

Online Research @ Cardiff

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/103267/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Dubois, Didier, Prade, Henri and Schockaert, Steven ORCID:
<https://orcid.org/0000-0002-9256-2881> 2017. Generalized possibilistic logic:
foundations and applications to qualitative reasoning about uncertainty.
Artificial Intelligence 252 , pp. 139-174. 10.1016/j.artint.2017.08.001 file

Publishers page: <https://doi.org/10.1016/j.artint.2017.08.001>
<<https://doi.org/10.1016/j.artint.2017.08.001>>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies.

See

<http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Generalized Possibilistic Logic: Foundations and Applications to Qualitative Reasoning about Uncertainty

Didier Dubois^a, Henri Prade^a, Steven Schockaert^b

^aToulouse University, Université Paul Sabatier, IRIT, CNRS, 118 Route de Narbonne, 31062 Toulouse Cedex 09, France

^bCardiff University, School of Computer Science & Informatics, 5 The Parade, Cardiff CF24 3AA, UK

Abstract

This paper introduces generalized possibilistic logic (GPL), a logic for epistemic reasoning based on possibility theory. Formulas in GPL correspond to propositional combinations of assertions such as “it is certain to degree λ that the propositional formula α is true”. As its name suggests, the logic generalizes possibilistic logic (PL), which at the syntactic level only allows conjunctions of the aforementioned type of assertions. At the semantic level, PL can only encode sets of epistemic states encompassed by a single least informed one, whereas GPL can encode any set of epistemic states. This feature makes GPL particularly suitable for reasoning about what an agent knows about the beliefs of another agent, e.g., allowing the former to draw conclusions about what the other agent does not know. We introduce an axiomatization for GPL and show its soundness and completeness w.r.t. possibilistic semantics. Subsequently, we highlight the usefulness of GPL as a powerful unifying framework for various knowledge representation formalisms. Among others, we show how comparative uncertainty and ignorance can be modelled in GPL. We also exhibit a close connection between GPL and various existing formalisms, including possibilistic logic with partially ordered formulas, a logic of conditional assertions in the style of Kraus, Lehmann and Magidor, answer set programming and a fragment of the logic of minimal belief and negation as failure. Finally, we analyse the computational complexity of reasoning in GPL, identifying decision problems at the first, second, third and fourth level of the polynomial hierarchy.

Keywords: Possibilistic logic, Epistemic reasoning, Non-monotonic reasoning

1. Introduction

Possibilistic logic [1] (PL) is a logic for reasoning with uncertain propositional formulas. Formulas in PL take the form (α, λ) where α is a propositional formula and λ is a certainty degree taken from the unit interval, or from another linear scale. Contrary to probabilistic logics, possibilistic logic models accepted beliefs in the sense that if two propositions are believed to a certain level, so is their conjunction. In many applications, a PL knowledge

Email addresses: dubois@irit.fr (Didier Dubois), prade@irit.fr (Henri Prade), schockaerts1@cardiff.ac.uk (Steven Schockaert)

base encodes the epistemic state of an agent. We then assume that all the agent knows are the formulas contained in the knowledge base and their logical consequences, with the weights referring to the degree of epistemic entrenchment [2] or the strength of belief. However, in its standard form, possibilistic logic has limitations as a tool for epistemic reasoning, i.e., reasoning about uncertainty, in at least two respects.

First, given that a knowledge base encodes a single epistemic state, PL does not allow us to encode incomplete information about the epistemic state of an agent. For example, assume that this agent privately flips a coin and looks at the result without revealing it. Then either the agent knows that the result was tails, which could be encoded as $\{(tails, 1)\}$, where 1 indicates complete certainty, or the agent knows that the result was heads, which could be encoded as $\{(\neg tails, 1)\}$. However, all an outside agent knows is that one of these two situations holds, and in particular this other agent knows that the first agent is not ignorant about the outcome of the coin flip. To express this situation in PL, we would need to write a disjunction $(tails, 1) \vee (\neg tails, 1)$ which is not allowed in the language. In this paper, we propose a generalized possibilistic logic (GPL) in which such disjunctions can be expressed. This brings PL syntax closer to the one of modal logics for epistemic reasoning, and, to emphasize this, we will use a slightly different notation and write $\mathbf{N}_1(tails) \vee \mathbf{N}_1(\neg tails)$ instead.

Second, PL does not allow us to explicitly encode information about the absence of knowledge. Instead, in practice, we must rely on a kind of closed-world assumption, i.e., assume that the agent does not know whether α is true if neither α nor its negation can be derived from the given knowledge base representing what is known about this agent's beliefs. When reasoning about beliefs as revealed by an agent, this assumption is hard to keep and we need to distinguish between situations where we (the outside agent) know that the agent is ignorant about α and situations where we do not know whether the agent knows α or not. In GPL, this can be achieved by putting a negation in front of PL formulas: $\neg\mathbf{N}_1(\alpha)$ expresses that we know that the agent does not believe in the truth of α ,¹ whereas situations where we have no such knowledge are encoded by GPL theories which have models in which $\mathbf{N}_1(\alpha)$ is true and models in which $\mathbf{N}_1(\alpha)$ is false.

GPL is closely related to modal logics for epistemic reasoning such as KD45 and S5. However, it is essentially a two-tiered propositional logic, and, instead of using Kripke frames, the semantics we propose for GPL is based on possibility distributions, which explicitly represent epistemic states. Our ability to directly interpret the modality \mathbf{N} as a constraint on a necessity measure results from the fact that we do not allow the modality \mathbf{N} to be nested. Furthermore, by not allowing objective formulas, we can naturally interpret each GPL formula as a constraint on the possible epistemic states (i.e., possibility distributions) of an agent. Compared to existing epistemic modal logics [3], we thus trade some expressiveness for a more intuitive way of capturing revealed beliefs. Among others, the use of possibility distributions has the advantage that (strength of) belief can be naturally encoded as a graded notion and that existing concepts from possibility theory such as minimal specificity and guaranteed possibility can be exploited to model ignorance in a natural way. This will

¹It means the agent either believes its negation or ignores the truth status of α .

enable us to encode various forms of non-monotonic reasoning in GPL. For instance, we will show how GPL can be used to model the semantics of answer set programming [4] (ASP) without relying on a fixpoint construction, unlike most existing characterizations of ASP, and how default rules in the sense of System P [5] can be modelled by taking advantage of the fact that GPL can express comparative uncertainty.

The paper is structured as follows. First, we recall some basic notions from possibility theory and possibilistic logic. In Section 3 we define the language of GPL and a corresponding semantics in terms of possibility distributions. We then provide an axiomatization which is sound and complete w.r.t. this latter semantics. In Section 4 we analyze how GPL can be used to reason about the ignorance of another agent, focusing on the role of minimal specificity and an extension to the language of GPL related to the notion of “only knowing” [6]. In Section 5 we then focus on the ability of GPL to model comparative uncertainty (e.g., α is more certain than β), showing how GPL can be used to encode a variant of possibilistic logic with partially ordered formulas [7], and how, as a result, a conditional logic based on System P [5] can be embedded in GPL. Subsequently, in Section 6 we explain in more detail how GPL relates to a number of existing formalisms for non-monotonic reasoning that are based on the notion of negation as failure. Section 7 discusses a number of computational issues, including the complexity of the main reasoning tasks. We also propose a reduction to SAT, allowing for a straightforward implementation of the reasoning tasks at the first level of the polynomial hierarchy. Finally, we present our conclusions.

This paper aggregates and significantly extends parts of [8] and [9]. In particular, in [8] we introduced the syntax, semantics and axiomatization of GPL, whereas in [9] we studied methods for modelling ignorance in GPL, introduced a new proof of the completeness of the axiomatization, and discussed some of the complexity results from Section 7. The results in Sections 5 and 6 are entirely new (although the encodings in Section 6 are similar in spirit to the encoding of equilibrium logic in [8]).

2. Preliminaries from possibility theory

Consider a variable X which has an unknown value from some finite universe \mathcal{U} . In possibility theory [10, 11, 12], available knowledge about the value of X is encoded as a mapping $\pi : \mathcal{U} \rightarrow [0, 1]$, which is called a possibility distribution. The intended interpretation of $\pi(u) = 1$ is that $X = u$ is fully compatible with all available information, while $\pi(u) = 0$ means that $X = u$ can be excluded based on available information. Note that the special case where we have no information about X is encoded using the vacuous possibility distribution, defined as $\pi(u) = 1$ for all $u \in \mathcal{U}$. Usually, we require that $\pi(u) = 1$ for some $u \in \mathcal{U}$, which corresponds to the assumption that the available information is consistent. If the possibility distribution π satisfies this condition, it is called normalized.

In general, the value of $\pi(u)$ can be interpreted in terms of degrees of potential surprise: the smaller the value of $\pi(u)$, the more we would be surprised to find out that $X = u$. This interpretation goes back to Shackle [13] and supports a purely qualitative interpretation of the possibility degrees $\pi(u)$. In such a case, we could replace the unit interval $[0, 1]$ by another linear scale (although an involutive order-reversing mapping is also needed). Other

interpretations of possibility degrees relate a possibility distribution to a family of probability distributions [14], to a family of likelihood functions [15], to Shafer belief functions [16], or to Spohn ordinal conditional functions [2, 17] and thus to infinitesimal probabilities [18], among others.

2.1. Set functions in possibility theory

A possibility distribution π induces a possibility measure Π , defined for $A \subseteq \mathcal{U}$ as [10]:

$$\Pi(A) = \max_{u \in A} \pi(u).$$

A dual measure N , called the necessity measure, is defined for $A \subseteq \mathcal{U}$ as [11]:

$$N(A) = 1 - \Pi(\mathcal{U} \setminus A) = \min_{u \notin A} (1 - \pi(u)).$$

Intuitively, $\Pi(A)$ reflects to what extent it is possible, given the available knowledge, that the value of X is among those in A , while $N(A)$ reflects to what extent the available knowledge entails that the value of X must necessarily be among those in A . Two other measures that can be introduced are the guaranteed possibility measure Δ and the potential necessity measure ∇ , defined for $A \subseteq \mathcal{U}$ as [12]:

$$\begin{aligned} \Delta(A) &= \min_{u \in A} \pi(u); \\ \nabla(A) &= 1 - \Delta(\mathcal{U} \setminus A) = \max_{u \in A} (1 - \pi(u)). \end{aligned}$$

Intuitively, $\Delta(A)$ reflects the extent to which all values in A are considered possible, while $\nabla(A)$ reflects the extent to which some value outside A is impossible. Note that for all $A \neq \emptyset$

$$\Delta(A) \leq \Pi(A); \quad N(A) \leq \nabla(A).$$

If π is normalized, we have $\Pi(A) = 1$ or $N(A) = 0$, and thus in particular:

$$N(A) \leq \Pi(A).$$

If $\pi(u) = 0$ for some $u \in \mathcal{U}$, we have $\Delta(A) = 0$ or $\nabla(A) = 1$, and thus:

$$\Delta(A) \leq \nabla(A).$$

Finally, note that Π and N are monotone w.r.t. set inclusion while Δ and ∇ are antitone, i.e., for $A \subseteq B$ we have

$$\Pi(A) \leq \Pi(B); \quad N(A) \leq N(B); \quad \Delta(A) \geq \Delta(B); \quad \nabla(A) \geq \nabla(B).$$

2.2. Possibilistic logic

A formula in propositional possibilistic logic [1] (PL for short) is an expression of the form (α, λ) , where $\lambda \in]0, 1]$ is a certainty degree and α is a propositional formula, built from a set of atomic formulas At using the connectives conjunction \wedge , negation \neg , disjunction \vee , implication \rightarrow , and equivalence \equiv in the usual way. Let Ω be the set of all interpretations of At and let \mathcal{L} be the set of all propositional formulas built from At . The semantics of possibilistic logic is defined in terms of possibility distributions over Ω . Specifically, a possibility distribution π over Ω satisfies the formula (α, λ) iff $N(\llbracket \alpha \rrbracket) \geq \lambda$, where $\llbracket \alpha \rrbracket$ denotes the set of all (classical) models of α . As π represents an epistemic state (it is a fuzzy set of classical models), we call it an epistemic model of (α, λ) , or an e-model for short. For the ease of presentation, we will write $N(\alpha)$ instead of $N(\llbracket \alpha \rrbracket)$ throughout this paper.

