger interpretations have to be considered due to grounding. When considering data complexity, i.e., $\mathcal{T}$ and $\mathcal{P}$ are fixed modulo ground literals (facts), then also the arity of predicates is fixed, resulting in exponentially lower complexity. In comparison with (two-valued) reasoning on equilibrium models observe that brave and cautious reasoning are one level up in the exponential hierarchy. Deciding pseq-model existence has the same complexity as equilibrium model existence if the rules part is normal (non-disjunctive), and does not increase for disjunctive $\mathcal{P}$ (as apposed to equilibrium model existence, which increases to the second level in this case). We eventually remark that if $\mathcal{T} \models_{\mathcal{A}} \varphi_{Bool}$ is in coNEXP (resp., data complexity in coNP), then the results for disjunctive rule parts do not change, while the upper bounds for normal $\mathcal{P}$ cease to hold. However, we postpone a more detailed analysis of this case. Likewise, due to space constraints, we refer to Fink (2012) and Fink et al. (2011) for a more formal account of the complexity results reported, and conclude this section with a succinct statement of these results.

Theorem 3 *Given a weakly $\mathcal{T}$ -safe hybrid theory $(\mathcal{T}, \mathcal{P})$, combined and data complexity for CONSISTENCY, MODELCHECK, BRAVECONS, and CAUTIOUSCONS is complete for the classes given in Table 1.*

Related Work and Conclusion

We have addressed an open problem for combinations of rules and ontologies under ASP semantics, how to deal with inconsistency, by developing a paraconsistent and paracoherent semantics for such hybrid theories. The semantics generalizes previous approaches for the individual parts, in particular paraconsistent semantics for DLs (Ma, Hitzler, and Lin 2007) and paracoherent semantics for ASP programs (Eiter, Fink, and Moura 2010) such that benign properties carry over. A study of semantic properties established these objectives and has been complemented by a complexity analysis of important corresponding reasoning tasks.

Besides the works generalized, most closely related to our work is (Huang, Li, and Hitzler 2011) establishing a paraconsistent semantics for MKNF knowledge bases. Based on a nonmonotonic modal logic (as opposed to QHT), it addresses inconsistency tolerance for tight couplings with different syntax and semantics (and sets of interpretations as the main semantic structures). Disregarding that in general we consider a broader class of hybrid theories (concerning syntactical restrictions and domain assumptions) and confining to a common fragment that is syntactically translatable, the approaches still differ semantically. In general, paraconsistent MKNF semantics is weaker than pseq-semantics, which is also reflected in a lower worstcase complexity. The main reason is that, it does not take incoherence into account, while like pseq-semantics it aims at resolving contradiction wrt. conflicting literals. For instance, our motivating example expressed as a hybrid MKNF theory would not have a paraconsistent MKNF model. For coherent (but not necessarily consistent) theories however, both semantics coincide on cautious consequences, in particular for

$\mathcal{P} = \emptyset$ since both generalize classical four-valued logic, and for $\mathcal{T} = \emptyset$ (and coherent $\mathcal{P}$) because both generalize PAS semantics. In this sense, pseq-semantics may be considered a generalization of paraconsistent MKNF semantics as well, tackling incoherence in addition to inconsistency (suitably reflected in the semantic structures, i.e., without mixing these concepts).

Another form of tolerating inconsistency in combinations has been considered in Lembo et al. (2011), and Rosati (2011). Taking a database perspective, inconsistency there arises wrt. a dedicated set of formulas representing integrity constraints. Tolerance, i.e., consistent query answering, is achieved by considering variants of the extensional data (the ABox, resp. logic program facts) as possible repairs and reasoning over them.

Balduccini and Gelfond (2003) have developed consistency-restoring rules (CR-Prolog) as a means to tolerate inconsistency in ASP programs using abductive logic programming. Since it leaves the task of writing appropriate cr-rules and specifying preferences to the programmer, CR-Prolog provides a flexible framework to realize inconsistency tolerant semantics. However, it cannot be applied to realize pseq-semantics for hybrid theories in a principled way, in other words, pseq-semantics cannot be characterized in terms of CR-Prolog. Intuitively, the main reasons are that PAS cannot be captured faithfully and that the minimization of abducibles is in contrast with the weaker four-valued treatment of structural literals.

