

Stable Models in Generalized Possibilistic Logic

Didier Dubois

Université Paul Sabatier, CNRS
IRIT, Toulouse, France
dubois@irit.fr

Henri Prade

Université Paul Sabatier, CNRS
IRIT, Toulouse, France
prade@irit.fr

Steven Schockaert

Cardiff University, School of Computer
Science & Informatics, Cardiff, UK
s.schockaert@cs.cardiff.ac.uk

Abstract

Possibilistic logic is a well-known logic for reasoning under uncertainty, which is based on the idea that the epistemic state of an agent can be modeled by assigning to each possible world a degree of possibility, taken from a totally ordered, but essentially qualitative scale. Recently, a generalization has been proposed that extends possibilistic logic to a meta-epistemic logic, endowing it with the capability of reasoning about epistemic states, rather than merely constraining them. In this paper, we further develop this generalized possibilistic logic (GPL). We introduce an axiomatization showing that GPL is a fragment of a graded version of the modal logic KD, and we prove soundness and completeness w.r.t. a semantics in terms of possibility distributions. Next, we reveal a close link between the well-known stable model semantics for logic programming and the notion of minimally specific models in GPL. More generally, we analyze the relationship between the equilibrium logic of Pearce and GPL, showing that GPL can essentially be seen as a generalization of equilibrium logic, although its notion of minimal specificity is slightly more demanding than the notion of minimality underlying equilibrium logic.

Introduction

In its basic form, possibilistic logic (Dubois, Lang, and Prade 1994) is a logic for reasoning about uncertain propositional knowledge, in which formulas take the form of pairs (α, λ) with α a propositional formula and $\lambda \in]0, 1]$ a degree of certainty. Its semantics are defined in terms of possibility distributions, which are mappings π from the set of all propositional interpretations (or *worlds*) Ω to $[0, 1]$. For $\omega \in \Omega$, $\pi(\omega)$ reflects the degree to which world ω is *possible*, i.e. to what extent it is compatible with available knowledge. Given a possibility distribution π , the associated measures of possibility Π and necessity N reflect to what extent a proposition is possible and necessary, respectively:

$$\Pi(\alpha) = \sup\{\pi(\omega) \mid \omega \models \alpha\} \quad (1)$$

$$N(\alpha) = 1 - \Pi(\neg\alpha) = \inf\{1 - \pi(\omega) \mid \omega \models \neg\alpha\} \quad (2)$$

A possibility distribution π is a model of (α, λ) if $N(\alpha) \geq \lambda$ for N the associated necessity measure.

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

An important aspect of possibilistic logic is that models correspond to epistemic states, rather than to propositional interpretations, which forms a natural basis for epistemic reasoning. However, possibilistic logic only takes sets of formulas of the form (α, λ) into account, which we could interpret as conjunctions of assertions of the form $N(\alpha) \geq \lambda$. In some applications, on the other hand, we may want to link such assertions using different propositional connectives. In logic programming, for instance, a (negation-free) rule such as $\alpha \rightarrow \beta$ intuitively means that whenever α is known to be true, we should accept β to be true as well. This could be expressed using necessity measures and material implication as $N(\alpha) \geq 1 \Rightarrow N(\beta) \geq 1$. This implication, however, cannot be expressed in possibilistic logic, an observation which stands in stark contrast to the expressivity of modal logics for epistemic reasoning. In (Banerjee and Dubois 2009), a so-called Meta-Epistemic Logic (MEL) was introduced as a first step to bridge this gap, in the form of a simple modal logic with a semantics in terms of Boolean possibility distributions (i.e. possibility distributions π such that $\pi(\omega) \in \{0, 1\}$ for every $\omega \in \Omega$). Essentially, MEL is a fragment of the modal logic KD, in which neither the nesting of modalities nor the occurrence of non-modal propositional formulas is allowed. Recently, generalized possibilistic logic (GPL) was introduced as a graded version of MEL (Dubois, Prade, and Schockaert 2011; Dubois and Prade 2011), developing an original proposal of (Dubois and Prade 2007).

The aim of this paper is twofold. On the one hand, we further develop GPL by introducing an axiomatization for it and proving soundness and completeness. On the other hand, we show how the notion of *stable models* from the semantics of logic programming (Gelfond and Lifschitz 1988) can naturally be captured using GPL. As equilibrium logic constitutes one of the most general approaches to the stable model semantics (Lifschitz 2008), we use it as our starting point, and show how equilibrium logic can be simulated in GPL to a large extent. We furthermore advocate that GPL has important advantages over equilibrium logic as a general framework for capturing the stable model semantics in an intuitive and general way. First, as equilibrium logic essentially uses the syntax of classical logic, but with a different underlying semantics, the intuitive meaning of equilibrium logic formulas is not always easy to grasp. When

expressing these formulas in GPL, however, the underlying meaning becomes explicit, as GPL only uses classical connectives, together with modalities that have an intuitive meaning. Second, as will become clear below, equilibrium logic corresponds to a limited fragment of GPL, in which modalities only occur in front of literals. This means that assertions such as “it is known that either a or b holds” cannot be expressed in equilibrium logic, which is only able to express that “either a is known or b is known”. As is discussed in (Bauters et al. 2011), a similar limitation also hampers answer set programming (ASP). Due to its increased syntactic freedom, GPL has the potential to further generalize the stable model semantics, while also being more intuitive to express pieces of knowledge.

The idea to use a modal logic to capture the stable model semantics of logic programming has been proposed early on. In particular, (Lifschitz and Schwarz 1993) proposes a way to interpret answer set programs as theories in autoepistemic logic (Moore 1985). While this work shares several of the advantages of our approach, the translation from answer set programming to autoepistemic logic does not readily generalize to arbitrary propositional theories. Although the semantics of the modality occurring in autoepistemic logic is closely related to the semantics of MEL in terms of subsets of interpretations, autoepistemic logic does not allow the use of intermediary certainty degrees. As we show in this paper, the ability to discriminate between fully certain statements and statements that are only somewhat certain plays a crucial role. Another attempt to provide the stable model semantics with an epistemic flavor has been presented in (Loyer and Straccia 2006), where an algebraic approach is taken, interpreting rule bases using an immediate consequence operator that assigns to every atom a state from a particular bilattice. While this work offers a unifying view of different logic programming semantics, it is not primarily aimed at generalizing the stable model semantics, nor clarifying its intuitions, and is thus different in spirit from our work. Although related, this paper is also different in spirit from existing work on possibilistic ASP (Nicolas et al. 2006; Bauters et al. 2010), which is primarily concerned with adding uncertainty to ASP rules.

The paper is structured as follows. In the following section we recall the basics of equilibrium logic. Then, we provide the syntax and semantics of GPL, introduce an axiomatization for this logic, and prove its soundness and completeness. In the subsequent section, we show how much of equilibrium logic can be encapsulated in GPL. In particular, we propose a translation from equilibrium logic to GPL theories, such that the equilibrium models of the equilibrium logic theory correspond to a particular class of models of the GPL theory. Then, we show how this transformation can be simplified in the case of disjunctive answer set programs, before presenting complexity results and concluding remarks.

Equilibrium Logic

Equilibrium logic was introduced by Pearce with the aim of extending the notion of answer set to general propositional theories (Pearce 1997; 2006). The formulation of this logic is based on an extension of *the logic of here-and-there* with

strong negation. The logic of here-and-there, also known as Smetanich logic, is known to be the strongest intermediate logic that is properly included in classical logic (Chagrov and Zakharyashev 1997; Pearce 1997). This logic can semantically be characterized as a three-valued logic. Alternatively, however, it can also be characterized in terms of Kripke frames, using a two-valued valuation in two worlds, called h (here) and t (there).

The semantics of equilibrium logic is also based on these two worlds, by considering a three-valued valuation in both worlds (Pearce 1997). In particular, a *valuation* V is defined as a mapping from $\{h, t\} \times At$ to $\{-1, 0, 1\}$, where At is the set of all atoms in the language, such that

$$(V(h, a) \neq 0) \Rightarrow (V(h, a) = V(t, a)) \quad (3)$$

The intuition is that $V(s, a) = 1$ means that a is known to be true in world s , $V(s, a) = -1$ means that a is known to be false in world s , and $V(s, a) = 0$ means that the truth of a is unknown in world s . Furthermore, the there-world is assumed to be a refinement of the here-world, i.e., atoms whose truth value is unknown ‘here’ may have a known truth value ‘there’, but whenever the truth value of a is already known ‘here’ it has to be the same ‘there’. Since t is a refinement of h , there are five possibilities for the valuation of an atom a . Hence, the logic defined in this way is actually a five-valued logic, which is called $N5$ in (Pearce 2006)¹. When there may be cause for confusion, we will refer to $N5$ valuations to denote $\{h, t\} \times At \rightarrow \{-1, 0, 1\}$ mappings V that satisfy (3). Let Lit be the set of all literals, i.e., $l \in Lit = At \cup \{\sim a \mid a \in At\}$, where $\sim$ stands for a strong negation (reflecting a negation with classical, involutive behavior). For a valuation V , let V_h and V_t be the sets of literals that are *true* in worlds h and t :

$$V_h = \{l \in Lit \mid V(h, l) = 1\} \quad V_t = \{l \in Lit \mid V(t, l) = 1\}$$

Clearly, the property $V_h \subseteq V_t$ holds, due to (1). So an $N5$ model can be interpreted as a nested pair (V_h, V_t) of literals.

By defining $\leq$ as the binary relation $\{(h, h), (t, t), (h, t)\}$, valuations are extended to arbitrary formulas as follows:

$$\begin{aligned} V(s, \sim \alpha) &= -V(s, \alpha) \\ V(s, \alpha \otimes \beta) &= \min(V(s, \alpha), V(s, \beta)) \\ V(s, \alpha \oplus \beta) &= \max(V(s, \alpha), V(s, \beta)) \\ V(s, \alpha \triangleright \beta) &= \begin{cases} 1 & \text{if } \forall s' \geq s. (V(s', \alpha) = 1) \Rightarrow (V(s', \beta) = 1) \\ -1 & \text{if } V(s, \alpha) = 1 \text{ and } V(s, \beta) = -1 \\ 0 & \text{otherwise} \end{cases} \\ V(s, \text{not } \alpha) &= \begin{cases} 1 & \text{if } \forall s' \geq s. V(s', \alpha) < 1 \\ -1 & \text{if } V(s, \alpha) = 1 \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

where we write $\otimes$, $\oplus$ and $\triangleright$ for conjunction, disjunction and implication, to avoid confusion with the corresponding operators in classical logic. Moreover, as opposed to the involutive strong negation $\sim$, negation as failure is denoted by *not*.