A possibility distribution π is an e-model of a set of PL formulas K iff π is an e-model of every formula in K . K generally has multiple e-models, but they can be partially ordered by the specificity ordering, whereby π_1 is less specific than π_2 , written $\pi_1 \preceq \pi_2$, if $\pi_1(\omega) \geq \pi_2(\omega)$ for every $\omega \in \Omega$. It can be shown that the set of e-models of a set of PL formulas K has a unique least element π_K w.r.t. \preceq , which is called the least specific e-model of K . It can be expressed, for all $\omega \in \Omega$ as [1]:

$$\pi_K(\omega) = 1 - \max\{\lambda \mid (\alpha, \lambda) \in K, \omega \not\models \alpha\}$$

where we assume $\max \emptyset = 0$. Intuitively, the more certain the formulas that are violated by ω , the less plausible ω is considered to be.

The following inference rules are valid in PL:

$$\text{if } (\alpha, \lambda) \in K \text{ then } K \vdash_{PL} (\alpha, \lambda) \tag{1}$$

$$\text{if } \vdash \alpha \text{ then } K \vdash_{PL} (\alpha, 1) \tag{2}$$

$$\text{if } \lambda_1 \geq \lambda_2 \text{ and } K \vdash_{PL} (\alpha, \lambda_1) \text{ then } K \vdash_{PL} (\alpha, \lambda_2) \tag{3}$$

$$\text{if } K \vdash_{PL} (\alpha \vee \beta, \lambda_1) \text{ and } K \vdash_{PL} (\neg \alpha \vee \gamma, \lambda_2) \text{ then } K \vdash_{PL} (\beta \vee \gamma, \min(\lambda_1, \lambda_2)) \tag{4}$$

Let us write $K \models_{PL} (\alpha, \lambda)$ if every e-model of K is an e-model of (α, λ) . If there is no cause for confusion we also write \models_{PL} as \models and \vdash_{PL} as \vdash . It is possible to show that the following statements are all equivalent for a set of PL formulas K (see e.g., [19]):

1. $K \vdash_{PL} (\alpha, \lambda)$ can be derived from (1)–(4).
2. $K \models_{PL} (\alpha, \lambda)$
3. The least specific e-model π_K of K is an e-model of (α, λ) .

Inference in possibilistic logic thus remains close to inference in propositional logic. In particular, let the c -cut K_c of K be the propositional theory $K_c = \{\alpha \mid (\alpha, \lambda) \in K \text{ and } \lambda \geq c\}$. Then we have that $K \models_{PL} (\alpha, \lambda)$ iff $K_\lambda \cup \{\neg \alpha\}$ is unsatisfiable. It follows that entailment checking in possibilistic logic is coNP-complete and that efficient reasoners can easily be implemented on top of off-the-shelf SAT solvers.

Possibilistic logic can be seen as a tool for specifying a ranking on propositional formulas. As such, it is closely related to the notion of epistemic entrenchment [20], as has been pointed

out in [2]. This makes PL a natural vehicle for implementing strategies for belief revision [21] and managing inconsistency [22]. Along similar lines, there are close connections between PL and default reasoning in the sense of System P [5], which can be exploited to implement several forms of reasoning about rules with exceptions [23].

Syntactically, propositional possibilistic logic is similar to the propositional fragment of Markov logic [24]. Semantically, however, the certainty weights in Markov logic are interpreted probabilistically. In particular, a set $M = \{(\alpha_1, w_1), \dots, (\alpha_n, w_n)\}$ of (propositional) Markov logic formulas defines the probability distribution p_M defined as follows ($\omega \in \Omega$):

$$p_M(\omega) = \frac{1}{Z} \exp \left(\sum_{i=1}^n \{w_i \mid \omega \models \alpha_i\} \right) \quad (5)$$

where Z is a normalization constant. This probabilistic semantics makes Markov logic particularly useful in machine learning settings. Note that we can equivalently define p_M as follows

$$p_M(\omega) = \frac{1}{Z'} \exp \left(\sum_{i=1}^n \{-w_i \mid \omega \not\models \alpha_i\} \right) \quad (6)$$

where the new normalization constant Z' is given by $Z' = \frac{Z}{\exp(\sum_i w_i)}$. This alternative formulation highlights the close relationship between the propositional fragment of Markov logic and the so-called penalty logic [25]. The two main differences are that negative weights are not considered in penalty logic² and that the penalty associated with an interpretation is not normalized. This lack of normalization makes penalty logic somewhat closer in spirit to possibilistic logic. Attaching a positive weight w to a formula α in penalty logic is similar to attaching a degree of necessity $1 - \exp(-w)$ to this formula in possibilistic logic. Thus the main difference between penalty logic and possibilistic logic is that in the former case the product is used to combine certainty degrees while in the latter case the minimum is used.³

However, we can also view Markov logic, penalty logic and possibilistic logic as equivalent frameworks for defining rankings of possible worlds. Indeed, as was shown in [27], given a Markov logic knowledge base M , we can always construct a possibilistic logic knowledge base K such that M and K define the same ranking of possible worlds, and vice versa. In fact, any ranking of interpretations can be represented by a possibilistic knowledge base.

3. Generalized possibilistic logic

While PL is useful to encode a single epistemic state, our aim is to develop GPL as a logic for reasoning about the epistemic state of an agent from its revealed beliefs. A GPL

²Note however that in Markov logic, we can replace (α, w) by $(\neg\alpha, -w)$ thanks to the use of the normalization constant, so allowing negative weights does not increase the expressivity of propositional Markov logic.

³Moreover, it is worth noticing that (6) defines the probability of an interpretation by using a possibility distribution which is renormalized by dividing each possibility degree by their sum. See [26] for a discussion of this type of possibility-probability transformation.

knowledge base then encodes the set of epistemic states that are compatible with these revealed beliefs. The aim of this section is to define the syntax and semantics of GPL, and to introduce an axiomatization for this logic. We will use α, β , etc. to denote propositions in standard propositional logic, formed with the connectives, \wedge and \neg . As usual, we will also use the abbreviations $\alpha \vee \beta = \neg(\neg\alpha \wedge \neg\beta)$, $\alpha \rightarrow \beta = \neg(\alpha \wedge \neg\beta)$ and $\alpha \equiv \beta = (\alpha \rightarrow \beta) \wedge (\beta \rightarrow \alpha)$. Let \mathcal{L} be the language of all propositional formulas over a finite set of atomic propositions At . Unless stated otherwise, we restrict the set of certainty degrees to the finite subset $\Lambda_k = \{0, \frac{1}{k}, \frac{2}{k}, \dots, 1\}$ of the unit interval, with $k \in \mathbb{N} \setminus \{0\}$ and let $\Lambda_k^+ = \Lambda_k \setminus \{0\}$.

3.1. Syntax

We define the language \mathcal{L}_{GPL}^k of generalized possibilistic logic with $k + 1$ certainty levels as follows:

- If $\alpha \in \mathcal{L}$ and $\lambda \in \Lambda_k^+$, then $\mathbf{N}_\lambda(\alpha) \in \mathcal{L}_{GPL}^k$.
- If $\Phi \in \mathcal{L}_{GPL}^k$ and $\Psi \in \mathcal{L}_{GPL}^k$, then $\neg\Phi$ and $\Phi \wedge \Psi$ are also in \mathcal{L}_{GPL}^k .

The corresponding logic will be referred to as GPL_k . When k is clear from the context we will also refer to this logic as GPL, and to the corresponding language as \mathcal{L}_{GPL} . Note that GPL is a graded version of the logic called MEL (Meta-Epistemic, or yet Minimal Epistemic, Logic), which was introduced in [28]. The MEL language is a special case of GPL where $k = 1$. Whereas MEL uses a standard modal logic syntax ($\square = \mathbf{N}_1$), we use a modality which refers to the necessity measure N to emphasize the link with possibility theory. Furthermore note that we view \mathcal{L}_{GPL}^k as a language with k different modalities $\mathbf{N}_{\frac{1}{k}}, \dots, \mathbf{N}_1$, rather than a language with a single modality and constants denoting certainty degrees.

In the following, we will also use the following abbreviation:

$$\mathbf{\Pi}_\lambda(\alpha) = \neg\mathbf{N}_{\nu(\lambda)}(\neg\alpha) \quad (7)$$

where we write $\nu(\lambda)$ as an abbreviation for $1 - \lambda + \frac{1}{k}$. Semantically the modality $\mathbf{\Pi}_\lambda$ will correspond to a lower bound on a possibility measure, namely (7) is the counterpart of the duality between a possibility and a necessity measure on a finite scale, where we have to shift from one level for moving from a strict inequality to an inequality in the broad sense.

Let us define a meta-atom as an expression of the form $\mathbf{N}_\lambda(\alpha)$, and a meta-literal as an expression of the form $\mathbf{N}_\lambda(\alpha)$ or $\neg\mathbf{N}_\lambda(\alpha)$. A meta-clause is an expression of the form $\Phi_1 \vee \dots \vee \Phi_n$ with each Φ_i a meta-literal. A meta-term is an expression of the form $\Phi_1 \wedge \dots \wedge \Phi_n$ with each Φ_i a meta-literal.

3.2. Semantics

The semantics of GPL are defined in terms of normalized possibility distributions over propositional interpretations, encoding epistemic states, where possibility degrees are, by duality, of the form $1 - \lambda, \forall \lambda \in \Lambda_k$.⁴ Let \mathcal{P}_k be the set of all such possibility distributions. An e-model of a GPL formula is any possibility distribution π from \mathcal{P}_k , namely:

⁴In our conventions, it comes down to using Λ_k as both a certainty and a possibility scale.

- π is an e-model of $\mathbf{N}_\lambda(\alpha)$ iff $N(\alpha) \geq \lambda$;
- π is an e-model of $\Phi_1 \wedge \Phi_2$ iff π is an e-model of Φ_1 and of Φ_2 ;
- π is an e-model of $\neg\Phi_1$ iff π is not an e-model of Φ_1 ;

where N is the necessity measure induced by π . As usual, π is called an e-model of a set of GPL formulas K , written $\pi \models_{GPL}^k K$, if it is an e-model of each formula in K . It is called a minimally specific e-model of K if there is no e-model $\pi' \neq \pi$ of K such that $\pi'(\omega) \geq \pi(\omega)$ for each possible world ω . We write $K \models_{GPL}^k \phi$, for K a set of GPL formulas and ϕ a GPL formula, if every e-model of K is also an e-model of ϕ . When k is clear from the context, we will sometimes write \models_{GPL}^k as \models_{GPL} ; furthermore, if there is no cause for confusion, we will also write \models_{GPL}^k as \models .

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 which suggests that α is true, and none that suggests that it is false. Note that we can distinguish between complete and partial certainty only if $k \geq 2$. 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 $\Pi(\neg\alpha) \geq 1 - \lambda + \frac{1}{k}$. The abbreviation introduced in (7) thus corresponds to a syntactic counterpart of the duality between necessity and possibility measures. Note how the use of a finite scale makes it possible to express strict inequalities, even though we only use inequalities in the wide sense in the interpretation of graded modalities. Intuitively $\mathbf{\Pi}_1(\alpha)$ means that α is fully compatible with our available beliefs (i.e., nothing prevents α from being true), while $\mathbf{\Pi}_\lambda(\alpha)$ with $\lambda < 1$ means that α cannot be fully excluded ($\Pi(\alpha) \geq \lambda$).