Ongoing and future work comprises an extension of the concept of weak $\mathcal{T}$ -safety for hybrid theories to sentences of a more liberal form by a suitable adaption of the concept of safety in Cabalar, Pearce, and Valverde (2009), and Lee, Lifschitz, and Palla (2008). A further interesting theoretical issue is to investigate whether and how more general properties of paraconsistent extensions of monotonic logics, studied, e.g., by Arieli, Avron, and Zamansky (2011b), carry over to nonmonotonic settings. Of course, developing algorithms for implementing paraconsistent reasoners by extending existing solvers, for instance SAT-based ASP solvers or genuine reasoners for hybrid knowledge bases like Ontobroker, is another next step on our agenda.

References

- Alcântara, J.; Damásio, C. V.; and Pereira, L. M. 2004. A declarative characterization of disjunctive paraconsistent answer sets. In de Mántaras, R. L., and Saitta, L., eds., *ECAI*, 951–952. IOS Press.
- Almukdad, A., and Nelson, D. 1984. Constructible falsity and inexact predicates. *J. Symb. Logic* 49(1):231–233.
- Arieli, O.; Avron, A.; and Zamansky, A. 2011a. Ideal paraconsistent logics. *Studia Logica* 99(1-3):31–60.
- Arieli, O.; Avron, A.; and Zamansky, A. 2011b. What is an ideal logic for reasoning with inconsistency? In Walsh, T., ed., *IJCAI*, 706–711. AAAI Press.
- Balduccini, M., and Gelfond, M. 2003. Logic programs with consistency-restoring rules. In Doherty, P.; McCarthy, J.; and Williams, M.-A., eds., *International Symposium on*