¹This logic is called $N2$ in (Pearce 1997).

These connectives are perhaps more easily understood by recursively extending the sets of literals V_h and V_t to sets $\mathcal{V}_h$ and $\mathcal{V}_t$ of more general formulas as follows

- $\alpha \otimes \beta \in \mathcal{V}_s$ if and only if $\alpha \in \mathcal{V}_s$ and $\beta \in \mathcal{V}_s, s = h, t$;
- $\alpha \oplus \beta \in \mathcal{V}_s$ if and only if $\alpha \in \mathcal{V}_s$ or $\beta \in \mathcal{V}_s, s = h, t$;
- $\alpha \triangleright \beta \in \mathcal{V}_h$ if and only if $\alpha \in \mathcal{V}_h$ implies $\beta \in \mathcal{V}_h$;
and $\alpha \in \mathcal{V}_t$ implies $\beta \in \mathcal{V}_t$;
- $\alpha \triangleright \beta \in \mathcal{V}_t$ if and only if $\alpha \in \mathcal{V}_t$ implies $\beta \in \mathcal{V}_t$;
- $\text{not } \alpha \in \mathcal{V}_h$ if and only if $\alpha \notin \mathcal{V}_t$ (and thus $\alpha \notin \mathcal{V}_h$);
- $\sim \text{not } \alpha \in \mathcal{V}_h$ if and only if $\alpha \in \mathcal{V}_h$;

$\sim \text{not } \alpha \in \mathcal{V}_t \setminus \mathcal{V}_h$ if and only if $\alpha \in \mathcal{V}_t \setminus \mathcal{V}_h$, etc.

An (N5) valuation V is called an (N5) model of a set of formulas Θ , written $V \models \Theta$, if for each $\alpha \in \Theta$, it holds that $V(h, \alpha) = V(t, \alpha) = 1$ (i.e., $\alpha \in \mathcal{V}_h$).

Equilibrium logic is obtained from N5 logic by restricting attention to particular N5 models, called *equilibrium models*. A model is called *h-minimal* if its *here* world is as little committing as possible, given its particular *there* world².

Definition 1. (Pearce 2006) Let the ordering $\preceq$ be defined for two N5 valuations V and V' as $V \preceq V'$ iff $V_t = V'_t$ and $V_h \subseteq V'_h$. An N5 model V of a set of formulas Θ is then called *h-minimal* if it is minimal w.r.t. $\preceq$ among all models of Θ , i.e., for every other model V' of Θ it holds that either $V_t \neq V'_t$ or $V_h \not\subseteq V'_h$.

Note that minimality refers to the set of N5 literals that are verified by a valuation, and not to the set of atoms. The notion of *h-minimality* makes the connection with ASP more explicit: what is true ‘there’ can intuitively be understood as a guess of what can be derived from available knowledge, whereas what is true ‘here’ can actually be derived. Recall that in ASP we are interested in the case where the guess about what can be derived coincides with what can actually be derived. Accordingly, equilibrium models are *h-minimal* models whose valuation in *h* and *t* coincides.

Definition 2. (Pearce 1997) A *h-minimal* model V of a set of formulas Θ is called an *equilibrium model* if $V_h = V_t$.

A set of N5 formulas Θ corresponds to a (disjunctive) answer set program if it consists of formulas of the form

$$l_1 \otimes \dots \otimes l_s \otimes \text{not } p_1 \otimes \dots \otimes \text{not } p_m \triangleright q_1 \oplus \dots \oplus q_n \quad (4)$$

and facts of the form

$$l_1 \oplus \dots \oplus l_n \quad (5)$$

If there are no occurrences of *not* in Θ (i.e., $m = 0$ in all rules of the form (4)), the answer set of Θ is the unique minimal set of literals A such that $\{l_1, \dots, l_n\} \cap A \neq \emptyset$ for every fact of the form (5) and $\{q_1, \dots, q_n\} \cap A \neq \emptyset$ whenever $\{l_1, \dots, l_s\} \subseteq A$ for every rule of the form (4). In the general case, the notion of answer set is defined in terms of the Gelfond-Lifschitz reduct. Let A be a set of literals, then the Gelfond-Lifschitz reduct Θ^A w.r.t. A is obtained from

²In (Pearce 1997) the notion of *h-minimality* is defined in a slightly different way. The difference is irrelevant, however, w.r.t. the definition of equilibrium models.

Θ by removing all rules of the form (4) for which $p_i \in A$ for some $i \in \{1, \dots, m\}$, and replacing all other rules by $l_1 \otimes \dots \otimes l_s \otimes \triangleright q_1 \oplus \dots \oplus q_n$. Then A is called an answer set of Θ iff A is the answer set of the reduct Θ^A . The following result shows that equilibrium logic properly extends answer set programming.

Proposition 1. (Pearce 1997) Let Θ be an equilibrium logic theory which corresponds to an answer set program. Furthermore, let S be a consistent set of literals (i.e., a and $\sim a$ cannot be both in S , for any atom in At). Then S is an answer set of Θ iff Θ has an equilibrium model V such that $S = V_t$.

Example 1. Let us consider the following equilibrium logic theory Θ :

$$\Theta = \{a, (a \otimes \text{not } b \triangleright c) \oplus (a \triangleright d)\}$$

It is easy to verify that an N5 valuation V is an N5 model of Θ iff $V(h, a) = 1$ and either $V(h, c) = 1$, or $V(t, b) = 1$, or $V(h, d) = 1$. This means that there are three *h-minimal* models V_1, V_2, V_3 of Θ :

	here				there			
	a	b	c	d	a	b	c	d
V_1	1	0	1	0	1	0	1	0
V_2	1	0	0	0	1	1	0	0
V_3	1	0	0	1	1	0	0	1

Of these three *h-minimal* models, only V_1 and V_3 are equilibrium models, as $V_2(h, b) \neq V_2(t, b)$. Note that the model corresponding to (1, 0, 1, 1) ‘here’ and (1, 0, 1, 1) ‘there’ is not *h-minimal*, as there is another model which coincides ‘there’ but corresponds to (1, 0, 1, 0) ‘here’.

Generalized possibilistic logic

Generalized possibilistic logic (Dubois and Prade 2011) extends possibilistic logic in the sense that assertions of the form (α, λ) can be combined using any propositional connective, rather than only conjunction. The aim of this section is to familiarize the reader with the semantics of GPL, and to introduce an axiomatization for this logic. In this section α, β , etc. denote propositions in standard propositional logic, formed with standard connectives, $\wedge, \neg$, and we use abbreviations $\alpha \vee \beta \equiv \neg(\neg\alpha \wedge \neg\beta)$ and $\alpha \rightarrow \beta \equiv \neg(\alpha \wedge \neg\beta)$.

To emphasize that the semantics of GPL are based on possibility theory, in this paper we use a slightly different notation than the standard modal logic syntax of MEL (Banerjee and Dubois 2009). Moreover, our notation also differs from (Dubois and Prade 2011), which uses a notation close to standard possibilistic logic. We improve readability w.r.t. the latter paper while emphasizing the link with modal logic. Let $\Lambda = \{0, \frac{1}{k}, \frac{2}{k}, \dots, 1\}$ with $k \in \mathbb{N} \setminus \{0\}$ be the set of certainty degrees under consideration, $\Lambda^+ = \Lambda \setminus \{0\}$. Well-formed formulas in generalized possibilistic logic (GPL) are defined as follows:

- If α is a propositional formula over the set of atoms At and $\lambda \in \Lambda^+$, then $N_\lambda(\alpha)$ is a well-formed formula.
- If γ and δ are well-formed formulas, then $\neg\gamma$ and $\gamma \wedge \delta$ are also well-formed formulas.

Intuitively, $\mathbf{N}_1(\alpha)$ means that it is completely certain that α is true, whereas $\mathbf{N}_\lambda(\alpha)$ with $\lambda < 1$ means that there is evidence that suggests that α is true, and none that suggests that it is false (i.e. it is considered more plausible that α is true than that α is false). Note that we need $k \geq 2$, ensuring that there are at least three certainty levels, to distinguish between complete and partial certainty. Formally, an agent asserting $\mathbf{N}_\lambda(\alpha)$ has an epistemic state π such that $N(\alpha) \geq \lambda > 0$. Hence $\neg \mathbf{N}_\lambda(\alpha)$ stands for $N(\alpha) < \lambda$, which means $\Pi(\neg\alpha) \geq 1 - \lambda + \frac{1}{k} \in \Lambda^+$ (using (2)). Let us introduce some useful abbreviations: $\nu(\lambda) = 1 - \lambda + \frac{1}{k}, \forall \lambda \in \Lambda^+$, $\Pi_\lambda(\alpha) \equiv \neg \mathbf{N}_{\nu(\lambda)}(\neg\alpha)$. Then $\Pi_1(\alpha)$ means that α is fully compatible with our available beliefs (i.e. nothing prevents α from being true), while $\Pi_\lambda(\alpha)$ with $\lambda < 1$ means that α cannot be fully excluded ($\Pi(\alpha) \geq \lambda$).

The semantics of GPL are defined in terms of normalized possibility distributions (i.e., s.t. $\exists \omega \pi(\omega) = 1$) over propositional interpretations, where possibility degrees are limited to Λ . A model of a GPL formula is such a possibility distribution which satisfies:

- π is a model of $\mathbf{N}_\lambda(\alpha)$ iff $N(\alpha) \geq \lambda$;
- π is a model of $\gamma_1 \wedge \gamma_2$ iff π is a model of γ_1 and a model of γ_2 ;
- π is a model of $\neg\gamma_1$ iff π is not a model of γ_1 ;

where N is the necessity measure induced by π . As usual, π is called a model of a set of GPL formulas K , written $\pi \models K$ if it is a model of each of the formulas in the set. It is called a minimally specific model, if there is no model $\pi' \neq \pi$ such that $\pi'(\omega) \geq \pi(\omega)$ for each possible world ω . We write $K \models \phi$, for K a set of GPL formulas and ϕ a GPL formula, iff every model of K is also a model of ϕ . We call a possibilistic model π Boolean if $\pi(\omega) \in \{0, 1\}$ for every possible world ω .