This formalism is similar to an autoepistemic logic [29, 6]. However the latter aims to capture how an agent reasons about its own beliefs. One crucial difference, which has been pointed out in [30], is that when reasoning about one's own beliefs, it should not be possible to state $\mathbf{N}_1(\alpha) \vee \mathbf{N}_1(\beta)$ without either stating $\mathbf{N}_1(\alpha)$ or $\mathbf{N}_1(\beta)$. Indeed, if we accept that an agent is aware of its epistemic state, the agent can tell, for each propositional formula, whether or not it is believed. Accordingly, in standard possibilistic logic, we cannot encode $\mathbf{N}_1(\alpha) \vee \mathbf{N}_1(\beta)$. We can just encode $\mathbf{N}_1(\alpha)$ or $\mathbf{N}_1(\beta)$, or their conjunction. However, we will be able to overcome this limitation in GPL. More generally, in a graded setting, if the agent is aware of its epistemic state, it can tell which of two propositional formulas it considers to be most certain. This is again in accordance with possibilistic logic, whereas in GPL we will be able to encode the case where we are ignorant about which of two formulas is most certain for an external agent. This suggests that while standard possibilistic logic offers a natural setting for reasoning with one's own beliefs, GPL naturally lends itself to reasoning about another agent's beliefs. For this reason, we could say that GPL is an "alter-epistemic" logic.

As to the possible kinds of conclusions that can be inferred from a GPL base K regarding a propositional formula α , if $k = 2$, one can distinguish between the following five cases:

- $K \models \mathbf{N}_1(\alpha)$ means that we know that the agent knows that α is true.
- $K \models \mathbf{N}_1(\neg\alpha)$ means that we know that the agent knows that α is false.

- $K \models \mathbf{N}_1(\alpha) \vee \mathbf{N}_1(\neg\alpha)$, $K \not\models \mathbf{N}_1(\alpha)$ and $K \not\models \mathbf{N}_1(\neg\alpha)$ means that we know that the agent knows whether α is true or false, but we do not know which it is.
- $K \models \mathbf{\Pi}_1(\alpha) \wedge \mathbf{\Pi}_1(\neg\alpha)$ means that we know that the agent is ignorant about whether α is true or false.
- $K \not\models \mathbf{N}_1(\alpha) \vee \mathbf{N}_1(\neg\alpha)$ and $K \not\models \mathbf{\Pi}_1(\alpha) \wedge \mathbf{\Pi}_1(\neg\alpha)$ means that we are ignorant about whether the agent is ignorant about α .

This is in contrast with the only three situations that can be distinguished in classical logic (and in PL), i.e., we know that α is true, we know that α is false, or we do not know whether α is true or false. When $k > 2$, we can consider graded counterparts of the five aforementioned cases. Moreover, a GPL base can then also express comparative uncertainty. For example:

- $K \models \bigvee_{i=1}^k \mathbf{N}_{\frac{i}{k}}(\alpha) \wedge \neg\mathbf{N}_{\frac{i}{k}}(\beta)$: we know that the agent is more certain that α holds than that β holds, noticing that it is equivalent to $\exists i, N(\alpha) \geq \frac{i}{k} > N(\beta)$.
- $K \models \bigvee_{i=1}^k \mathbf{\Pi}_{\frac{i}{k}}(\alpha) \wedge \neg\mathbf{\Pi}_{\frac{i}{k}}(\beta)$: we know that the agent would be less surprised to learn that α is true than to learn that β is true, noticing that it is equivalent to $\exists i, \Pi(\alpha) \geq \frac{i}{k} > \Pi(\beta)$.
- $K \models \bigvee_{i=1}^k (\mathbf{N}_{\frac{i}{k}}(\alpha) \vee \mathbf{N}_{\frac{i}{k}}(\neg\alpha)) \wedge \neg\mathbf{N}_{\frac{i}{k}}(\beta) \wedge \neg\mathbf{N}_{\frac{i}{k}}(\neg\beta)$: we know that the agent is more certain about the truth or the falsity of α than about β , but we may not know with which certainty degree the agent knows the truth value of α , nor to what extent this certainty degree is greater than the certainty degree about the truth or the falsity of β .
- $K \models \bigvee_{i=1}^k (\mathbf{N}_{\frac{i}{k}}(\alpha) \wedge \neg\mathbf{N}_{\frac{i}{k}}(\beta)) \vee (\mathbf{N}_{\frac{i}{k}}(\beta) \wedge \neg\mathbf{N}_{\frac{i}{k}}(\alpha))$: we know that the agent considers one of α, β more certain than the other, but we may not know which.
- $K \models \bigwedge_{i=1}^k (\mathbf{N}_{\frac{i}{k}}(\alpha) \rightarrow \mathbf{N}_{\frac{i}{k}}(\beta))$ expresses that the agent is at least as certain about β as about α .

Example 1. *The six nations championship is a rugby competition consisting of 5 rounds. In each round, every team plays against one of the other 5 teams, so that over 5 rounds all teams have played once against each other. Let us write $plays_i(x, y)$ to denote that x and y have played against each other in round i , and $won_i(x)$ to denote that team x has won its game in round i . Let $T = \{eng, fra, ire, ita, sco, wal\}$. To express that an agent knows the rules of the championship, we can consider formulas such as, among others:*

$$\mathbf{N}_1(\bigvee \{plays_i(x, u) \mid u \neq x, u \in T\}) \quad (8)$$

where $x \in T$. A formula such as $\mathbf{N}_{\frac{3}{4}}(won_1(wal))$ means that the agent strongly believes, but is not fully certain, that Wales (wal) has won its first round game, while $\mathbf{\Pi}_{\frac{3}{4}}(won_1(wal))$

means that the agent does not exclude that Wales has won its first round game, without evidence as to the contrary. The following formula expresses that the agent considers it more plausible that Wales has won its first game than that England (eng) has won its first game

$$\bigvee_{i=1}^k \Pi_{\frac{i}{k}}(won_1(wal)) \wedge \neg \Pi_{\frac{i}{k}}(won_1(eng)) \quad (9)$$

Recall that the certainty degrees in GPL are typically only assumed to have an ordinal meaning. Saying that the necessity of a formula is $\frac{3}{4}$ then does not have any intrinsic meaning, other than the fact that this formula is considered e.g., more certain than a formula with necessity $\frac{1}{2}$ and less certain than a formula with necessity $\frac{7}{8}$. The above example illustrates two alternative ways in which applications can deal with such ordinal certainty degrees. One idea is to use a small number of categories that are meaningful to a user, such as e.g., ‘completely certain’, ‘very certain’, ‘quite certain’, ‘somewhat certain’, and map these categories to the available elements from Λ_k (e.g., ‘very certain’ could correspond to a necessity of $\frac{3}{4}$). The second idea would be to avoid assigning certainty degrees, and only express certainty in a comparative way, as is illustrated in (9). This second approach will be discussed in more detail in Section 5.

3.3. Axiomatization

We consider the following axiomatization, which closely parallels the one of MEL [28]:

(PL) The axioms of classical logic for meta-formulas.

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

(N) $\mathbf{N}_1(\alpha)$ whenever $\alpha \in \mathcal{L}$ is a classical tautology.

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

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

If Φ can be derived from a set of GPL formulas K using the axioms **(PL)**, **(K)**, **(N)**, **(D)**, **(W)** and modus ponens, we write $K \vdash_{GPL} \Phi$; if there is no cause for confusion we also write $K \vdash \Phi$. Note in particular that when λ is fixed we get a fragment of the modal logic KD. In particular, the axioms entail that $\mathbf{N}_\lambda(\alpha \wedge \beta)$ is equivalent to $\mathbf{N}_\lambda(\alpha) \wedge \mathbf{N}_\lambda(\beta)$. It is easy to see that if α and β are logically equivalent formulas, then $\mathbf{N}_\lambda(\alpha)$ and $\mathbf{N}_\lambda(\beta)$ are also equivalent. Indeed, in that case, $(\alpha \rightarrow \beta) \wedge (\beta \rightarrow \alpha)$ holds, and by applying **(N)**, **(W)**, **(K)**, **(D)** we get both $\mathbf{N}_\lambda(\alpha) \rightarrow \mathbf{N}_\lambda(\beta)$ and $\mathbf{N}_\lambda(\beta) \rightarrow \mathbf{N}_\lambda(\alpha)$. Also note that from **(N)** and **(W)** we can derive a graded version of the necessitation rule, i.e., if $\vdash \alpha$ then $\vdash_{GPL} \mathbf{N}_\lambda(\alpha)$ for any $\lambda \in \Lambda_k$. Finally note that in the case where $k = 1$, GPL coincides with the logic MEL. In this latter case, we have $\Pi_1(\alpha) = \neg \mathbf{N}_1(\neg \alpha)$ whereas in general we only have $\Pi_1(\alpha) = \neg \mathbf{N}_{\frac{1}{k}}(\neg \alpha)$. As we will see in Section 6, the ability to differentiate between full possibility for α and the lack of full certainty for $\neg \alpha$ is crucial when using GPL to provide a semantics for negation as failure.

Proposition 1 (Soundness and completeness). *Let K be a set of GPL formulas and Φ a GPL formula. It holds that $K \models_{GPL} \Phi$ iff $K \vdash_{GPL} \Phi$.*

Proof. The proof is presented in Appendix A. □

The main idea behind the proof is that we can see formulas in GPL as propositional formulas which are built from the set of atomic formulas of \mathcal{L}_{GPL}^k . Given a knowledge base K in GPL, we construct a propositional base K^* made of formulas of K plus axioms of GPL, viewed as propositional formulas as well. We then show that there exists a bijection between the set of propositional models of K^* (seen as a propositional logic knowledge base) and the set of e-models of K (seen as a GPL knowledge base). A very similar strategy has been used, among others, in [44], [82] and [45, 46], in the context of multi-valued modal logics for reasoning about necessity (see Section 3.4).

Proposition 1 remains valid even if the set At of atomic propositions is countably infinite. On the other hand, the completeness result no longer holds if infinitely many certainty degrees are allowed in the language, as e.g. $\{\mathbf{N}_\lambda(a) \mid \lambda < \frac{1}{2}\} \models_{GPL} \mathbf{N}_{\frac{1}{2}}(a)$, for $a \in At$ but $\{\mathbf{N}_\lambda(a) \mid \lambda < \frac{1}{2}\} \not\models_{GPL} \mathbf{N}_{\frac{1}{2}}(a)$. This is not a real restriction, since knowledge bases only have finitely many formulas in practice, which means that only finitely many certainty levels actually need to be used, and since the semantics of GPL is based on the relative ordering of the certainty degrees, we can then always map these certainty degrees to Λ_k for some k . In Section 5, however, we will discuss an extension of GPL in which we can express comparative uncertainty statements, where it will be desirable to allow an unbounded number of certainty degrees at the semantic level.