- Logical Formalization of Commonsense Reasoning*, AAAI 2003 Spring Symposium Series, 9–18.
- Bertossi, L. E.; Hunter, A.; and Schaub, T., eds. 2005. *Inconsistency Tolerance [result from a Dagstuhl seminar]*, volume 3300 of *Lecture Notes in Computer Science*. Springer.
- Cabalar, P.; Pearce, D.; and Valverde, A. 2009. A revised concept of safety for general answer set programs. In Erdem, E.; Lin, F.; and Schaub, T., eds., *LPNMR*, volume 5753 of *Lecture Notes in Computer Science*, 58–70. Springer.
- de Bruijn, J.; Pearce, D.; Polleres, A.; and Valverde, A. 2007. Quantified equilibrium logic and hybrid rules. In Marchiori, M.; Pan, J. Z.; and de Sainte Marie, C., eds., *RR*, volume 4524 of *Lecture Notes in Computer Science*, 58–72. Springer.
- Eiter, T.; Ianni, G.; Lukasiewicz, T.; Schindlauer, R.; and Tompits, H. 2008. Combining answer set programming with description logics for the semantic web. *Artif. Intell.* 172(12-13):1495–1539.
- Eiter, T.; Fink, M.; and Moura, J. 2010. Paracoherent answer set programming. In Lin, F.; Sattler, U.; and Truszczynski, M., eds., *KR*, 486–496. AAAI Press.
- Fink, M.; El Ghali, A.; Chniti, A.; Korf, R.; Schwichtenberg, A.; Lévy, F.; Pührer, J.; and Eiter, T. 2011. D2.3 Consistency maintenance. Final report. Technical report, ONTORULE IST-2009-231875 Project. Forthcoming.
- Fink, M. 2012. A paraconsistent semantics for hybrid theories. Technical Report INFSYS RR-1843-12-02, Institut für Informationssysteme, Technische Universität Wien, Austria.
- Gelfond, M., and Lifschitz, V. 1991. Classical negation in logic programs and disjunctive databases. *New Generation Comput.* 9(3/4):365–386.
- Grosof, B. N.; Horrocks, I.; Volz, R.; and Decker, S. 2003. Description logic programs: combining logic programs with description logic. In *WWW*, 48–57.
- Heymans, S.; de Bruijn, J.; Predoiu, L.; Feier, C.; and Nieuwenborgh, D. V. 2008. Guarded hybrid knowledge bases. *TPLP* 8(3):411–429.
- Heymans, S.; Korf, R.; Erdmann, M.; Pührer, J.; and Eiter, T. 2010. F-logic#: Loosely coupling f-logic rules and ontologies. In Huang, J. X.; King, I.; Raghavan, V. V.; and Rueger, S., eds., *Web Intelligence*, 248–255. IEEE.
- Huang, S.; Li, Q.; and Hitzler, P. 2011. Paraconsistent semantics for hybrid mknf knowledge bases. In Rudolph, S., and Gutierrez, C., eds., *RR*, volume 6902 of *Lecture Notes in Computer Science*, 93–107. Springer.
- Hunter, A. 1998. *Paraconsistent logics*. Norwell, MA, USA: Kluwer Academic Publishers. 11–36.
- Kifer, M. 2005. Rules and ontologies in f-logic. In Eisinger, N., and Maluszynski, J., eds., *Reasoning Web*, volume 3564 of *Lecture Notes in Computer Science*, 22–34. Springer.
- Lee, J.; Lifschitz, V.; and Palla, R. 2008. Safe formulas in the general theory of stable models (preliminary report). In de la Banda M. G., and Pontelli, E., eds., *ICLP*, volume 5366 of *Lecture Notes in Computer Science*, 672–676. Springer.
- Lembo, D.; Lenzerini, M.; Rosati, R.; Ruzzi, M.; and Savo, D. F. 2011. Query rewriting for inconsistent dl-lite ontologies. In Rudolph, S., and Gutierrez, C., eds., *RR*, volume 6902 of *Lecture Notes in Computer Science*, 155–169. Springer.
- Lifschitz, V.; Pearce, D.; and Valverde, A. 2007. A characterization of strong equivalence for logic programs with variables. In Baral, C.; Brewka, G.; and Schlipf, J. S., eds., *LPNMR*, volume 4483 of *Lecture Notes in Computer Science*, 188–200. Springer.
- Lukasiewicz, T. 2007. A novel combination of answer set programming with description logics for the semantic web. In Franconi, E.; Kifer, M.; and May, W., eds., *ESWC*, volume 4519 of *Lecture Notes in Computer Science*, 384–398. Springer.
- Ma, Y.; Hitzler, P.; and Lin, Z. 2007. Algorithms for paraconsistent reasoning with owl. In Franconi, E.; Kifer, M.; and May, W., eds., *ESWC*, volume 4519 of *Lecture Notes in Computer Science*, 399–413. Springer.
- Ma, Y.; Hitzler, P.; and Lin, Z. 2008. Paraconsistent reasoning for expressive and tractable description logics. In Baader, F.; Lutz, C.; and Motik, B., eds., *Description Logics*, volume 353 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Motik, B., and Rosati, R. 2010. Reconciling description logics and rules. *J. ACM* 57(5).
- Odintsov, S. P., and Pearce, D. 2005. Routley semantics for answer sets. In Baral, C.; Greco, G.; Leone, N.; and Terracina, G., eds., *LPNMR*, volume 3662 of *Lecture Notes in Computer Science*, 343–355. Springer.
- Pearce, D., and Valverde, A. 2005. A first order nonmonotonic extension of constructive logic. *Studia Logica* 80(2-3):321–346.
- Pearce, D., and Valverde, A. 2008. Quantified equilibrium logic and foundations for answer set programs. In de la Banda M. G., and Pontelli, E., eds., *ICLP*, volume 5366 of *Lecture Notes in Computer Science*, 546–560. Springer.
- Rosati, R. 2005b. Semantic and computational advantages of the safe integration of ontologies and rules. In Fages, F., and Soliman, S., eds., *PPSWR*, volume 3703 of *Lecture Notes in Computer Science*, 50–64. Springer.
- Rosati, R. 2006. DL+log: Tight integration of description logics and disjunctive datalog. In Doherty, P.; Mylopoulos, J.; and Welty, C. A., eds., *KR*, 68–78. AAAI Press.
- Rosati, R. 2008. On combining description logic ontologies and nonrecursive datalog rules. In Calvanese, D., and Lausen, G., eds., *RR*, volume 5341 of *Lecture Notes in Computer Science*, 13–27. Springer.
- Rosati, R. 2011. On the complexity of dealing with inconsistency in description logic ontologies. In Walsh, T., ed., *IJCAI*, 1057–1062. AAAI Press.
- Routley, R. 1974. Semantical analyses of propositional systems of fitch and nelson. *Studia Logica* 33(3):283–398.
- Sakama, C., and Inoue, K. 1995. Paraconsistent stable semantics for extended disjunctive programs. *J. Log. Comput.* 5(3):265–285.
- van Dalen, D. 1983. *Logic and Structure*. Springer.