We consider the following axiomatization, which closely parallels the one of MEL (Banerjee and Dubois 2009).

The axioms of classical logic (6)

$\mathbf{N}_\lambda(\alpha) \rightarrow \mathbf{N}_\lambda(\beta)$, if $\alpha \rightarrow \beta$ is a (classical) tautology (7)

$\mathbf{N}_\lambda(\alpha \wedge \beta) \rightarrow \mathbf{N}_\lambda(\alpha) \wedge \mathbf{N}_\lambda(\beta)$ (8)

$\mathbf{N}_\lambda(\alpha) \wedge \mathbf{N}_\lambda(\beta) \rightarrow \mathbf{N}_\lambda(\alpha \wedge \beta)$ (9)

$\mathbf{N}_1(\top)$ (10)

$\mathbf{N}_\lambda(\alpha) \rightarrow \Pi_1(\alpha)$ (11)

$\mathbf{N}_{\lambda_1}(\alpha) \rightarrow \mathbf{N}_{\lambda_2}(\alpha)$, if $\lambda_1 \geq \lambda_2$ (12)

and the modus ponens rule. Note in particular that from (6), (9) and (7), we can derive a weighted version of axiom **K**:

$$\mathbf{N}_\lambda(\alpha \rightarrow \beta) \rightarrow (\mathbf{N}_\lambda(\alpha) \rightarrow \mathbf{N}_\lambda(\beta))$$

In fact (7) can be replaced by this version of **K**, as when λ is fixed we get a fragment of the modal logic KD.

Lemma 1. *The following inference rules can be proved:*

- $\{\mathbf{N}_{\lambda_1}(\alpha), \mathbf{N}_{\lambda_2}(\alpha \rightarrow \beta)\} \vdash \mathbf{N}_{\min(\lambda_1, \lambda_2)}(\beta)$
- $\{\Pi_{\lambda_1}(\alpha), \mathbf{N}_{\lambda_2}(\alpha \rightarrow \beta)\} \vdash \Pi_{\lambda_1}(\beta)$ if $\lambda_2 > 1 - \lambda_1$

The first one is obtained using (12) on both premises (weakening them to level $\min(\lambda_1, \lambda_2)$), then getting

$\mathbf{N}_{\min(\lambda_1, \lambda_2)}(\alpha) \rightarrow \mathbf{N}_{\min(\lambda_1, \lambda_2)}(\beta)$ using modus ponens on **K** and $\mathbf{N}_{\min(\lambda_1, \lambda_2)}(\alpha \rightarrow \beta)$, and then modus ponens again.

The second one is obtained by proving $\mathbf{N}_{\nu(\lambda_1)}(\neg\alpha)$ from $\mathbf{N}_{\lambda_2}(\alpha \rightarrow \beta)$ and $\mathbf{N}_{\nu(\lambda_1)}(\neg\beta)$ in the same way (just rewriting $\alpha \rightarrow \beta$ as $\neg\beta \rightarrow \neg\alpha$). However, we need to assume $\nu(\lambda_1) \leq \lambda_2$ in order to weaken $\mathbf{N}_{\lambda_2}(\alpha \rightarrow \beta)$ into $\mathbf{N}_{\nu(\lambda_1)}(\alpha \rightarrow \beta)$. And $\nu(\lambda_1) \leq \lambda_2$ is equivalent to $1 - \lambda_1 + 1/k \leq \lambda_2$, i.e., $\lambda_2 > 1 - \lambda_1$ ³.

Proposition 2 (Soundness). *Let K be a GPL theory, i.e., a set of GPL formulas. Assume that the GPL formula ϕ can be derived from K using modus ponens and the axioms (6)–(12). It holds that $K \models \phi$.*

Proposition 3 (Completeness). *Let K be a GPL theory and ϕ a GPL formula. If $K \models \phi$ then it holds that ϕ can be derived from K using modus ponens and the axioms (6)–(12).*

Relationship with equilibrium logic

In (Dubois, Prade, and Schockaert 2011), it has been shown how answer set programming rules can be expressed in GPL in a semantics-preserving way. The main idea underlying this translation is to interpret a rule such as $a \triangleright b$ (where we use the same atoms a, b in propositional logic and answer set programming) as stating that when a is certain (to some degree), then also b is certain (to some degree). As such, we may relate the rule $a \triangleright b$ to the set of GPL formulas $\{\mathbf{N}_\lambda(a) \rightarrow \mathbf{N}_\lambda(b) \mid \lambda \in \Lambda^+\}$. Furthermore, following a proposal in (Bauters et al. 2010) to interpret negation as failure in a possibilistic logic setting, it is considered that *not a* corresponds to $\Pi_1(\neg a)$. A rule such as *not a* $\triangleright b$ is then taken to correspond to the GPL formula $\Pi_1(\neg a) \rightarrow \mathbf{N}_1(b)$, translating the intuition that as soon as it is consistent to assume that $\neg a$ holds, we may derive b with (full) certainty.

However, it is clear that the full generality of GPL is not needed to capture the semantics of logic programming or equilibrium logic. In particular, as degrees of certainty do not occur in equilibrium logic, we may be inclined to let $\Lambda^+ = \{1\}$. In this case, however, we have that $\neg \mathbf{N}_1(b) \equiv \Pi_1(\neg b)$ and thus that $\Pi_1(\neg a) \rightarrow \mathbf{N}_1(b)$ is equivalent to $\neg \mathbf{N}_1(a) \rightarrow \neg \Pi_1(\neg b)$ and, by contraposition, also to $\Pi_1(\neg b) \rightarrow \mathbf{N}_1(a)$. This would mean, under the above view, that the rules *not b* $\triangleright a$ and *not a* $\triangleright b$ would be equivalent, which does not agree with the stable model semantics. In the translation from disjunctive ASP to autoepistemic logic, proposed in (Lifschitz and Schwarz 1993), this issue is tackled by translating *not b* $\triangleright a$ to the formula $\neg K(b) \rightarrow (K(a) \wedge a)$, where K is a modal operator which plays a role similar to $\mathbf{N}_1$. In GPL, however, expressions such as $\mathbf{N}_1(a) \wedge a$ are not allowed, as it is hard to provide an intuitive (epistemic) semantics for them. As will become clear, however, adding one intermediary certainty level, i.e. choosing $\Lambda^+ = \{1/2, 1\}$, is sufficient to enable us to capture the semantics of rules, and more generally equilibrium logic formulas, within GPL. In this case, we can discriminate between propositions in which we are fully certain

³If $\nu(\lambda_1) > \lambda_2$, the weakening axiom (12) leads us to derive $\mathbf{N}_{\lambda_2}(\neg\alpha)$, whose negation is weaker than the premiss $\Pi_{\lambda_1}(\alpha)$.

(strong necessity), and propositions which we merely consider to be more plausible than not (weak necessity). To emphasize the qualitative nature of the certainty scale, we will write $\mathbf{N}_w(\cdot)$, $\mathbf{N}_s(\cdot)$, $\mathbf{\Pi}_w(\cdot)$ and $\mathbf{\Pi}_s(\cdot)$ instead of $\mathbf{N}_{1/2}(\cdot)$, $\mathbf{N}_1(\cdot)$, $\mathbf{\Pi}_{1/2}(\cdot)$, $\mathbf{\Pi}_1(\cdot)$ respectively. Now we only have the equivalence $\neg\mathbf{N}_s(b) \equiv \mathbf{\Pi}_w(\neg b)$.

The fact that we need an intermediate certainty level is also closely related to the observation that we can then model 5 different epistemic statuses for each atom, which correspond to the 5 truth values that can be assigned to atoms in $N5$ logic. The intuition is that what is already true *here* is treated as necessarily true, whereas what is only true *there* is treated as more plausible than not, but not fully certain. Under this view, and using again the same atoms in $N5$ and in propositional logic, we assume that a valuation V for which

- $V(h, a) = V(t, a) = 1$ ($a \in V_h$) corresponds to a possibility distribution which satisfies $\mathbf{N}_s(a)$,
- $V(h, a) = V(t, a) = -1$ ($\sim a \in V_h$) corresponds to a possibility distribution which satisfies $\mathbf{N}_s(\neg a)$,
- $V(h, a) = 0$ and $V(t, a) = 1$ ($a \in V_t \setminus V_h$) corresponds to a possibility distribution which satisfies $\mathbf{N}_w(a) \wedge \neg\mathbf{N}_s(a)$, or equivalently $\mathbf{N}_w(a) \wedge \mathbf{\Pi}_w(\neg a)$,
- $V(h, a) = 0$ and $V(t, a) = -1$ ($\sim a \in V_t \setminus V_h$) corresponds to a possibility distribution which satisfies $\mathbf{N}_w(\neg a) \wedge \neg\mathbf{N}_s(\neg a)$, or equivalently $\mathbf{N}_w(\neg a) \wedge \mathbf{\Pi}_w(a)$,
- $V(h, a) = V(t, a) = 0$ ($a, \sim a \notin V_t$) corresponds to a possibility distribution which satisfies $\neg\mathbf{N}_w(a) \wedge \neg\mathbf{N}_w(\neg a)$, or equivalently $\mathbf{\Pi}_s(\neg a) \wedge \mathbf{\Pi}_s(a)$.

Note that the $N5$ strong negation $\sim$ exactly translates into the standard propositional negation $\neg$ here. By identifying the assignment of truth values of $N5$ logic to atoms with corresponding GPL formulas, we emphasize the epistemic nature of asserting formulas in $N5$.