Using Proposition 1, and some well-known properties on necessity and possibility measures, it follows that the following formulas are theorems in GPL:

$$\begin{aligned} \mathbf{N}_\lambda(\alpha) \wedge \mathbf{N}_\lambda(\beta) &\equiv \mathbf{N}_\lambda(\alpha \wedge \beta) \\ \mathbf{\Pi}_\lambda(\alpha \wedge \beta) &\rightarrow \mathbf{\Pi}_\lambda(\alpha) \wedge \mathbf{\Pi}_\lambda(\beta) \\ \mathbf{N}_\lambda(\alpha) \vee \mathbf{N}_\lambda(\beta) &\rightarrow \mathbf{N}_\lambda(\alpha \vee \beta) \\ \mathbf{\Pi}_\lambda(\alpha) \vee \mathbf{\Pi}_\lambda(\beta) &\equiv \mathbf{\Pi}_\lambda(\alpha \vee \beta) \end{aligned}$$

Next is a counterpart to the modus ponens rule in PL (4):

$$\mathbf{N}_{\lambda_1}(\alpha) \wedge \mathbf{N}_{\lambda_2}(\alpha \rightarrow \beta) \rightarrow \mathbf{N}_{\min(\lambda_1, \lambda_2)}(\beta) \tag{10}$$

To show that this is a theorem in GPL, thanks to Proposition 1, it suffices to note that every necessity measure N satisfying $N(\alpha) \geq \lambda_1$ and $N(\neg\alpha \vee \beta) \geq \lambda_2$ also satisfies $N(\beta) \geq \min(\lambda_1, \lambda_2)$, which is equivalent to the usual modus ponens in PL, a special case of (4). To see how (10) can be derived from the axioms of GPL, note that the deduction theorem is valid in GPL, and it thus suffices to show that $\mathbf{N}_{\min(\lambda_1, \lambda_2)}(\beta)$ can be derived from $\{\mathbf{N}_{\lambda_1}(\alpha), \mathbf{N}_{\lambda_2}(\alpha \rightarrow \beta)\}$. Starting from this latter set of premises, we apply **(W)** to obtain $\mathbf{N}_{\min(\lambda_1, \lambda_2)}(\alpha)$ and $\mathbf{N}_{\min(\lambda_1, \lambda_2)}(\alpha \rightarrow \beta)$. Applying modus ponens on axiom **(K)** and $\mathbf{N}_{\min(\lambda_1, \lambda_2)}(\alpha \rightarrow \beta)$, we obtain $\mathbf{N}_{\min(\lambda_1, \lambda_2)}(\alpha) \rightarrow \mathbf{N}_{\min(\lambda_1, \lambda_2)}(\beta)$. Using modus ponens on the latter formula and $\mathbf{N}_{\min(\lambda_1, \lambda_2)}(\alpha)$ we obtain $\mathbf{N}_{\min(\lambda_1, \lambda_2)}(\beta)$.

The following theorem is the counterpart of a hybrid modus ponens rule introduced in [31]:

$$\mathbf{\Pi}_{\lambda_1}(\alpha) \wedge \mathbf{N}_{\lambda_2}(\alpha \rightarrow \beta) \rightarrow \mathbf{\Pi}_{\lambda_1}(\beta), \text{ if } \lambda_2 > 1 - \lambda_1 \quad (11)$$

Again a direct proof can be given, using the deduction theorem, 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 $1 - \lambda_1 + \frac{1}{k} \leq \lambda_2$, i.e., $\lambda_2 > 1 - \lambda_1$ ⁵.

Resolution rules in possibilistic logic [31], extending (10) and (11), can be proved likewise in GPL or, alternatively, by using the decomposability of $\mathbf{N}_\lambda(\cdot)$ w.r.t. conjunction.

3.4. Related work

Although possibility theory has been the basis of an original theory of approximate reasoning [32], it was not introduced as a logical setting for epistemic reasoning, strictly speaking. Nonetheless, in the setting of his representation language PRUF [33], Zadeh discusses the representation of statements of the form “ X is A ” (meaning that the possible values of the single-valued variable X are fuzzily restricted by fuzzy set A), linguistically qualified in terms of truth, probability, or possibility. Interestingly, the representation of possibility-qualified statements led to possibility distributions over possibility distributions, but certainty-qualified statements, first considered in [34] (see also [11]), and used as the basic building blocks of possibilistic logic, were not considered at all, just because necessity measures as the dual of possibility measures were playing almost no role in Zadeh’s view (with the exception of half a page in [35]). Possibility-qualified statements were exploited in [31] in relation with a weighted resolution principle extending the inference rule (11), whose formal analogy with an inference rule existing in modal logic was stressed.

The similarity between possibility theory (including necessity measures) and modal logic should not come as a surprise since the analogy between the duality property $N(A) = 1 - \Pi(\Omega \setminus A)$ in possibility theory and the definition of $\Diamond p$ as $\neg\Box\neg p$ is striking, and has been known for a long time [36]. Likewise, the axiom $\Box p \rightarrow \Diamond p$ (axiom D in modal logic systems) may encode the inequality $N(A) \leq \Pi(A)$, and the characteristic axiom of necessity measures $N(A \cap B) = \min(N(A), N(B))$ corresponds to the theorem $(\Box p \wedge \Box q) \leftrightarrow \Box(p \wedge q)$ which is valid in modal system K. Nevertheless, no formally established connection between modal logic and possibility theory existed until the late 1980s.

This striking parallel between possibility theory and modal logic eventually led to proposals for a modal analysis and encoding of possibility theory. For instance, L. Fariñas and A. Herzig [37] proposed such an encoding by heavily relying on Lewis’ conditional logics of comparative possibility [38], as indeed the only numerical counterparts of Lewis possibility relations are possibility measures [39]. Another attempt was later made by Boutilier [40], in the scope of non-monotonic inference based on a plausibility relation over possible worlds. The idea was to use this ordinal counterpart of a possibility distribution as an accessibility

⁵If $\nu(\lambda_1) > \lambda_2$, the weakening axiom (**W**) leads us to derive $\mathbf{N}_{\lambda_2}(\neg\alpha)$, whose negation is weaker than the premise $\mathbf{\Pi}_{\lambda_1}(\alpha)$.

relation and to construct modalities from it. Another, more semantically-oriented trend was to build specific accessibility relations agreeing with possibility theory [41, 42].

A major difference with GPL is that the semantics of the above logics relies on accessibility relations. GPL can be embedded into a multimodal logic, but it is actually just a two-level propositional logic since its semantics is based on graded epistemic states, viewed as higher-order interpretations, not relying on accessibility relations. This point was discussed in [43]: relational semantics of epistemic logics may make sense in the scope of introspective reasoning, but appears more difficult to justify for modelling partial knowledge about the epistemic state of an external agent. In GPL, any agent is supposed to be aware of its own epistemic state, so it can model its own beliefs using a complete GPL base (see Section 4 on this point). Also, formally, GPL is a complexification of propositional logic, adding weighted modalities in front of propositional formulas only, and, at the semantic level, moving from usual interpretations to sets thereof, while simple epistemic logics like S5 or KD45 are constructed as a simplification of a complex logic allowing nested modalities naturally interpreted via accessibility relations, and need introspection axioms to simplify complex formulas into equivalent ones of depth at most 1. So beyond the formal analogies between modal logic and GPL, the motivations and the construction method are radically different.

A proposal closer to GPL is the one of Hájek [44], where possibility theory is cast into a many-valued logic setting, using many-valued modal formulas. The main difference with GPL, from a formal point of view, is that necessity is expressed as a single multi-valued modality, rather than a set of classical modalities in GPL. This implies that necessity statements need to be combined using a fuzzy logic, rather than classical propositional logic in GPL. A number of related logics are studied in [45, 46], which are using variants of Łukasiewicz logic both for the formulas inside the modalities and for combining the multi-valued modalities. In case these variants of Łukasiewicz are finite-valued (or e.g., include the Baaz Δ connective [47]), it is easy to see that GPL can be framed as a fragment of such a multi-valued modal logic. A general completeness result for such two-tiered (multi-valued) model logics has been introduced in [48]. Liao and Lin [49] have also studied a modal logic which is very similar to GPL, albeit using $[0, 1]$ as a possibility scale (which forces them to introduce additional multimodal formulas to deal with strict inequalities). Their tableau-based proof methods could be of interest to develop inference techniques for GPL.

While from a formal point of view, GPL is close to some of these aforementioned logics, our focus in this paper is rather different. Specifically, our main aim is to study what is gained, in terms of the kinds of epistemic reasoning scenarios that can be modelled, from the increase in syntactic freedom compared to standard possibilistic logic. Among others, we will analyse several ways in which partial ignorance can be modelled, study the relation between GPL and logics of comparative uncertainty, and show how different forms of non-monotonic reasoning can naturally be modelled using GPL. To the best of our knowledge, these links with possibilistic logic (or the related multi-valued modal logics) have not been studied in previous work.

4. Reasoning about ignorance in GPL

Possibility theory offers a number of tools for modelling limitations on what is known. These tools can be used in GPL to explicitly model what we know that an external agent does not know. In particular, Section 4.1 proposes a method based on the guaranteed possibility measure, which is subsequently refined in Section 4.2. In Section 4.3, we then analyse how the principle of minimal specificity can be applied to reason about what an external agent does not know.

4.1. Ignorance as guaranteed possibility

Using the modalities \mathbf{N} and $\mathbf{\Pi}$ we can model constraints of the form $N(\alpha) \geq \lambda$, $N(\alpha) \leq \lambda$, $\Pi(\alpha) \geq \lambda$ and $\Pi(\alpha) \leq \lambda$. So far, however, we have not considered the guaranteed possibility measure Δ and potential necessity measure ∇ . Counterparts of these measures can be introduced as abbreviations in the language, by noting that $\Delta(\alpha) = \min_{\omega \in \llbracket \alpha \rrbracket} \Pi(\{\omega\})$. For a propositional interpretation ω let us write $conj_\omega$ for the conjunction of all literals made true by ω , i.e., $conj_\omega = \bigwedge_{\omega \models a} a \wedge \bigwedge_{\omega \models \neg a} \neg a$. Then we define:

$$\Delta_\lambda(\alpha) = \bigwedge_{\omega \in \llbracket \alpha \rrbracket} \Pi_\lambda(conj_\omega) \qquad \nabla_\lambda(\alpha) = \neg \Delta_{\nu(\lambda)}(\neg \alpha) \quad (12)$$

In fact, since $\Pi(\alpha) = \max_{\omega \in \llbracket \alpha \rrbracket} \Delta(\{\omega\})$, another strategy we could have taken is to axiomatize a logic based on guaranteed possibility, and to define the modality \mathbf{N} as an abbreviation. In particular, such a logic could be axiomatized by using the following graded version of the data logic of Dubois, Hájek and Prade [50]:

(PL) The axioms of classical logic for meta-formulas.

(K $^\Delta$) $\Delta_\lambda(\alpha \wedge \neg \beta) \rightarrow (\Delta_\lambda(\neg \alpha) \rightarrow \Delta_\lambda(\neg \beta))$.

(Δ) $\Delta_1(\alpha)$ whenever $\neg \alpha \in \mathcal{L}$ is a tautology.

(D $^\Delta$) $\Delta_\lambda(\alpha) \rightarrow \nabla_1(\alpha)$.

(W $^\Delta$) $\Delta_{\lambda_1}(\alpha) \rightarrow \Delta_{\lambda_2}(\alpha)$, if $\lambda_1 \geq \lambda_2$.

and the modus ponens rule. We could then also introduce the following abbreviations:

$$\mathbf{\Pi}_\lambda(\alpha) = \bigvee_{\omega \in \llbracket \alpha \rrbracket} \Delta_\lambda(conj_\omega) \quad (13)$$

$$\mathbf{N}_\lambda(\alpha) = \neg \mathbf{\Pi}_{\nu(\lambda)}(\neg \alpha) \quad (14)$$

The resulting logic is very similar to GPL. However, for **(D $^\Delta$)** to be sound, we need to restrict e-models to possibility distributions π for which $\pi(\omega) = 0$ for at least one propositional interpretation ω . Similarly, for these axioms to be complete, we need to drop the requirement that $\pi(\omega) = 1$ for at least one interpretation. In fact, the soundness and completeness result

from Proposition 1 can straightforwardly be adapted to a logic centered on the Δ modality, by taking advantage of the following duality:

$$\pi \models \mathbf{N}_\lambda(\alpha) \quad \text{iff} \quad \bar{\pi} \models \mathbf{\Delta}_\lambda(\neg\alpha) \quad (15)$$

where the possibility distribution $\bar{\pi}$ is defined as $\bar{\pi}(\omega) = 1 - \pi(\omega)$ for all $\omega \in \Omega$. This duality can be readily verified using the definitions of the N and Δ measures in possibility theory (see Section 2.1).