To develop this idea, it is useful to represent an $N5$ valuation V as a mapping σ from formulas to $\{-2, -1, 0, 1, 2\}$ such that $\sigma(\alpha) = V(h, \alpha) + V(t, \alpha)$ (Pearce 2006). Note that because the set of literals V_h is constrained to be a subset of V_t , such a mapping σ unambiguously defines a valuation V . We can then consider formulas of the form $\alpha \geq i$ or $\alpha \leq j$ where α is an $N5$ formula, $i \in \{-1, 0, 1, 2\}$ and $j \in \{-2, -1, 0, 1\}$. Such a formula is satisfied by a valuation V iff the corresponding mapping σ is such that $\sigma(\alpha) \geq i$ or $\sigma(\alpha) \leq j$ respectively. Under the above view, we can associate a GPL formula $\Phi(\alpha \geq i)$ or $\Phi(\alpha \leq j)$ with each assertion of the form $\alpha \geq i$ or $\alpha \leq j$. In particular, we easily prove, based on the above encoding of the five truth-values as modalities in GPL:

$$\begin{aligned} \Phi(a \geq 2) &= \mathbf{N}_s(a) & \Phi(a \geq 1) &= \mathbf{N}_w(a) \\ \Phi(a \geq 0) &= \mathbf{\Pi}_s(a) & \Phi(a \geq -1) &= \mathbf{\Pi}_w(a) \\ \Phi(a \leq -2) &= \mathbf{N}_s(\neg a) & \Phi(a \leq -1) &= \mathbf{N}_w(\neg a) \\ \Phi(a \leq 0) &= \mathbf{\Pi}_s(\neg a) & \Phi(a \leq 1) &= \mathbf{\Pi}_w(\neg a) \end{aligned}$$

More generally, we can use this idea to express that an arbitrary $N5$ formula has a given truth value. For example, it is not hard to see that $a \otimes b \geq i$ iff $a \geq i$ and $b \geq i$. Therefore, the assertion $a \otimes b \geq i$ can be expressed in GPL

as $\Phi(a \geq i) \wedge \Phi(b \geq i)$. In general, for $i \in \{-1, 0, 1, 2\}$ and α and β two $N5$ formulas, we can prove the following translation rules are correct:

$$\begin{aligned} \Phi(\alpha \otimes \beta \geq i) &= \Phi(\alpha \geq i) \wedge \Phi(\beta \geq i) \\ \Phi(\alpha \oplus \beta \geq i) &= \Phi(\alpha \geq i) \vee \Phi(\beta \geq i) \\ \Phi(\sim \alpha \geq i) &= \Phi(\alpha \leq -i) \\ \Phi(\text{not } \alpha \geq i) &= \begin{cases} \Phi(\alpha \leq 0) & \text{if } i \geq 0 \\ \Phi(\alpha \leq 1) & \text{if } i = -1 \end{cases} \\ \Phi(\alpha \triangleright \beta \geq 2) &= (\Phi(\alpha \geq 2) \rightarrow \Phi(\beta \geq 2)) \\ &\quad \wedge (\Phi(\alpha \geq 1) \rightarrow \Phi(\beta \geq 1)) \\ \Phi(\alpha \triangleright \beta \geq 1) &= \Phi(\alpha \geq 1) \rightarrow \Phi(\beta \geq 1) \\ \Phi(\alpha \triangleright \beta \geq 0) &= \Phi(\alpha \geq 1) \rightarrow \Phi(\beta \geq 0) \\ \Phi(\alpha \triangleright \beta \geq -1) &= \Phi(\alpha \geq 2) \rightarrow \Phi(\beta \geq -1) \end{aligned}$$

and for $i \in \{-2, -1, 0, 1\}$, it holds

$$\Phi(\alpha \leq i) = \neg\Phi(\alpha \geq i + 1)$$

Example 2. Consider the $N5$ -formula $\text{not}(a \triangleright b)$, whose intuition is rather difficult to grasp, when formulated in this way. Asserting that this formula is true can be expressed in GPL as follows:

$$\begin{aligned} \Phi(\text{not}(a \triangleright b) \geq 2) &= \Phi(a \triangleright b \leq 0) \\ &= \neg\Phi(a \triangleright b \geq 1) \\ &= \neg(\Phi(a \geq 1) \rightarrow \Phi(b \geq 1)) \\ &= \neg(\mathbf{N}_w(a) \rightarrow \mathbf{N}_w(b)) \end{aligned}$$

It is easy to see that the latter formula is equivalent to $\mathbf{N}_w(a) \wedge \mathbf{\Pi}_s(\neg b)$, whose intuition is clear: a is considered to be somewhat but not fully certain (i.e. a is true in all most plausible classical interpretations compatible with the GPL formula), whereas $\neg b$ is fully possible (i.e. epistemic states in which b is known to be true are excluded). Alternatively we can say that a is true there while b is not true there.

Every $N5$ formula α in an equilibrium logic base asserts that α is fully true and should be translated by the GPL formula $\Phi(\alpha \geq 2)$. Interestingly, due to its recursive definition starting from $N5$ literals, this translation is such that the GPL modalities only occur in front of literals. This observation essentially explains why $N5$ logic can be truth functional: as the equivalent of GPL formulas such as $\mathbf{N}_s(\alpha \vee \beta)$ cannot be expressed in $N5$, the truth value of $N5$ formulas only depends on what is supposed to hold for individual atoms. In contrast, to decide whether $\mathbf{N}_s(\alpha \vee \beta)$ holds in a given epistemic state π , it does not suffice to consider what we know about α and what we know about β in isolation. Hence the GPL translation of $N5$ will be in a specific fragment of GPL.

We now more closely investigate the precise link between a set of $N5$ formulas $\Theta = \{\alpha_1, \dots, \alpha_n\}$ and the GPL theory $T = \{\Phi(\alpha_1 \geq 2), \dots, \Phi(\alpha_n \geq 2)\}$, which intuitively asserts that the formulas in Θ are fully certain. First, for each $N5$ model V we can define a 3-valued possibility distribution

π_V such that (Lit is the set of literals)

$$\pi_V(\omega) = \begin{cases} 0 & \text{if } \exists l \in Lit. (\omega \models l) \wedge V(h, l) = -1 \\ 1/2 & \text{if } \exists l \in Lit. (\omega \models l) \\ & \wedge V(h, l) = 0 \wedge V(t, l) = -1 \\ 1 & \text{otherwise} \end{cases}$$

In terms of the valuation σ_V , it says

$$\pi_V(\omega) = \begin{cases} 0 & \text{if } \exists l \in Lit. (\omega \models l) \wedge \sigma_V(l) = -2 \\ 1/2 & \text{if } \exists l \in Lit. (\omega \models l) \wedge \sigma_V(l) = -1 \\ 1 & \text{otherwise} \end{cases}$$

In order to shed a model-theoretic light on the construction of π_V , let $E_h(V)$ and $E_t(V)$ be the Boolean epistemic states defined by the set of standard models of the propositional bases $B_h = \{a \in V_h\} \cup \{\neg a : \sim a \in V_h\}$ and $B_t = \{a \in V_t\} \cup \{\neg a : \sim a \in V_t\}$ respectively (they are known as partial models in partial logic (Blamey 1988)). Define the possibility distribution $\pi_{(h,t)}$ as follows:

$$\pi_{(h,t)}(\omega) = \begin{cases} 0 & \text{if } \omega \notin E_h(V) \\ 1/2 & \text{if } \omega \in E_h(V) \setminus E_t(V) \\ 1 & \text{if } \omega \in E_t(V). \end{cases} \quad (13)$$

Then it is easy to check that $\pi_{(h,t)} = \pi_V$. Indeed

- $\pi_{(h,t)}(\omega) = 0$ if and only if $\omega \notin E_h$, that is $\exists l \in B_h, \omega \not\models l$, i.e. $\omega \models \neg l$; so, letting $l' = \neg l$, $\omega \models l'$ and $\sigma(l') = -2$.
- $\pi_{(h,t)}(\omega) = 1/2$ if and only if $\omega \in E_h, \omega \notin E_t$, that is $\forall l \in B_h, \omega \models l$ and $\exists l \in B_t \setminus B_h, \omega \not\models l$, so $\exists l, \omega \models \neg l, l \notin B_h, l \in B_t$, hence $\sigma_V(l) = -1$.

Proposition 4. *Let α be a formula in N5-logic. For each N5 model V of α , the possibility distribution π_V is a model of $\Phi(\alpha \geq 2)$.*

Note that the obtained possibility distribution π_V has a special shape: its core and support are partial models.

Proposition 5. *Let α be a formula in N5-logic. For each model π of $\Phi(\alpha \geq 2)$, the N5 valuation V_π defined by*

$$V_\pi(h, a) = \begin{cases} 1 & \text{if } \pi \models \mathbf{N}_s(a) \\ -1 & \text{if } \pi \models \mathbf{N}_s(\neg a) \\ 0 & \text{otherwise} \end{cases}$$

$$V_\pi(t, a) = \begin{cases} 1 & \text{if } \pi \models \mathbf{N}_w(a) \\ -1 & \text{if } \pi \models \mathbf{N}_w(\neg a) \\ 0 & \text{otherwise} \end{cases}$$

is a model of α .

In the above proposition π is any possibility distribution, but it is not necessarily equal to π_{V_π} , the possibility distribution built from V_π following (13). In fact, we can prove that $\pi_{V_\pi} \geq \pi$ that is

Proposition 6. *Given an N5 model V the least specific possibility distribution π such that $\forall a \in At$,*

- $\pi \models \mathbf{N}_s(a)$ if $\sigma_V(a) = 2$.
- $\pi \models \mathbf{N}_s(\neg a)$ if $\sigma_V(a) = -2$.

- $\pi \models \mathbf{N}_w(a) \wedge \neg \mathbf{N}_s(a)$ if $\sigma_V(a) = 1$.
- $\pi \models \mathbf{N}_w(\neg a) \wedge \neg \mathbf{N}_s(\neg a)$ if $\sigma_V(a) = -1$.
- $\pi \models \neg \mathbf{N}_w(a) \wedge \neg \mathbf{N}_w(\neg a)$ if $\sigma_V(a) = 0$.

exists and is equal to π_V

In general, minimally specific models of $\Phi(\alpha \geq 2)$ are not unique. As it turns out, every such minimally specific model which is moreover Boolean corresponds to an equilibrium model of α .