However it is straightforward to show that (\mathbf{K}^Δ) , (Δ) and (\mathbf{W}^Δ) are valid in GPL. We can furthermore show that the following formulas are valid in GPL:

$$\begin{aligned} \mathbf{\Delta}_\lambda(\alpha) \wedge \mathbf{\Delta}_\lambda(\beta) &\equiv \mathbf{\Delta}_\lambda(\alpha \vee \beta) & \nabla_\lambda(\alpha \vee \beta) &\rightarrow \nabla_\lambda(\alpha) \wedge \nabla_\lambda(\beta) \\ \mathbf{\Delta}_\lambda(\alpha) \vee \mathbf{\Delta}_\lambda(\beta) &\rightarrow \mathbf{\Delta}_\lambda(\alpha \wedge \beta) & \nabla_\lambda(\alpha) \vee \nabla_\lambda(\beta) &\equiv \nabla_\lambda(\alpha \wedge \beta) \end{aligned}$$

and

$$\mathbf{\Delta}_{\lambda_1}(\alpha \wedge \beta) \wedge \mathbf{\Delta}_{\lambda_2}(\neg\alpha \wedge \gamma) \rightarrow \mathbf{\Delta}_{\min(\lambda_1, \lambda_2)}(\beta \wedge \gamma) \quad (16)$$

$$\nabla_{\lambda_1}(\alpha \wedge \beta) \wedge \mathbf{\Delta}_{\lambda_2}(\neg\alpha \wedge \gamma) \rightarrow \nabla_{\lambda_1}(\beta \wedge \gamma), \text{ if } \lambda_2 \geq \nu(\lambda_1) \quad (17)$$

Note that (16) is the counterpart of a basic inference rule of the logic of accumulated data [50].

For any possibility distribution π over Ω , we can easily define a GPL knowledge base which has π as its only e-model, using the modality Δ . In particular, let $\alpha_1, \dots, \alpha_k$ be propositional formulas such that $\llbracket \alpha_i \rrbracket = \{\omega \mid \pi(\omega) \geq \frac{i}{k}\}$. Then we define the knowledge base Φ_π as:

$$\Phi_\pi = \bigwedge_{i=1}^k \mathbf{N}_{\nu(\frac{i}{k})}(\alpha_i) \wedge \mathbf{\Delta}_{\frac{i}{k}}(\alpha_i). \quad (18)$$

A formula of the form Φ_π defines a GPL base which is complete in the following sense.

Proposition 2. $\forall \alpha \in \mathcal{L}, \lambda \in \Lambda, \Phi_\pi \vdash \mathbf{N}_\lambda(\alpha)$ or $\Phi_\pi \vdash \neg \mathbf{N}_\lambda(\alpha)$

Proof. In Equation (18), the degree of possibility of each $\omega \in \llbracket \alpha_i \rrbracket$ is defined by inequalities from above and from below. Indeed, $\mathbf{\Delta}_{\frac{i}{k}}(\alpha_i)$ means that $\pi(\omega) \geq \frac{i}{k}$ for all $\omega \in \llbracket \alpha_i \rrbracket$, whereas, $\mathbf{N}_{\nu(\frac{i}{k})}(\alpha_i)$ means $\pi(\omega) \leq \frac{i-1}{k}$ for all $\omega \notin \llbracket \alpha_i \rrbracket$. It follows that $\pi(\omega) = 0$ if $\omega \notin \llbracket \alpha_1 \rrbracket$, $\pi(\omega) = \frac{i}{k}$ if $\omega \in \llbracket \alpha_i \rrbracket \setminus \llbracket \alpha_{i+1} \rrbracket$ (for $i < k$) and $\pi(\omega) = 1$ if $\omega \in \llbracket \alpha_k \rrbracket$. In other words, π is indeed the only e-model of Φ_π . Since we clearly have $N(\alpha) \geq \lambda$ or $\neg(N(\alpha) \geq \lambda)$ for any necessity measure, it follows that $\Phi_\pi \vdash \mathbf{N}_\lambda(\alpha)$ or $\Phi_\pi \vdash \neg \mathbf{N}_\lambda(\alpha)$. \square

If we view the epistemic state of an agent as a possibility distribution, this means that every epistemic state can be modelled using a GPL knowledge base. Conceptually, the construction of Φ_π relates to the notion of “only knowing” from Levesque [6]. For example, assume that we want to model that all the agent knows is that β is true with certainty $\frac{j}{k}$.

Then we have $\pi(\omega) = 1$ for $\omega \in \llbracket \beta \rrbracket$ and $\pi(\omega) = \frac{k-j}{k}$ for $\omega \notin \llbracket \beta \rrbracket$. This means that in the notation of (18), $\alpha_{k-j+1} = \dots \alpha_k = \beta$ and we obtain $\Phi_\pi = \Delta_1(\beta) \wedge \mathbf{N}_{\nu(\frac{k-j+1}{k})}(\beta) \wedge \Delta_{\frac{k-j}{k}}(\top)$. In the case when $k = 1$, Equation (18) reads $\mathbf{N}_1(\alpha) \wedge \Delta_1(\alpha)$ and isolates a single crisp e-model corresponding to the set of classical models of α as already pointed out in [28]. It expresses that we precisely know the epistemic state of the external agent, namely that (s)he only knows that α is true.

In practice, we will often have incomplete knowledge about the epistemic state of this agent. Suppose we only know that the epistemic state is among those in $S \subseteq \mathcal{P}_k$. This can be encoded as a GPL knowledge base $\Phi_S = \bigvee_{\pi \in S} \Phi_\pi$ with Φ_π defined as above. As a consequence, any GPL knowledge base is semantically equivalent to a formula of the form Φ_S , and any subset of epistemic states can be captured by a GPL knowledge base.

Since the modality Δ was introduced as an abbreviation, allowing this modality has no impact on the expressiveness of the language or on the completeness of the axiomatization. However, the formula $\Delta_\lambda(\alpha)$ abbreviates a GPL formula which may be of exponential size, and allowing the modality Δ in the language is thus essential if we want to capture our knowledge about an agent's epistemic state in a compact way. As we will see in Section 7, this is reflected in an increase in computational complexity.

4.2. Contextual ignorance as restricted guaranteed possibility

The modality Δ allows us to express limitations on what an agent knows. However, it does not readily allow us to explicitly encode the ignorance of an agent on a particular topic.

Example 2. Consider again the scenario from Example 1 and suppose we want to encode that “all the agent knows about the games in round 3 is that Wales has won its game”. We cannot represent this as $\mathbf{N}_1(\text{won}_3(\text{wal})) \wedge \Delta_1(\text{won}_3(\text{wal}))$, as that would entail e.g., $\neg \mathbf{N}_1(\text{won}_2(\text{wal}))$, which is not warranted.

To encode limitations on the knowledge of the agent on a particular topic, understood as a set of propositional variables $\mathcal{X} \subseteq \text{At}$, we propose the following variant of the Δ modality:

$$\Delta_\lambda^\mathcal{X}(\alpha) = \bigwedge_{\omega \in \llbracket \alpha \rrbracket} \Pi_\lambda(\text{conj}_\omega^\mathcal{X})$$

where $\text{conj}_\omega^\mathcal{X}$ is the restriction of conj_ω to those literals about variables in \mathcal{X} , i.e., $\text{conj}_\omega^\mathcal{X} = \bigwedge \{x \mid x \in \mathcal{X}, \omega \models x\} \wedge \bigwedge \{\neg x \mid x \in \mathcal{X}, \omega \models \neg x\}$. Note that $\models_{GPL} \Delta_\lambda(\alpha) \equiv \Delta_\lambda^{\text{At}}(\alpha)$. For example, in the scenario from Example 2, instead of asserting $\Delta_1(\text{won}_3(\text{wal}))$, we can assert $\Delta_1^\mathcal{X}(\text{won}_3(\text{wal}))$, with $\mathcal{X} = \{\text{plays}_3(x, y) \mid x, y \in T\} \cup \{\text{won}_3(x) \mid x \in T\}$ the set of all atomic formulas about round 3 of the championship. As we will see in Section 7, allowing this refinement of the Δ modality leads to a further increase in computational complexity.

4.3. Ignorance as minimal specificity

The less specific than relation \preceq defines a partial order on the set of e-models of a GPL knowledge base K in a natural way, which allows us to introduce two non-monotonic entailment relations:

- We say that Φ is a brave consequence of K , written $K \models_b \Phi$ iff Φ is satisfied by a minimally specific e-model of K .
- We say that Φ is a cautious consequence of K , written $K \models_c \Phi$ iff Φ is satisfied by all minimally specific e-models of K .

In standard possibilistic logic, every knowledge base K has a least specific e-model π_K . As a result, in standard possibilistic logic, the entailment relations \models , \models_b and \models_c coincide. In GPL, this is no longer the case.

Example 3. Let $u, v, w \in At$. The formula $\mathbf{N}_1(u) \vee \mathbf{N}_1(v)$ has two minimally specific e-models π_u and π_v defined as:

$$\pi_u(\omega) = \begin{cases} 0 & \text{if } \omega \models \neg u \\ 1 & \text{otherwise} \end{cases} \quad \pi_v(\omega) = \begin{cases} 0 & \text{if } \omega \models \neg v \\ 1 & \text{otherwise} \end{cases} \quad (19)$$

This already shows that \models and \models_b do not coincide, as e.g., $\mathbf{N}_1(u) \vee \mathbf{N}_1(v) \models_b \mathbf{N}_1(u)$ while clearly $\mathbf{N}_1(u) \vee \mathbf{N}_1(v) \not\models \mathbf{N}_1(u)$. To see why \models and \models_c do not coincide, note that since u, v and w are logically independent, $\mathbf{N}_1(u) \vee \mathbf{N}_1(v) \models_c \mathbf{\Pi}_1(w) \wedge \mathbf{\Pi}_1(\neg w)$ while $\mathbf{N}_1(u) \vee \mathbf{N}_1(v) \not\models \mathbf{\Pi}_1(w) \wedge \mathbf{\Pi}_1(\neg w)$.

Reasoning about what is true in all minimally specific e-models, as opposed to all e-models, is similar to making a kind of meta-closed-world assumption. Intuitively, it amounts to assuming that the agent is ignorant about a formula α unless it has been asserted that the agent knows whether α is true or false. For example, in the scenario from Example 2, we can simply assert $\mathbf{N}_1(\text{won}_3(\text{wal}))$, as the knowledge that the agent is ignorant about anything else related to round 3 is implicit in the fact that no other knowledge has been asserted. However, even under this assumption, there may be situations in which we are ignorant about whether the agent knows whether α is true, as illustrated in the next example.

Example 4. Consider again the formula $\mathbf{N}_1(u) \vee \mathbf{N}_1(v)$ and its minimally specific e-models π_u and π_v from Example 3. It holds that $\mathbf{N}_1(u) \vee \mathbf{N}_1(v) \not\models_c \mathbf{N}_1(u) \vee \mathbf{N}_1(\neg u)$ (since $\pi_v \not\models \mathbf{N}_1(u)$ and $\pi_v \not\models \mathbf{N}_1(\neg u)$), i.e., we cannot conclude that the agent knows about u . However, we also have $\mathbf{N}_1(u) \vee \mathbf{N}_1(v) \not\models_c \mathbf{\Pi}_1(u) \wedge \mathbf{\Pi}_1(\neg u)$ (since $\pi_u \models \mathbf{N}_1(u)$), i.e., we cannot conclude that the agent is ignorant about u either.