Proposition 7. *Let α be a formula in N5-logic and let π be a minimally specific model of $\Phi(\alpha \geq 2)$. It holds that the N5 valuation V defined in Proposition 5 is an h-minimal model of α . If moreover π satisfies the constraint that $\pi(\omega) \neq 1/2$ for every interpretation ω , it holds that V is an equilibrium model.*

Despite this correspondence between N5 models and possibilistic models, and despite the correspondence between answer sets and minimally specific possibilistic models, the converse of this last property does not always hold, as is illustrated by the following counterexample.

Example 3. *Let $\alpha = (\text{not } a) \oplus a$. The N5 models are those valuations V that assign to a the values $(V(h, a), V(t, a))$ given by $(1, 1)$, $(0, 0)$, $(-1, -1)$ and $(0, -1)$. The valuations corresponding to $(1, 1)$, $(0, 0)$ and $(0, -1)$ are h-minimal. Hence α has two equilibrium models V_1 and V_2 , defined by $V_1(h, a) = V_1(t, a) = 1$ and $V_2(h, a) = V_2(t, a) = 0$.*

On the other hand, $\Phi(\alpha \geq 2) \equiv \Pi_s(\neg a) \vee \mathbf{N}_s(a)$. There are only two interpretations $\omega = a, \neg a$. The equilibrium models V_1 and V_2 correspond to the possibility distributions π_1 and π_2 , defined by $\pi_1(a) = 1$ and $\pi_1(\neg a) = 0$ and $\pi_2(a) = \pi_2(\neg a) = 1$. However, only π_2 is minimally specific.

One reason for this difference is that equilibrium logic is not systematically concerned with minimal specificity. Equilibrium models corresponds to the idea of a good guess. Minimizing V_h in (V_h, V_t) corresponds to a form of minimal specificity (maximally enlarging the support of π_V), but there is no request to minimize the set T of literals that are true ‘there’ as well among the total models. However, as we will show next, this counterexample does not affect theories that correspond to disjunctive answer set programs.

Relationship with answer set programming

Example 3 illustrates that some equilibrium models are not minimally specific possibility distributions, but the formula α on which the example is built does not have a clear intuition. We may wonder whether there might be counterexamples beyond such pathological cases. The following proposition suggests that the answer is negative: in any equilibrium logic theory which corresponds to a disjunctive answer set program, the equilibrium models are always minimally specific possibility distributions.

Proposition 8. *Let $\Theta = \{\alpha_1, \dots, \alpha_n\}$ be an N5 logic theory that corresponds to a disjunctive ASP program P . If V is an equilibrium model of Θ , then the possibility distribution π obtained from V using the procedure from Proposition 4 is a minimally specific model of $T = \{\Phi(\alpha_1 \geq 2), \dots, \Phi(\alpha_n \geq 2)\}$ and satisfies $\pi(\omega) \in \{0, 1\}$ for every possible world ω .*

Corollary 1. *Let Θ and T be as in Proposition 8. It holds that V has an equilibrium model iff T has a minimally specific model which is Boolean.*

Note that our translation of ASP is different from the translation in autoepistemic logic (Lifschitz and Schwarz 1993), as we do not need objective formulas. Moreover, it turns out that in the case of disjunctive answer set programs, the translation to GPL can be somewhat simplified. Indeed, a fact of the form (5) corresponds to the following GPL formula:

$$\begin{aligned}\Phi(l_1 \oplus \dots \oplus l_n \geq 2) &\equiv \Phi(l_1 \geq 2) \vee \dots \vee \Phi(l_n \geq 2) \\ &\equiv \mathbf{N}_s(l_1) \vee \dots \vee \mathbf{N}_s(l_n)\end{aligned}$$

whereas a rule of the form (4) corresponds to

$$\begin{aligned}\Phi(l_1 \otimes \dots \otimes \text{not } p_m \triangleright q_1 \oplus \dots \oplus q_n \geq 2) \\ \equiv (\Phi(l_1 \otimes \dots \otimes \text{not } p_m \geq 2) \rightarrow \Phi(q_1 \oplus \dots \oplus q_n \geq 2)) \\ \wedge (\Phi(l_1 \otimes \dots \otimes \text{not } p_m \geq 1) \rightarrow \Phi(q_1 \oplus \dots \oplus q_n \geq 1)) \\ \equiv (\mathbf{N}_s(l_1) \wedge \dots \wedge \mathbf{I}_s(\neg p_m) \rightarrow \mathbf{N}_s(q_1) \vee \dots \vee \mathbf{N}_s(q_n)) \\ \wedge (\mathbf{N}_w(l_1) \wedge \dots \wedge \mathbf{I}_s(\neg p_m) \rightarrow \mathbf{N}_w(q_1) \vee \dots \vee \mathbf{N}_w(q_n))\end{aligned}$$

In other words, a rule corresponds to the conjunction of two implications. However, as facts are modelled in terms of strong necessity, the second implication is irrelevant, i.e. from fully certain facts, we can derive fully certain conclusions.

Proposition 9. *Let Θ and T be as in Proposition 8, and let T' be the GPL theory that is obtained by converting every rule of the form (4) in Θ to a GPL formula $\mathbf{N}_s(l_1) \wedge \dots \wedge \mathbf{N}_s(l_n) \wedge \mathbf{I}_s(\neg p_1) \wedge \dots \wedge \mathbf{I}_s(\neg p_m) \rightarrow \mathbf{N}_s(q_1) \vee \dots \vee \mathbf{N}_s(q_n)$ and every fact of the form (5) to a GPL formula $\mathbf{N}_s(l_1) \vee \dots \vee \mathbf{N}_s(l_n)$. It holds that a Boolean possibility distribution π is a minimally specific model of T iff it is a minimally specific model of T' .*

This latter proposition reveals the connection between the approach from (Dubois, Prade, and Schockaert 2011) and our general translation of equilibrium logic theories. Moreover, the characterization of stable models as Boolean minimally specific possibility distributions is considerably simpler than the characterization proposed in (Dubois, Prade, and Schockaert 2011), where stable models were identified with maximally consistent sets of assumptions.

Computational complexity

Recall that Σ_2^P is the class of all problems that can be solved in polynomial time on a non-deterministic machine using an NP-oracle, while Π_2^P is the class of problems whose complement is in Σ_2^P . In disjunctive ASP, the problem of checking whether a program has at least one answer set containing a given literal l is Σ_2^P -complete, while the problem of checking whether all answer sets contain l is Π_2^P -complete (Eiter and Gottlob 1993). In GPL, the corresponding problems are as follows, for T and ϕ respectively a set of GPL formulas and a GPL formula:

SAT Check whether T has a minimally specific model π which is Boolean and for which $N(\phi) = 1$.

ENT Check whether for every least specific model π of T which is Boolean, it holds that $N(\phi) = 1$.

where in both cases N is the necessity measure induced by π . Because we know from Propositions 7 and 8 that the aforementioned decision problems in disjunctive ASP correspond to a special case of **SAT** and **ENT**, we know that these latter decision problems are Σ_2^P -hard and Π_2^P -hard respectively. We can moreover show the following result.

Proposition 10. *The complexity of **SAT** is in Σ_2^P and the complexity of **ENT** is in Π_2^P :*

In other words, we find that **SAT** is Σ_2^P -complete and **ENT** is Π_2^P -complete.

Concluding remarks

In this paper we have axiomatized the generalized possibilistic logic (GPL) and investigated its relation to equilibrium logic. In particular, we have shown how a set of $N5$ formulas Θ can be translated to a set of GPL formulas T such that the $N5$ models of Θ correspond to the possibilistic models of T . We have furthermore shown that the minimally specific models of T which are Boolean correspond to equilibrium models of Θ . However, some equilibrium models of Θ may not correspond to minimally specific models of T , i.e. the notion of h -minimality from equilibrium logic is weaker than the notion of minimal specificity from possibility theory. In spite of the fact that GPL offers more syntactic freedom, we have shown that the computational complexity remains the same as in equilibrium logic.

Since its introduction, equilibrium logic has mainly been popular because of three different reasons. First, it extends disjunctive ASP in the sense of Proposition 1. Second, an important characterization of strong equivalence⁴ has been obtained (Lifschitz, Pearce, and Valverde 2001): P and Q are strongly equivalent iff their sets of $N5$ models coincide. Third, due to the syntactic freedom, equilibrium logic can be used to give a declarative semantics to ASP extensions, such as aggregates. From the results presented in this paper, it follows that GPL shares the same advantages, while offering even more syntactic freedom. Indeed, while conceptually speaking, in equilibrium logic the modalities only occur in front of literals, in GPL they can occur in front of arbitrary propositional formulas. Moreover, as shown in Proposition 8, when restricted to the syntax of disjunctive ASP, the Boolean minimally specific models of a GPL theory correspond to the equilibrium models, and thus to the answer sets. Finally, strong equivalence can still be characterized in GPL as $N5$ models were shown to correspond to minimally specific possibilistic models in Propositions 4 and 5.

In those cases where equilibrium models do not correspond to minimally specific possibilistic models, it is not at all clear that equilibrium logic is closer to intuition. Indeed, one of the main disadvantages of equilibrium logic is that the intuitive meaning of many formulas is highly unclear, e.g. how should negation-as-failure in front of a rule

⁴Recall that two answer set programs P and Q are strongly equivalent, if for every program R it holds that $P \cup R$ and $Q \cup R$ have the same answer sets.

behave? By using a syntax in which modalities are written explicitly, GPL has the potential of being more intuitive to use, albeit with a slightly less compact syntax.

Proofs

Proof of Proposition 3

Using the proposed axioms, we can ensure that K only contains occurrences of $\wedge$ and $\vee$ at the meta-level. By distributivity of the connectives $\wedge$ and $\vee$, and because a set of formulas is equivalent to the conjunction of these formulas, we can moreover assume that K is of the form $\theta_1 \vee \dots \vee \theta_n$ where θ_i is a GPL formula in which $\wedge$ is the only connective occurring at the meta-level. It is clear that $K \models \phi$ is then equivalent with asserting that $\theta_i \models \phi$ for each i . From the axioms of propositional logic, we thus find that it is sufficient to show that ϕ can be derived from each θ_i . For the same reasons, we can assume that ϕ is of the form $\phi_1 \wedge \dots \wedge \phi_m$ where each ϕ_j is a GPL formula in which $\vee$ is the only connective occurring at the meta-level. What we then need to show is that $\theta_i \models \phi_j$ for every i and j .

We will treat θ_i as a set of formulas of the form $\mathbf{N}_\lambda(\alpha)$ and $\mathbf{\Pi}_\lambda(\alpha)$, with α a propositional formula. In the following, we let $\theta_i^N = \{\mathbf{N}_\lambda(\alpha) \mid \mathbf{N}_\lambda(\alpha) \in \theta_i\}$. From classical possibilistic logic, we know that θ_i^N has a unique minimally specific model π^* . It is not hard to see that either π^* is also the unique minimally specific model of θ_i or θ_i has no models. Indeed, for this epistemic model, and $\mathbf{\Pi}_\mu(\beta) \in \theta_i$, either $\mathbf{\Pi}^*(\beta) \geq \mu$ for all such formulas and π^* is an epistemic model of θ_i , or $\mathbf{\Pi}^*(\beta) < \mu$ for a specific $\mathbf{\Pi}_\mu(\beta) \in \theta_i$ and π^* is not an epistemic model of θ_i , nor will any more specific model π be so (since $\mathbf{\Pi}(\beta) < \mu$ will hold as well). Now we consider two cases:

1. First assume that π^* is the unique minimally specific model of θ_i . As axioms (8) and (9) imply that $\mathbf{\Pi}_\lambda(\alpha \vee \beta) \equiv \mathbf{\Pi}_\lambda(\alpha) \vee \mathbf{\Pi}_\lambda(\beta)$, and ϕ_j is a disjunction of $\mathbf{\Pi}_\lambda(\gamma)$'s and $\mathbf{N}_\lambda(\delta)$'s, it can be put in the form

$$\mathbf{\Pi}_{\lambda_1}(\gamma_1) \vee \dots \vee \mathbf{\Pi}_{\lambda_p}(\gamma_p) \vee \mathbf{N}_{\mu_1}(\delta_1) \vee \dots \vee \mathbf{N}_{\mu_q}(\delta_q)$$

where $\lambda_1 > \lambda_2 > \dots > \lambda_p$.

- If $\theta_i \models \mathbf{N}_{\mu_1}(\delta_1) \vee \dots \vee \mathbf{N}_{\mu_q}(\delta_q)$, then $\theta_i^N \models \mathbf{N}_{\mu_k}(\delta_k)$ for some k that is, $\mathbf{N}_{\mu_k}(\delta_k)$ can be derived from θ_i^N using the completeness of the inference rule for standard possibilistic logic, and the fact that this inference rule can be simulated using the proposed axioms and modus ponens. We can then deduce ϕ_j by weakening θ_i^N using the axioms of propositional logic.
- Suppose $\theta_i \not\models \mathbf{N}_{\mu_1}(\delta_1) \vee \dots \vee \mathbf{N}_{\mu_q}(\delta_q)$, but we have that $\theta_i \models \mathbf{N}_{\frac{1}{k}}(\gamma_1 \vee \dots \vee \gamma_p)$. Then we find that $\mathbf{N}_{\frac{1}{k}}(\gamma_1 \vee \dots \vee \gamma_p)$ can be derived from θ_i , and using (11) we can then also derive $\mathbf{\Pi}_1(\gamma_1 \vee \dots \vee \gamma_p)$, which means that $\mathbf{\Pi}_1(\gamma_1) \vee \dots \vee \mathbf{\Pi}_1(\gamma_s)$ can be derived. Now, $\mathbf{\Pi}_1(\gamma_l)$ is that same as $\neg \mathbf{N}_{\frac{1}{k}}(\neg \gamma_l)$. Using contraposition from classical logic and (12), this allows us to derive $\neg \mathbf{N}_{\nu(\lambda_l)}(\neg \gamma_l)$ from $\mathbf{\Pi}_1(\gamma_l)$, which is the same as $\mathbf{\Pi}_{\lambda_l}(\gamma_l)$. This means that we can also derive $\mathbf{\Pi}_{\lambda_1}(\gamma_1) \vee \dots \vee \mathbf{\Pi}_{\lambda_p}(\gamma_p)$.

- If, on the other hand, $\theta_i \not\models \mathbf{N}_{\frac{1}{k}}(\gamma_1 \vee \dots \vee \gamma_p)$, we have $\mathbf{N}^*(\gamma_1 \vee \dots \vee \gamma_p) = 0$ and thus $\mathbf{\Pi}^*(\neg \gamma_1 \wedge \dots \wedge \neg \gamma_p) = 1$. Since any refinement of a model of θ_i^N is still a model of θ_i^N , we can then always construct a model π' of θ_i^N in which $\mathbf{\Pi}'(\gamma_l) < \lambda_l$ for every l . Indeed, let π' be the interpretation defined by $\pi'(\omega) = \pi^*(\omega)$ if $\omega \models \neg \gamma_1 \wedge \dots \wedge \neg \gamma_p$ and $\pi'(\omega) = \min(\lambda_l - \frac{1}{k}, \pi^*(\omega))$ if $\omega \models \gamma_l \wedge \neg \gamma_{l+1} \wedge \dots \wedge \neg \gamma_p$. Note that because $\mathbf{\Pi}^*(\neg \gamma_1 \wedge \dots \wedge \neg \gamma_p) = \mathbf{\Pi}'(\neg \gamma_1 \wedge \dots \wedge \neg \gamma_p) = 1$, we are guaranteed that π' is still normalized. Hence, when $\theta_i \not\models \mathbf{N}_{\frac{1}{k}}(\gamma_1 \vee \dots \vee \gamma_p)$, we can construct a model π' of θ_i^N which is not a model of $\mathbf{\Pi}_{\lambda_1}(\gamma_1) \vee \dots \vee \mathbf{\Pi}_{\lambda_l}(\gamma_l)$. This means that either $\pi' \not\models \theta_i$ or $\pi' \models \mathbf{N}_{\mu_k}(\delta_k)$ for some k .
- Assume that $\pi' \not\models \theta_i$. As $\pi^* \models \theta_i$ implies $\pi' \models \theta_i^N$, $\pi' \not\models \theta_i$ means that there must be a formula $\mathbf{\Pi}_\mu(\beta)$ in $\theta_i \setminus \theta_i^N$ violated by π' , i.e. $\mathbf{\Pi}'(\beta) < \mu$, which implies that there exists an l such that $\mu = \lambda_l$. As only the possibility degree of γ_l -worlds may have changed from a value above μ to a value below μ , any β -world ω for which $\pi^*(\omega) \geq \mu$ is also a γ_l -world. So, $\mathbf{\Pi}(\beta \wedge \neg \gamma_l) < \lambda_l$. This means that $\theta_i^N \models \mathbf{N}_{\nu(\lambda_l)}(\beta \rightarrow \gamma_l)$. By induction, we find that $\mathbf{N}_{\nu(\lambda_l)}(\beta \rightarrow \gamma_l)$ is derivable. We must show that from $\mathbf{N}_{\nu(\lambda_l)}(\beta \rightarrow \gamma_l)$ and $\mathbf{\Pi}_{\lambda_l}(\beta)$, $\mathbf{\Pi}_{\lambda_l}(\gamma_l)$ can be derived. Actually this is an instance of the second GPL inference rule proved in Lemma 1, since the condition $\nu(\lambda_l) > 1 - \lambda_l$ is clearly satisfied.
- Assume $\pi' \models \mathbf{N}_{\mu_k}(\delta_k)$. By construction, π' is the least specific model of $\theta_i^N \cup \{\neg \mathbf{\Pi}_{\lambda_1}(\gamma_1), \dots, \neg \mathbf{\Pi}_{\lambda_s}(\gamma_p)\}$, which is equivalent to the standard possibilistic logic theory $\theta_i^N \cup \{\mathbf{N}_{\nu(\lambda_1)}(\neg \gamma_1), \dots, \mathbf{N}_{\nu(\lambda_p)}(\neg \gamma_p)\}$. Given the completeness of the inference rule of standard possibilistic logic w.r.t. what is true for the least specific model, and the fact that we can simulate this inference rule using the proposed axioms, we find that $\mathbf{N}_{\mu_k}(\delta_k)$ can be derived from $\theta_i^N \cup \{\neg \mathbf{\Pi}_{\lambda_1}(\gamma_1), \dots, \neg \mathbf{\Pi}_{\lambda_s}(\gamma_p)\}$, which means that $\neg \mathbf{\Pi}_{\lambda_1}(\gamma_1), \dots, \neg \mathbf{\Pi}_{\lambda_s}(\gamma_p) \rightarrow \mathbf{N}_{\mu_k}(\delta_k)$ can be derived from θ_i^N , which implies ϕ_j can.

2. Assume that θ_i has no models. This means that there is some formula of the form $\mathbf{\Pi}_\lambda(\gamma)$ in θ_i which is violated by π^* , and therefore by all models of θ_i^N . This means that $\theta_i^N \models \neg \mathbf{\Pi}_\lambda(\gamma)$, in other words $\theta_i^N \models \mathbf{N}_{\nu(\lambda)}(\neg \gamma)$. By syntactic inference in standard possibilistic logic, we can therefore infer $\mathbf{N}_{\nu(\lambda)}(\neg \gamma)$, and thus derive a logical inconsistency at the meta-level. Using the axioms from classical logic, this means that we can derive anything from θ_i , and in particular ϕ_j .