In [30], it is argued that the epistemic state of an agent can be modelled as a propositional formula α , although through introspection the agent knows more than what is encoded by α directly. For example, if $\alpha \not\models \beta$, the agent knows that it does not know β in this setting. In GPL, we can characterize what the agent knows (with full certainty) as the consequences of $\mathbf{N}_1(\alpha)$ under the inference relation \models_c , which coincides with \models_b in this case, since $\mathbf{N}_1(\alpha)$ has a unique least specific e-model. In particular, the argument of [30], translated to the terminology of this paper, is that theories which model the epistemic state of an agent should have such a unique least specific e-model. Formulas for which this is not the case (e.g., $\mathbf{N}_1(a) \vee \mathbf{N}_1(b)$) are called dishonest by [30]. In our setting, however, we should not

Table 1: Overview of the considered variants of GPL.

logic	meta-atoms	axioms	e-models
GPL	$\mathbf{N}_\lambda(\alpha)$	(PL) , (K) , (N) , (D) , (W)	possibility degrees from Λ_k
$\text{GPL}_{\succ}^{\text{core}}$	$\alpha \succ \beta$	(PL) , (Ax₁) – (Ax₅)	possibility degrees from $[0, 1] \cap \mathbb{Q}$
GPL_{\succ}	$\mathbf{N}_\lambda(\alpha), \alpha \succ \beta$	(PL) , (W) , (Ax₁) – (Ax₈)	possibility degrees from $[0, 1] \cap \mathbb{Q}$
GPL_c	$c(\alpha \sim \beta)$	(PL) , (RE) , (LLE) , (RW) , (OR) , (CM) , (CUT) , (WRM) , (INC)	possibility degrees from $[0, 1] \cap \mathbb{Q}$

exclude GPL formulas which have multiple minimally specific e-models, if our aim is to reason about the revealed beliefs of another agent. Indeed, such situations can easily arise if we have incomplete information about the epistemic state of an external agent. For example, we may know that this agent has received and read the notification email on a submitted conference paper, while not knowing whether the paper has been accepted or rejected. In that case we know $\mathbf{N}_1(\textit{accept}) \vee \mathbf{N}_1(\neg\textit{accept})$, i.e., all we know is that the external agent knows the result with certainty.

5. Reasoning about comparative uncertainty

In this section, we analyze in more detail how GPL can be used for reasoning about comparative uncertainty. First, in Section 5.1, we introduce the logic GPL_{\succ} , which extends GPL with formulas of the form $\alpha \succ \beta$, expressing that α is strictly more certain than β . We will also consider a fragment $\text{GPL}_{\succ}^{\text{core}}$ of this logic, in which only (propositional combinations of) such comparative certainty assertions are allowed. In Section 5.2 we then propose an axiomatization of $\text{GPL}_{\succ}^{\text{core}}$, by extending the axiomatization of a logic for reasoning about partially ordered formulas [7]. This axiomatization is subsequently extended to an axiomatization of GPL_{\succ} in Section 5.3. Finally, in Section 5.4 we show how the ability of GPL to capture comparative uncertainty can be used to reason about propositional combinations of default rules in the sense of System P [5], and we introduce a variant of $\text{GPL}_{\succ}^{\text{core}}$, called GPL_c , which is aimed specifically at reasoning about propositional combinations of default rules.

Table 1 presents an overview of the considered logics. Like GPL itself, each of the logics introduced in this section is two-tiered, in the sense that formulas correspond to propositional combinations of meta-atoms, and each meta-atom corresponds to a modal operator applied to one or two propositional formulas. From a syntactic point of view, the only difference between the different logics lies in the considered modal operators. From a semantic point of view, an important difference between GPL and the variants considered in this section is that here we will allow infinitely many possibility degrees.

5.1. A logic for comparative uncertainty

Rather than expressing absolute degrees of certainty in the logic, one may find it more natural to simply express the fact that a proposition α is at least as (resp. more) certain than another one β . In possibility theory, this can be expressed as $N(\alpha) \geq N(\beta)$ (resp. $N(\alpha) > N(\beta)$). The statement $N(\alpha) \geq N(\beta)$ is equivalent to $\forall \lambda \in \Lambda_k^+, N(\beta) \geq \lambda$ implies

$N(\alpha) \geq \lambda$. In GPL, as already pointed out in Section 3.2, the latter expression can be syntactically described as

$$\bigwedge_{i=1}^k \mathbf{N}_{\frac{i}{k}}(\alpha) \vee \neg \mathbf{N}_{\frac{i}{k}}(\beta).$$

The statement $N(\alpha) > N(\beta)$ is the negation of $N(\beta) \geq N(\alpha)$, hence it can be encoded by

$$\bigvee_{i=1}^k \mathbf{N}_{\frac{i}{k}}(\alpha) \wedge \neg \mathbf{N}_{\frac{i}{k}}(\beta)$$

as was already illustrated in Example 1.

So far, we have assumed that the number of certainty levels k is fixed in advance. If k is not chosen sufficiently large, however, this can lead to some unwanted results. For example, if $k = 2$, the following statement, expressing $N(\alpha_1) > N(\alpha_2) > N(\alpha_3) > N(\alpha_4)$ is not satisfiable:

$$\left(\bigvee_{i=1}^k \mathbf{N}_{\frac{i}{k}}(\alpha_1) \wedge \neg \mathbf{N}_{\frac{i}{k}}(\alpha_2) \right) \wedge \left(\bigvee_{i=1}^k \mathbf{N}_{\frac{i}{k}}(\alpha_2) \wedge \neg \mathbf{N}_{\frac{i}{k}}(\alpha_3) \right) \wedge \left(\bigvee_{i=1}^k \mathbf{N}_{\frac{i}{k}}(\alpha_3) \wedge \neg \mathbf{N}_{\frac{i}{k}}(\alpha_4) \right)$$

In a purely ordinal setting, however, we would not expect a formula asserting $N(\alpha_1) > N(\alpha_2) > N(\alpha_3) > N(\alpha_4)$ to be always unsatisfiable. To address this issue, we introduce GPL_{\succ} , a logic for reasoning about comparative uncertainty in which an unbounded number of certainty levels can be used at the semantic level. The language \mathcal{L}_{\succ} of the logic GPL_{\succ} is defined as follows:

- If $\alpha \in \mathcal{L}$ and $\lambda \in \Lambda_k^+$, then $\mathbf{N}_{\lambda}(\alpha)$ belongs to \mathcal{L}_{\succ} ;
- if $\alpha, \beta \in \mathcal{L}$ then $\alpha \succ \beta$ belongs to \mathcal{L}_{\succ} ;
- if Φ and Ψ belong to \mathcal{L}_{\succ} , then $\neg \Phi$ and $\Phi \wedge \Psi$ are also in \mathcal{L}_{\succ} .

At the semantic level, e-models will be allowed to take arbitrary values from $\Lambda_* = [0, 1] \cap \mathbb{Q}$. As a result, $\mathbf{N}_{\lambda}(\alpha)$ cannot be defined as an abbreviation in GPL_{\succ} . We will use $\alpha \sim \beta$ as an abbreviation for $\neg(\alpha \succ \beta) \wedge \neg(\beta \succ \alpha)$ and $\alpha \succeq \beta$ as an abbreviation of $\neg(\beta \succ \alpha)$. Note that this reflects the fact that at the semantic level, \succeq will be a complete preordering with \sim as its equivalence part. Intuitively, $\alpha \succ \beta$ means that α is strictly more certain than β , $\alpha \succeq \beta$ means that α is at least as certain as β , and $\alpha \sim \beta$ means that α is equally certain as β .

Let π be a normalized possibility distribution over Ω such that $\pi(\omega) \in \Lambda_*$ for every $\omega \in \Omega$, and let N be the necessity measure induced by π . The notion of e-model for formulas from \mathcal{L}_{\succ} is defined as follows:

- π is an e-model of $\mathbf{N}_{\lambda}(\alpha)$ if $N(\alpha) \geq \lambda$;
- π is an e-model of $\alpha \succ \beta$ iff $N(\alpha) > N(\beta)$;

- π is an e-model of $\Phi_1 \wedge \Phi_2$ iff π is an e-model of Φ_1 and of Φ_2 ;
- π is an e-model of $\neg\Phi_1$ iff π is not an e-model of Φ_1 .

If π is an e-model of Φ we write $\pi \models_{\succ} \Phi$, and if every e-model of a set of GPL_{\succ} formulas K is also an e-model of Φ , we write $K \models_{\succ} \Phi$. If there is no cause for confusion, we also write \models_{\succ} as \models .

In the following, we will also consider the fragment $\text{GPL}_{\succ}^{\text{core}}$ of GPL_{\succ} , which is restricted to the language $\mathcal{L}_{\succ}^{\text{core}}$, defined as follows:

- if $\alpha, \beta \in \mathcal{L}$ then $\alpha \succ \beta$ belongs to $\mathcal{L}_{\succ}^{\text{core}}$;
- if Φ and Ψ belong to $\mathcal{L}_{\succ}^{\text{core}}$, then $\neg\Phi$ and $\Phi \wedge \Psi$ are also in $\mathcal{L}_{\succ}^{\text{core}}$.

In other words, $\text{GPL}_{\succ}^{\text{core}}$ is concerned only with propositional combinations of comparative certainty statements, while in GPL_{\succ} we also allow statements such as $\mathbf{N}_{\frac{1}{2}}(\alpha) \vee (\alpha \succ \beta)$.

5.2. Axiomatizing $\text{GPL}_{\succ}^{\text{core}}$

In this section, we introduce an axiomatization for the logic $\text{GPL}_{\succ}^{\text{core}}$. To this end, we start from the axiomatization of the relative certainty logic that was studied by Touazi et al. [7]. A knowledge base in this latter logic is a conjunction of statements of the form $\alpha \succ \beta$, i.e., the language they considered is the disjunction- and negation-free fragment of $\mathcal{L}_{\succ}^{\text{core}}$. In [7] axiomatization for the resulting logic was introduced, inspired by an earlier axiomatization that was proposed in [51]. It contains one axiom:

(Ax₁) $\alpha \succ \perp$ if α is a tautology

and three inference rules:

(RI₁) If $\beta \succ \alpha \wedge \chi$ and $\alpha \succ \beta \wedge \chi$ then $\beta \wedge \alpha \succ \chi$

(RI₂) If $\alpha \succ \beta$, $\alpha \vDash \alpha'$ and $\beta' \vDash \beta$ then $\alpha' \succ \beta'$

(RI₃) If $\alpha \succ \beta$ and $\beta \succ \alpha$ then \perp

Note that the reason why **(RI₁)**–**(RI₃)** were formulated as inference rules in [7], rather than axioms, is because propositional combinations of comparative certainty statements were not considered in [7]. **(RI₁)** is the so-called qualitiveness axiom [51] that ensures that if both α and β are more certain than χ then so is $\alpha \wedge \beta$, which is only compatible with necessity measures in the totally ordered case. **(RI₂)** is a natural mononicity assumption in agreement with logical entailment.

The axiom and inference rules are sufficient to ensure that \succ is a strict partial order. Indeed, **(RI₃)** encodes asymmetry, while the transitivity of \succ was established in [7], i.e., the following inference rule can be derived from **(Ax₁)** and **(RI₁)**–**(RI₃)**:

$$\text{If } \alpha \succ \beta \text{ and } \beta \succ \gamma \text{ then } \alpha \succ \gamma \quad (20)$$

A soundness and completeness result was also established in [7] with respect to a semantics in terms of partial orders between sets of models $\llbracket \alpha \rrbracket$ and $\llbracket \beta \rrbracket$ verifying obvious semantic counterparts of the axioms. Note in particular that, unlike in $\text{GPL}_{\succ}^{\text{core}}$, the semantics considered in [7] is not based on necessity measures.