Proof of Proposition 4

We show that for every $N5$ -logic formula α , it holds that

$$\begin{aligned} (V(h, \alpha) = 1) &\Rightarrow (\pi_V \models \Phi(\alpha \geq 2)) \\ (V(t, \alpha) = 1) &\Rightarrow (\pi_V \models \Phi(\alpha \geq 1)) \\ (V(h, \alpha) = -1) &\Rightarrow (\pi_V \models \Phi(\alpha \leq -2)) \end{aligned}$$

$$\begin{aligned}
(V(t, \alpha) = -1) &\Rightarrow (\pi_V \models \Phi(\alpha \leq -1)) \\
(V(h, \alpha) = 0) &\Rightarrow (\pi_V \models \Phi(\alpha \geq -1) \wedge \Phi(\alpha \leq 1)) \\
(V(t, \alpha) = 0) &\Rightarrow (\pi_V \models \Phi(\alpha \geq 0) \wedge \Phi(\alpha \leq 0))
\end{aligned}$$

It is easy to see that the proposition then follows, as $V \models \alpha$ means that $V(h, \alpha) = 1$. The proof proceeds by structural induction:

- If $\alpha = a$ is an $N5$ literal, the above results are obvious by construction, since $\Phi(\alpha \geq 2) = \mathbf{N}_s(a)$, $\Phi(\sim\alpha \geq 2)$ is the same as $\Phi(\alpha \leq -2)$ and is $\mathbf{N}_s(\neg a)$, etc., while $V(h, \alpha) = 1$ is the same as $E_h \models a$, $V(h, \alpha) = -1$ is the same as $E_h \models \neg a$, etc.
- The cases where α is of the form $\alpha_1 \otimes \alpha_2$, $\alpha_1 \oplus \alpha_2$ or $\sim\alpha_1$ are straightforward.
- Assume that α is of the form *not* α_1 . If $V(h, \text{not } \alpha_1) = 1$, we have $V(t, \alpha_1) \in \{-1, 0\}$ by definition, which means by induction that $\pi_V \models \Phi(\alpha_1 \leq -1) \vee (\Phi(\alpha_1 \geq 0) \wedge \Phi(\alpha_1 \leq 0))$, which is equivalent to $\pi_V \models \Phi(\alpha_1 \leq 0)$, or equivalently $\pi_V \models \Phi(\text{not } \alpha_1 \geq 2)$. Note that the case where $V(t, \text{not } \alpha_1) = 0$ is impossible. The remaining cases are analogous.
- Assume that α is of the form $\alpha_1 \triangleright \alpha_2$. If $V(h, \alpha_1 \triangleright \alpha_2) = 1$ we have that $V(h, \alpha_1) \neq 1$ or $V(h, \alpha_2) = 1$, and that $V(t, \alpha_1) \neq 1$ or $V(t, \alpha_2) = 1$. By induction, we have from $V(h, \alpha_1) \neq 1$ that $\pi_V \models \Phi(\alpha_1 \leq -2)$ or $\pi_V \models \Phi(\alpha_1 \geq -1) \wedge \Phi(\alpha_1 \leq 1)$, and thus in any case that $\pi_V \models \Phi(\alpha_1 \leq 1)$. From $V(h, \alpha_2) = 1$ we find by induction that $\pi_V \models \Phi(\alpha_2 \geq 2)$. Hence, in any case we have $\pi_V \models \Phi(\alpha_1 \leq 1) \vee \Phi(\alpha_2 \geq 2)$, which is equivalent to $\pi_V \models \neg\Phi(\alpha_1 \geq 2) \vee \Phi(\alpha_2 \geq 2)$, which is the same as $\pi_V \models \Phi(\alpha_1 \geq 2) \rightarrow \Phi(\alpha_2 \geq 2)$. Similarly, we find from $V(t, \alpha_1) \neq 1$ or $V(t, \alpha_2) = 1$ that $\pi_V \models \Phi(\alpha_1 \geq 1) \rightarrow \Phi(\alpha_2 \geq 1)$. Together, by definition this means that $\pi_V \models \Phi(\alpha_1 \triangleright \alpha_2 \geq 2)$. The other cases are analogous.

Proof of Proposition 5

We show that for any formula α , it holds that

$$\begin{aligned}
V(h, \alpha) &= \begin{cases} 1 & \text{if } \pi \models \Phi(\alpha \geq 2) \\ -1 & \text{if } \pi \models \Phi(\alpha \leq -2) \\ 0 & \text{otherwise} \end{cases} \\
V(t, \alpha) &= \begin{cases} 1 & \text{if } \pi \models \Phi(\alpha \geq 1) \\ -1 & \text{if } \pi \models \Phi(\alpha \leq -1) \\ 0 & \text{otherwise} \end{cases}
\end{aligned}$$

from which the proposition readily follows. We proceed by structural induction:

- The case where α is an atom follows immediately from the definition of V .
- The cases where α is of the form $\alpha_1 \otimes \alpha_2$, $\alpha_1 \oplus \alpha_2$ or $\sim\alpha_1$ are straightforward.
- If α is of the form *not* α_1 and $\pi \models \Phi(\text{not } \alpha_1 \geq 2)$, we have $\pi \models \Phi(\alpha_1 \leq 0)$. This means that $\pi \not\models \Phi(\alpha_1 \geq 1)$, and thus by induction that $V(t, \alpha_1) \neq 1$, which means that $V(h, \text{not } \alpha_1) = 1$. The other cases are similar.

- If α is of the form $\alpha_1 \triangleright \alpha_2$ and $\pi \models \Phi(\alpha_1 \triangleright \alpha_2 \geq 2)$, we find $\pi \models (\Phi(\alpha_1 \geq 2) \rightarrow \Phi(\alpha_2 \geq 2)) \wedge (\Phi(\alpha_1 \geq 1) \rightarrow \Phi(\alpha_2 \geq 1))$. By induction, we easily find that then $V(h, \alpha_1) = 1 \Rightarrow V(h, \alpha_2) = 1$ and $V(t, \alpha_1) = 1 \Rightarrow V(t, \alpha_2) = 1$, which means that $V(h, \alpha_1 \triangleright \alpha_2) = 1$. Again the other cases are similar.

Proof of Proposition 7

We already know from Proposition 5 that V is a model of α . Now suppose that V were not h -minimal, i.e. that there exists a model V' such that $V_t = V'_t$ and $V'_h \subset V_h$. Let π' be the possibility distribution corresponding to V' , according to the construction from Proposition 4. Note that by Proposition 4 we have that π' is a model of $\Phi(\alpha \geq 2)$. We now consider two cases:

- Assume that $\pi \neq \pi'$, and let ω_0 be an interpretation for which $\pi(\omega_0) \neq \pi'(\omega_0)$. It is not hard to see from the procedure of Proposition 4 that then $\pi(\omega_0) = 0$ and $\pi'(\omega_0) = 1/2$, as V and V' only differ in world h . But this would mean that π were not a minimally specific model, a contradiction.
- Assume that $\pi = \pi'$. Let a be an atom such that $V'(h, a) = 0$ and $V(h, a) \in \{1, -1\}$, and assume for instance that $V(h, a) = 1$ (the case where $V(h, a) = -1$ is entirely analogous). Note that this entails that $V(t, a) = V'(t, a) = 1$. Moreover, because V has been obtained from π using the procedure of Proposition 5, $V(h, a) = 1$ means that $\pi \models \mathbf{N}_s(a)$. On the other hand, since π' was obtained from V' using the procedure of Proposition 4 and $V'(h, a) = 0$, we have $\pi' \models \Pi_w(a) \wedge \Pi_w(\neg a)$, again a contradiction.

Hence we have shown that V is h -minimal. Now assume that $\pi(\omega) \neq 1/2$ for every interpretation ω . Clearly, we then have that $\pi \models \mathbf{N}_s(l)$ for any literal iff $\pi \models \mathbf{N}_w(l)$ (as $\pi(\omega) > 0$ is the same as $\pi(\omega) = 1$ under the given assumption). By construction of V , we then immediately find that $V_h = V_t$.

Proof of Proposition 8

We first show that π is a minimally specific model.

- Assume that no negation as failure occurs in P , and suppose that π were not a minimally specific model of T . Let $\pi' \neq \pi$ be a model of T such that $\pi'(\omega) \geq \pi(\omega)$ for all interpretations ω . Then there exists an interpretation ω_0 such that $\pi(\omega_0) = 1/2$ and $\pi'(\omega_0) = 1$, or $\pi(\omega_0) = 0$ and $\pi'(\omega_0) \in \{1/2, 1\}$. The former case would mean that there is a literal l such that $\omega_0 \models l$, $V(h, l) = 0$ and $V(t, l) = -1$, which is not possible since $V(h, l) = V(t, l)$ in every equilibrium model. Thus, we would have $\pi(\omega_0) = 0$ and $\pi'(\omega_0) \in \{1/2, 1\}$. Then there is a literal l such that $\omega \models l$ and $V(h, l) = -1$. Let V' be the $N5$ valuation obtained from π' using the procedure from Proposition 5. Then we know that V' is also an $N5$ model of Θ from Proposition 5. Since $\pi'(\omega_0) \neq 0$, we have that $N(\neg l) = 1 - \Pi(l) < 1$, i.e. $\pi' \not\models \mathbf{N}_s(\neg l)$, and by construction we thus find $V'(h, l) = 0$. Hence V' differs from V in that there are some atoms a for which $V(h, a) \neq 0$ while $V'(h, a) = 0$, and possibly also $V'(t, a) = 0$.

Now let V'' be defined as $V''(h, a) = V'(h, a)$ and $V''(t, a) = V(t, a)$ for every atom a . It is easy to see that, because there are no occurrences of negation as failure or nested rules, from $V(t, \alpha_i) = 1$ and $V'(h, \alpha_i) = 1$, it follows that $V''(h, \alpha_i) = 1$. This contradicts the assumption that V were h -minimal.