Clearly, the axiom (\mathbf{Ax}_1) and inference rules (\mathbf{RI}_1) – (\mathbf{RI}_3) are also sound for $\text{GPL}_{\succ}^{\text{core}}$. Note, however, that (\mathbf{RI}_1) – (\mathbf{RI}_3) can be formulated as axioms in $\text{GPL}_{\succ}^{\text{core}}$:

$$(\mathbf{Ax}_2) \quad (\beta \succ \alpha \wedge \chi) \wedge (\alpha \succ \beta \wedge \chi) \rightarrow (\beta \wedge \alpha \succ \chi)$$

$$(\mathbf{Ax}_3) \quad (\alpha \succ \beta) \rightarrow (\alpha' \succ \beta') \text{ if } \alpha \rightarrow \alpha' \text{ and } \beta' \rightarrow \beta \text{ are tautologies}$$

$$(\mathbf{Ax}_4) \quad \neg((\alpha \succ \beta) \wedge (\beta \succ \alpha))$$

Further axioms are needed to capture the fact that arbitrary propositional combinations of comparative certainty statements are allowed in $\text{GPL}_{\succ}^{\text{core}}$. Furthermore, to capture the fact that our semantics is based on necessity measures, we will need to impose that \succ is the complement of a weak order. Recall that $\alpha \sim \beta$ stands for $\neg(\alpha \succ \beta) \vee (\beta \succ \alpha)$. It is then obvious that the following formulas can be derived from (\mathbf{Ax}_1) – (\mathbf{Ax}_4) together with the axioms from propositional logic and modus ponens:

$$(\mathbf{S}_1) \quad (\alpha \succ \beta) \vee (\beta \succ \alpha) \vee (\alpha \sim \beta)$$

$$(\mathbf{S}_2) \quad \neg((\alpha \succ \beta) \wedge (\alpha \sim \beta))$$

$$(\mathbf{S}_3) \quad (\alpha \sim \beta) \rightarrow (\beta \sim \alpha)$$

To ensure that \sim is transitive, we augment the relative certainty logic from [7] with the following additional axiom:

$$(\mathbf{Ax}_5) \quad (\alpha \succ \beta) \wedge (\alpha \sim \alpha') \rightarrow (\alpha' \succ \beta)$$

For K a set of $\text{GPL}_{\succ}^{\text{core}}$ formulas and Φ a $\text{GPL}_{\succ}^{\text{core}}$ formula, we write $K \vdash_{\succ}^{\text{core}} \Phi$ to denote that Φ can be derived from K using (\mathbf{Ax}_1) – (\mathbf{Ax}_5) , modus ponens and the axioms of classical logic. When there is no cause for confusion, we also write $\vdash_{\succ}^{\text{core}}$ as \vdash . We now show some properties of \succ and \sim that follow from the proposed axioms and inference rules, and which will be useful for showing the soundness and completeness result.

Proposition 3. *The following theorems hold:*

$$\vdash_{\succ}^{\text{core}} \quad (\alpha \sim \beta) \wedge (\beta \sim \gamma) \rightarrow (\alpha \sim \gamma) \tag{21}$$

$$\vdash_{\succ}^{\text{core}} \quad \alpha \sim \alpha \tag{22}$$

$$\vdash_{\succ}^{\text{core}} \quad (\alpha \sim \beta) \rightarrow (\alpha' \sim \beta) \text{ if } \alpha \equiv \alpha' \text{ is a tautology} \tag{23}$$

$$\vdash_{\succ}^{\text{core}} \quad (\alpha \succ \beta) \wedge (\beta \sim \beta') \rightarrow (\alpha \succ \beta') \tag{24}$$

$$\vdash_{\succ}^{\text{core}} \quad (\alpha \succ \beta \wedge \chi) \rightarrow (\alpha \succ \beta) \vee (\alpha \succ \chi) \tag{25}$$

Proof. **Eq. 21** From (\mathbf{S}_1) we know that $\alpha \sim \gamma$ holds unless $\alpha \succ \gamma$ or $\gamma \succ \alpha$ holds. However, from $\alpha \succ \gamma$ and $\alpha \sim \beta$, we derive $\beta \succ \gamma$ using (\mathbf{Ax}_5) , which leads to a contradiction with $\beta \sim \gamma$ when using (\mathbf{S}_2) . Similarly, from $\gamma \succ \alpha$ and $\beta \sim \gamma$, we derive $\beta \succ \alpha$ using (\mathbf{S}_3) and (\mathbf{Ax}_5) , leading to a contradiction with $\alpha \sim \beta$ using (\mathbf{S}_3) and (\mathbf{S}_2) .

Eq. 22 From (\mathbf{S}_1) and the irreflexivity of \succ we immediately obtain (22).

Eq. 23 Note that from (\mathbf{S}_1) we know that $\alpha' \sim \beta$ holds unless $\alpha' \succ \beta$ or $\beta \succ \alpha'$ holds. From $\alpha' \succ \beta$ and $\models \alpha \equiv \alpha'$, we can derive $\alpha \succ \beta$ using (\mathbf{Ax}_3) , which leads to a contradiction with $\alpha \sim \beta$ when using (\mathbf{S}_2) . Similarly, from $\beta \succ \alpha'$ and $\models \alpha \equiv \alpha'$, we can derive $\beta \succ \alpha$ using (\mathbf{Ax}_3) , which leads to a contradiction with $\alpha \sim \beta$ when using (\mathbf{S}_3) and (\mathbf{S}_2) .

Eq. 24 From (\mathbf{S}_1) we know that $\alpha \succ \beta'$ holds unless $\alpha \sim \beta'$ or $\beta' \succ \alpha$ holds. From $\alpha \sim \beta'$ and $\beta \sim \beta'$, we can derive $\alpha \sim \beta$ using (\mathbf{S}_3) and (21), which leads to a contradiction with $\alpha \succ \beta$ when using (\mathbf{S}_2) . From $\beta' \succ \alpha$ and $\beta \sim \beta'$, we can derive $\beta \succ \alpha$ using (\mathbf{S}_3) and (\mathbf{Ax}_5) , which leads to a contradiction with $\alpha \succ \beta$ using (\mathbf{Ax}_4) .

Eq. 25 Suppose $\neg(\alpha \succ \beta) \wedge \neg(\alpha \succ \chi)$ holds. Then by (\mathbf{S}_1) we have

$$((\beta \succ \alpha) \vee (\beta \sim \alpha)) \wedge ((\chi \succ \alpha) \vee (\chi \sim \alpha))$$

Now suppose that $\alpha \succ \beta \wedge \chi$ were also the case. From either $\beta \succ \alpha$ or $\beta \sim \alpha$ we can then derive $\beta \succ \beta \wedge \chi$, using respectively (20) and (\mathbf{Ax}_5) . Similarly, from either $\chi \succ \alpha$ or $\chi \sim \alpha$ we can derive $\chi \succ \beta \wedge \chi$. Using (\mathbf{Ax}_3) we can derive $\beta \succ \chi \wedge \beta \wedge \chi$ and $\chi \succ \beta \wedge \beta \wedge \chi$, which using (\mathbf{Ax}_2) gives us $\beta \wedge \chi \succ \beta \wedge \chi$. Using (\mathbf{Ax}_3) we derive a contradiction, and thus we can conclude $\neg(\alpha \succ \beta \wedge \chi)$. \square

Note that it follows that \sim is an equivalence relation. Indeed, the symmetry of \sim is expressed in (\mathbf{S}_3) , while its transitivity and reflexivity are expressed in (21) and (22) respectively. It follows that the relation \succeq corresponds to a complete preorder (or weak order) on the language \mathcal{L} . In contrast, when using the partial order semantics of [7], the relation \sim is not transitive, and property (25) does not hold, namely we do not have that $\alpha > \beta \wedge \gamma$ implies one of $\alpha > \beta$ or $\alpha > \gamma$ in the comparative certainty logic of [7]. This clearly illustrates the difference with the relative certainty logic from [7]. Axiom (\mathbf{Ax}_5) is not derivable from $(\mathbf{Ax}_1\text{-}\mathbf{Ax}_4)$, since the latter only ensure that \succ is a partial order, hence in that case \sim also covers incomparability and is generally not transitive.

We can show the following soundness and completeness result for this extended set of axioms w.r.t. the $\text{GPL}_{\succeq}^{\text{core}}$ semantics.

Proposition 4. *Let Φ and Ψ be formulas in $\mathcal{L}_{\succeq}^{\text{core}}$. It holds that $\Phi \models_{\succ} \Psi$ iff $\Phi \vdash_{\succeq}^{\text{core}} \Psi$.*

Proof. The proof is given in Appendix B. \square

Finally, note that while we have focused on comparative necessity, in a similar way we could develop a logic for reasoning about comparative guaranteed possibility. Among other applications, it seems that such a logic could be useful for modelling and reasoning about desires [52].

5.3. Axiomatizing GPL_{\succ}

In GPL_{\succ} we can express formulas such as $\mathbf{N}_{\lambda}(\alpha) \wedge (\beta \succeq \alpha)$. Clearly, this formula entails $\mathbf{N}_{\lambda}(\beta)$. To capture such inferences at the syntactic level, we extend the axiomatization of GPL_{\succ}^{core} with the following axioms:

$$(\mathbf{Ax}_6) \quad \mathbf{N}_{\lambda}(\alpha) \wedge \neg \mathbf{N}_{\lambda}(\beta) \rightarrow (\alpha \succ \beta)$$

$$(\mathbf{Ax}_7) \quad \mathbf{N}_1(\alpha) \equiv (\alpha \sim \top)$$

$$(\mathbf{Ax}_8) \quad \mathbf{N}_{\lambda}(\alpha) \rightarrow (\alpha \succ \perp)$$

Note that (\mathbf{Ax}_6) is equivalent to $\mathbf{N}_{\lambda}(\alpha) \wedge (\beta \succeq \alpha) \rightarrow \mathbf{N}_{\lambda}(\beta)$. In addition to these axioms, we will also use the GPL axioms (\mathbf{PL}) (i.e. the axioms of classical logic) and (\mathbf{W}) . In particular, for K a set of GPL_{\succ} formulas and Φ a GPL_{\succ} formula (or set of GPL_{\succ} formulas), we write $K \vdash_{\succ} \Phi$ to denote that Φ can be derived from K using (\mathbf{Ax}_1) – (\mathbf{Ax}_8) , (\mathbf{PL}) , (\mathbf{W}) and modus ponens.

Note (\mathbf{K}) , (\mathbf{N}) and (\mathbf{D}) , which are axioms in GPL but not in GPL_{\succ} , can be derived as theorems in GPL_{\succ} . To show this for (\mathbf{K}) , we prove that $\mathbf{N}_{\lambda}(\alpha \rightarrow \beta) \wedge \mathbf{N}_{\lambda}(\alpha) \wedge \neg \mathbf{N}_{\lambda}(\beta)$ is inconsistent. By applying (\mathbf{Ax}_6) twice, we can derive $(\alpha \rightarrow \beta \succ \beta) \wedge (\alpha \succ \beta)$. This can be weakened to $(\alpha \rightarrow \beta \succ \alpha \wedge \beta) \wedge (\alpha \succ (\alpha \rightarrow \beta) \wedge \beta)$ using (\mathbf{Ax}_3) , from which we can derive $\alpha \wedge (\alpha \rightarrow \beta) \succ \beta$ using (\mathbf{Ax}_2) . The latter formula can be weakened to $\beta \succ \beta$ by applying (\mathbf{Ax}_3) again, which is inconsistent with (\mathbf{Ax}_4) . The GPL axiom (\mathbf{N}) can straightforwardly be derived from (\mathbf{Ax}_7) and (22). Finally, (\mathbf{D}) has to be expressed as $\mathbf{N}_{\lambda}(\alpha) \rightarrow \neg \mathbf{N}_{\frac{1}{k}}(\neg \alpha)$, since the abbreviation $\mathbf{\Pi}_1(\alpha)$ is not used in GPL_{\succ} . We show that $\mathbf{N}_{\lambda}(\alpha) \wedge \mathbf{N}_{\frac{1}{k}}(\neg \alpha)$ is inconsistent. To this end, we can derive $\mathbf{N}_{\frac{1}{k}}(\perp)$ using (\mathbf{K}) , (\mathbf{N}) , (\mathbf{W}) and (\mathbf{PL}) . Using (\mathbf{Ax}_8) , this leads to $\perp \succ \perp$, which is inconsistent with (\mathbf{Ax}_4) .