- In general, let $A = \{l \mid V(t, l) = 1\}$ be the answer set of P corresponding with V . Then we know that A is also an answer set of P^A . If we write $\Theta^A = \{\alpha_i^A, \dots, \alpha_n^A\}$ for the $N5$ logic theory corresponding with P^A , we have that V is an equilibrium model of Θ^A , and by the previous point, that π is a minimally specific model of $T^A = \{\Phi(\alpha_1^A \geq 2), \dots, \Phi(\alpha_n^A \geq 2)\}$. Now we consider a sequence Θ_i of $N5$ logic theories to go from $\Theta_0 = \Theta^A$ to $\Theta_m = \Theta$ and we show that after each step, π remains a minimally specific model of the corresponding GPL theory. We write T_i for the GPL theory that corresponds to Θ_i . Each step in the sequence corresponds to one occurrence of *not* l_i in the original program. There are two cases to consider:

- If $l_i \notin A$, then we need to add a conjunct of the form *not* l_i to the antecedent of some rules in Θ_{i-1} . By construction, if $l_i \notin A$ then $V(t, l_i) = V(h, l_i) \neq 1$ and $\pi \models \Pi_s(\neg l_i)$. Then clearly π is still a model. The only models of T_i that are no models of T_{i-1} are models π' satisfying $\pi' \not\models \Pi_s(\neg l_i)$. However, such models π' cannot be more minimal than π , as $\Pi(\neg l_i) = 1 > \Pi'(\neg l_i)$.
- If $l_i \in A$, Θ_i will contain rules that do not correspond to any rule in Θ_{i-1} . However, the antecedent of these rules contains an occurrence of *not* l_i , hence all of these rules will be satisfied by V and thus by π . Moreover, as adding formulas cannot increase the set of possibilistic models, π is still a minimally specific model of T_i .

By repeatedly applying these steps we arrive at $T_m = T$, showing that π is a minimally specific model of T .

Finally, we need to show that $\pi(\omega) \neq 1/2$ for any interpretation ω . Note that $\pi(\omega) = 1/2$ would mean that there is a literal l such that $\omega \models l$ and $V(h, l) = 0$ and $V(t, l) = -1$. However, this is not possible since V is an equilibrium model, and thus $V(t, l) = V(h, l)$ for any literal.

Proof of Proposition 9

Since π is Boolean, we have that $\pi \models N_w(l_i)$ iff $N(l_i) \geq 1/2$ iff $N(l_i) = 1$. In other words, as soon as $\pi \models N_w(l_1) \wedge \dots \wedge \Pi_s(\neg p_m)$ we also have that $\pi \models N_s(l_1) \wedge \dots \wedge \Pi_s(\neg p_m)$. As a result, we easily find that $\pi \models T$ iff $\pi \models T'$. Moreover, it is easy to see that when π is a minimally specific model of T' , we also have that π is minimally specific model of T . We now also show the converse.

Suppose that π were a minimally specific model of T , while there were a model $\pi' \neq \pi$ of T' such that $\pi(\omega) \leq \pi'(\omega)$ for every world ω . Then clearly, π' cannot be a model of T , hence there must be a rule of the form (4) in Θ such that

$$\begin{aligned} \pi' &\models N_s(l_1) \wedge \dots \wedge \Pi_s(\neg p_m) \rightarrow N_s(q_1) \vee \dots \vee N_s(q_n) \\ \pi' &\not\models N_w(l_1) \wedge \dots \wedge \Pi_s(\neg p_m) \rightarrow N_w(q_1) \vee \dots \vee N_w(q_n) \end{aligned}$$

Since $N_s(q_1) \vee \dots \vee N_s(q_n) \models N_w(q_1) \vee \dots \vee N_w(q_n)$, this is only possible if $\pi' \models N_w(l_1) \wedge \dots \wedge \Pi_s(\neg p_m)$ while $\pi' \not\models N_s(l_1) \wedge \dots \wedge \Pi_s(\neg p_m)$. Now let π'' be the possibility distribution defined by $\pi''(\omega) = 1$ if $\pi'(\omega) \geq 1/2$ and $\pi''(\omega) = \pi'(\omega) = 0$ otherwise. To complete the proof, it suffices to show that $\pi'' \models T$, as this contradicts our assumption that π were a minimally specific model of T .

Clearly, it holds that $\pi' \models N_s(l_i)$ iff $\pi' \models N_s(l_1)$. Hence any formula in T which corresponds to a fact will still be satisfied by π'' . Regarding rules, it is useful to note that when $\pi' \models N_w(l_1) \wedge \dots \wedge \Pi_s(\neg p_m)$ while $\pi' \not\models N_s(l_1) \wedge \dots \wedge \Pi_s(\neg p_m)$ we will have that $\pi'' \not\models N_w(l_1) \wedge \dots \wedge \Pi_s(\neg p_m)$. Hence any rules that are satisfied by π but not by π' will be satisfied by π'' . If $\pi' \not\models N_w(l_1) \wedge \dots \wedge \Pi_s(\neg p_m)$ we also have $\pi'' \not\models N_w(l_1) \wedge \dots \wedge \Pi_s(\neg p_m)$, as π'' is less specific than π' , hence any rule which is trivially satisfied by π' will also be satisfied by π'' . Finally, if $\pi' \models N_s(l_1) \wedge \dots \wedge \Pi_s(\neg p_m)$, we also have $\pi' \models N_s(q_1) \vee \dots \vee N_s(q_n)$ as π' is a model of T' by construction. Then we will also have $\pi'' \models N_s(q_1) \vee \dots \vee N_s(q_n)$.

Proof of Proposition 10

Let us consider the following algorithm in Σ_2^P :

1. For each disjunction occurring in T at the meta-level, guess which disjunct should be satisfied. As a conjunction of formulas at the meta-level can be considered as a set of formulas, and negations at the meta-level can be avoided by “bringing them inside the modalities”, what we are left with is a set T_0 of formulas of the form $N_s(\alpha)$, $N_w(\alpha)$, $\Pi_s(\alpha)$ and $\Pi_w(\alpha)$. Let $K_s = \{\alpha \mid N_s(\alpha) \in T_0\}$, $K_w = \{\alpha \mid N_w(\alpha) \in T_0\}$, $C_s = \{\alpha \mid \Pi_s(\alpha) \in T_0\}$ and $C_w = \{\alpha \mid \Pi_w(\alpha) \in T_0\}$. Note that K_s , K_w , C_s and C_w are sets of propositional formulas.
2. For each $\alpha \in C_s$, check that $K_w \cup K_s \not\models \neg \alpha$. Furthermore, for each $\alpha \in C_w$, check that $K_s \not\models \neg \alpha$. If either of these checks fails, the algorithm returns *fail*. Moreover, if $K_w \cup K_s$ is inconsistent, always return *fail*.
3. Check that $K_s \models K_w$.
4. Return *accept* if $K_s \models \phi$ and *fail* otherwise.

Clearly, this algorithm is in Σ_2^P . To see why it correctly verifies whether there exists a minimally specific model π which is Boolean and for which $N(\phi) = 1$, it is useful to note that there is a one-to-one mapping between consistent disjunction-free theories of the form T_0 and the minimally specific models of T . It is not hard to show that the first two steps correspond to checking that T_0 corresponds to a unique minimally specific model π of T . Moreover, checking whether π is Boolean is nothing else than verifying $K_s \models K_w$, which corresponds to the third step. The final step then verifies whether $N(\phi) = 1$ for N the necessity measure induced by π .

The complexity of decision problem **ENT** is shown entirely analogously.

References

Banerjee, M., and Dubois, D. 2009. A simple modal logic for reasoning about revealed beliefs. In *Proceedings of*

- the European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty. 805–816.
- Bauters, K.; Schockaert, S.; De Cock, M.; and Vermeir, D. 2010. Possibilistic answer set programming revisited. In *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence*.
- Bauters, K.; Schockaert, S.; De Cock, M.; and Vermeir, D. 2011. Weak and strong disjunction in possibilistic asp. In *Proceedings of the Fifth International Conference on Scalable Uncertainty Management*. 475–488.
- Blamey, S. 1988. Partial logic. In Gabbay, D., and Guentner, F., eds., *Handbook of Philosophical Logic*, volume 5. Kluwer Academic Publishers. 261–353.
- Chagrov, A., and Zakharyashev, M. 1997. *Modal Logic*. Oxford University Press.
- Dubois, D., and Prade, H. 2007. Toward multiple-agent extensions of possibilistic logic. In *Proceedings of the IEEE International Conference on Fuzzy Systems*, 187–192.
- Dubois, D., and Prade, H. 2011. Generalized possibilistic logic. In *Proceedings of the Fifth International Conference on Scalable Uncertainty Management*. 428–432.
- Dubois, D.; Lang, J.; and Prade, H. 1994. Possibilistic logic. In D. Gabbay, C. Hogger J. Robinson, D. N., ed., *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3. Oxford University Press. 439–513.
- Dubois, D.; Prade, H.; and Schockaert, S. 2011. Rules and meta-rules in the framework of possibility theory and possibilistic logic. *Scientia Iranica* 18:566–573.
- Eiter, T., and Gottlob, G. 1993. Complexity results for disjunctive logic programming and application to nonmonotonic logics. In *Proceedings of the International Logic Programming Symposium*, 266–278.
- Gelfond, M., and Lifschitz, V. 1988. The stable model semantics for logic programming. In *Proceedings of the Fifth International Conference and Symposium on Logic Programming*, 1081–1086.
- Lifschitz, V., and Schwarz, G. 1993. Extended logic programs as autoepistemic theories. In *Proceedings of the Second International Workshop on Logic Programming and Non-Monotonic Reasoning*, 101–114.
- Lifschitz, V.; Pearce, D.; and Valverde, A. 2001. Strongly equivalent logic programs. *ACM Transactions on Computational Logic* 2(4):526–541.
- Lifschitz, V. 2008. Twelve definitions of a stable model. In *Proceedings of the 24th International Conference on Logic Programming*, 37–51. Berlin, Heidelberg: Springer-Verlag.
- Loyer, Y., and Straccia, U. 2006. Epistemic foundation of stable model semantics. *Theory and Practice of Logic Programming* 6:355–393.
- Moore, R. C. 1985. Semantical considerations on nonmonotonic logic. *Artificial Intelligence* 25(1):75 – 94.
- Nicolas, P.; Garcia, L.; Stéphan, I.; and Lefèvre, C. 2006. Possibilistic uncertainty handling for answer set programming. *Annals of Mathematics and Artificial Intelligence* 47(1-2):139–181.
- Pearce, D. 1997. A new logical characterisation of stable models and answer sets. In *Proceedings of the Second International Workshop on Non-Monotonic Extensions of Logic Programming*, 57–70.
- Pearce, D. 2006. Equilibrium logic. *Annals of Mathematics and Artificial Intelligence* 47:3–41.