We can show the following soundness and completeness result.

Proposition 5. *Let Φ and Ψ be formulas in \mathcal{L}_{\succ} . It holds that $\Phi \models_{\succ} \Psi$ iff $\Phi \vdash_{\succ} \Psi$.*

Proof. The proof is given in Appendix C. □

5.4. Reasoning about conditionals

There are two rather distinct traditions in the field of non-monotonic reasoning. On the one hand, formalisms such as answer set programming, Reiter’s default logic [53], and Moore’s autoepistemic logic [54] allow us to explicitly make default assumptions of the form “unless there is evidence to the contrary, assume X ”. We will discuss such forms of non-monotonic reasoning in more detail in Section 6. On the other hand, there is a large class of approaches to reason about rules with exceptions, based on the view that in the case of

a conflict, priority should be given to more specific rules: from the information that birds generally fly, penguins generally cannot fly, and all penguins are birds, these approaches allow us to derive that Tweety, who is a penguin, cannot fly. In this section we show that we can encode such exception-tolerant rules in $\text{GPL}_{\succ}^{\text{core}}$.

Let us write $\alpha \sim \beta$ to encode the conditional “if α then generally β ”. Several approaches have been proposed to reason about such exception-tainted rules [55, 56, 57, 58]. One of the most important results in this field is that despite the different intuitions underlying these approaches, there is a consensus shared with virtually all of them about the minimal set of conditionals of the form $\alpha \sim \beta$ that should be entailed by a given rule base $R = \{\alpha_1 \sim \beta_1, \dots, \alpha_n \sim \beta_n\}$. This common core of conclusions is captured by the inference rules of System P [5]:

$\alpha \sim \alpha$	(Reflexivity)
If $\models \alpha \equiv \alpha'$ and $\alpha \sim \beta$ then $\alpha' \sim \beta$	(Left logical equivalence)
If $\beta \models \beta'$ and $\alpha \sim \beta$ then $\alpha \sim \beta'$	(Right weakening)
If $\alpha \sim \gamma$ and $\beta \sim \gamma$ then $\alpha \vee \beta \sim \gamma$	(OR)
If $\alpha \sim \beta$ and $\alpha \sim \gamma$ then $\alpha \wedge \beta \sim \gamma$	(Weak monotony)
If $\alpha \wedge \beta \sim \gamma$ and $\alpha \sim \beta$ then $\alpha \sim \gamma$	(CUT)

The last two inference rules correspond to the idea of cumulativity, whereby $\alpha \wedge \beta \sim \gamma$ and $\alpha \sim \gamma$ are equivalent if $\alpha \sim \beta$ is taken for granted. If we only consider conditionals $\alpha \sim \beta$ for which $\alpha \not\models \perp$, the conclusions that System P allows us to derive from the rule base R can be characterized using possibility theory. Specifically, let \mathcal{P}_R be the set of all possibility measures Π for which $\Pi(\alpha_i \wedge \beta_i) > \Pi(\alpha_i \wedge \neg \beta_i)$ for every $i \in \{1, \dots, n\}$. It can then be shown [59] that $\alpha \sim \beta$ can be derived from R using the axioms of System P iff $\Pi(\alpha \wedge \beta) > \Pi(\alpha \wedge \neg \beta)$ for every $\Pi \in \mathcal{P}_R$. Moreover, it holds that $\Pi(\alpha \wedge \beta) > \Pi(\alpha \wedge \neg \beta)$ for the unique least specific possibility measure (i.e., the possibility measure induced by the least specific possibility distribution relative to a finite but sufficiently large set of certainty levels) in \mathcal{P}_R iff $\alpha \sim \beta$ is in the rational closure of R [56], the latter being a well-known refinement of System P.

This means that both System P and the rational closure can naturally be characterized using $\text{GPL}_{\succ}^{\text{core}}$. In particular, we associate with each conditional $\alpha \sim \beta$ the $\text{GPL}_{\succ}^{\text{core}}$ formula $c(\alpha \sim \beta)$, stating that $\Pi(\alpha \wedge \beta) > \Pi(\alpha \wedge \neg \beta)$, or equivalently $N(\alpha \rightarrow \beta) > N(\alpha \rightarrow \neg \beta)$:

$$c(\alpha \sim \beta) = (\alpha \rightarrow \beta) \succ (\alpha \rightarrow \neg \beta)$$

Similar as before, we find that we can interpret \succ as an abbreviation in the language of GPL_k , provided that k is sufficiently large. In particular, we have the following proposition.

Proposition 6. *Let $c(R) = \{c(\alpha_i \sim \beta_i) \mid (\alpha_i \sim \beta_i) \in R\}$. Assume that $\alpha, \alpha_1, \dots, \alpha_n$ are consistent. It holds that:*

- $c(R) \models_{\text{GPL}}^k c(\alpha \sim \beta)$ iff $\alpha \sim \beta$ can be derived from R using the axioms of System P, provided that $k \geq |R| + 1$.

- $\alpha \sim \beta$ is in the rational closure of R iff $c(\alpha \sim \beta)$ is satisfied by the (unique) least specific GPL_k e-model of $c(R)$, provided that $k \geq |R|$.

Proof. If $k \geq |R| + 1$, it follows from Proposition 18 that $c(R) \wedge \neg c(\alpha \sim \beta)$ is inconsistent, or equivalently that $c(R) \models_{\succ} c(\alpha \sim \beta)$, iff every e-model of $c(R)$ also satisfies $c(\alpha \sim \beta)$. In other words, given the result from [59], $c(R) \models_{\succ} c(\alpha \sim \beta)$ holds iff $\alpha \sim \beta$ can be derived from R using the axioms of System P.

The least specific e-model of $c(R)$ corresponds to the least specific e-model of the possibilistic counterpart to the Z-ranking [58]. As this possibilistic counterpart corresponds to a knowledge base with at most $|R|$ levels, it is clear that least specific e-model of $c(R)$ will correspond to the rational closure of R if $k \geq |R|$. □

In contrast to System P, GPL_{\succ}^{core} can also be used to reason about propositional combinations of conditionals. For example, it holds that

$$c(a \sim b) \vee c(a \sim c) \models_{\succ} c(a \sim b \vee c) \quad (26)$$

Indeed, using **(Ax₃)** it follows from either of $(a \rightarrow b) \succ (a \rightarrow \neg b)$ and $(a \rightarrow c) \succ (a \rightarrow \neg c)$ that $(a \rightarrow b \vee c) \succ (a \rightarrow \neg(b \vee c))$. Hence, we find $((a \rightarrow b) \succ (a \rightarrow \neg b)) \vee ((a \rightarrow c) \succ (a \rightarrow \neg c)) \models (a \rightarrow b \vee c) \succ (a \rightarrow \neg(b \vee c))$, which is equivalent to (26). This means that we can use GPL_{\succ}^{core} to define a logic for reasoning about conditionals. Let us define the language \mathcal{L}_c as the following fragment of $\mathcal{L}_{\succ}^{core}$:

- If $\alpha, \beta \in \mathcal{L}$ and $\alpha \not\models \perp$, then $((\alpha \rightarrow \beta) \succ (\alpha \rightarrow \neg \beta))$ belongs to \mathcal{L}_c ;
- If Φ and Ψ belong to \mathcal{L}_c , then $\neg \Phi$ and $\Phi \wedge \Psi$ are also in \mathcal{L}_c .

We will refer to the corresponding logic as GPL_c . Its satisfaction relation \models_c is simply defined as the restriction of \models_{\succ} to the language fragment \mathcal{L}_c . To reason about formulas in \mathcal{L}_c we can rely on the axiomatization of $\mathcal{L}_{\succ}^{core}$ proposed in Section 5.2. Alternatively, as we show next, we can also define a syntactic inference relation \vdash_c that allows derivations to say within the language fragment \mathcal{L}_c . To this end, we will extend the inference rules of System P, all of which are theorems in GPL_{\succ}^{core} . In particular, it was shown in [7] that the following System P rules, here written as GPL_c formulas, can be derived from **(Ax₁)**–**(Ax₃)**:

$$\text{(OR)} \quad c(\alpha \sim \gamma) \wedge c(\beta \sim \gamma) \rightarrow c(\alpha \vee \beta \sim \gamma)$$

$$\text{(CM)} \quad c(\alpha \sim \beta) \wedge c(\alpha \sim \gamma) \rightarrow c(\alpha \wedge \beta \sim \gamma)$$

$$\text{(CUT)} \quad c(\alpha \wedge \beta \sim \gamma) \wedge c(\alpha \sim \beta) \rightarrow c(\alpha \sim \gamma)$$

Furthermore, it is easy to see that the following System P rules also follow from **(Ax₁)**–**(Ax₃)**:

$$\text{(RE)} \quad c(\alpha \sim \alpha), \text{ if } \alpha \text{ is consistent}$$

(LLE) $c(\alpha \sim \beta) \rightarrow c(\alpha' \sim \beta)$, if $\models (\alpha \equiv \alpha')$

(RW) $c(\alpha \sim \beta) \rightarrow c(\alpha \sim \beta')$, if $\beta \models \beta'$

In other words, all the inference rules of System P are theorems in $\text{GPL}_{\sim}^{\text{core}}$. To enable reasoning about propositional combinations of default rules in GPL_c , we will also use the following axioms:

(WRM) $c(\alpha \sim \gamma) \wedge \neg c(\alpha \sim \neg \beta) \rightarrow c(\alpha \wedge \beta \sim \gamma)$

(INC) $c(\alpha \sim \beta) \rightarrow \neg c(\alpha \sim \neg \beta)$

Note that **(INC)** follows directly from **(Ax₄)**, i.e., the GPL_c axiom **(INC)** is a theorem in $\text{GPL}_{\sim}^{\text{core}}$. The same holds for **(WRM)**, which follows from Proposition ?? below. The notation for the axiom **(WRM)** was introduced in [60], where a logic called NP^+ is discussed, in which disjunctions and negations can also be expressed. Note that **(WRM)** is similar to, but different from the rational monotonicity rule considered in [61]. The latter allows to derive $\alpha \wedge \beta \sim \gamma$ as soon as $\alpha \sim \gamma$ holds and $\alpha \sim \neg \beta$ cannot be established. In contrast, **(WRM)** requires that the negation of $\alpha \sim \neg \beta$ can be derived. The axiom **(INC)** is needed to make inconsistencies explicit at the propositional meta-level.

For K a set of GPL_c formulas and Ψ a GPL_c formulas, we write $K \vdash_c \Psi$ to denote that Ψ can be derived from K using **(RE)**, **(LLE)**, **(RW)**, **(OR)**, **(CM)**, **(CUT)**, **(WRM)**, **(INC)**, the axioms of propositional logic and modus ponens.

Proposition 7. *Let Φ and Ψ be formulas in \mathcal{L}_c . Then $\Phi \vdash_c \Psi$ iff $\Phi \models_c \Psi$.*

Proof. The proof is presented in Appendix D □

A similar result was obtained in [62], where possibility theory was also used to give a semantics to disjunctions of conditionals, albeit in a different context. Other logics in which propositional combinations of conditionals can be expressed include the logic NP^+ from [60], which has a semantics based on infinitesimal probabilities, the approach from [63], which is based on a three-valued semantics of conditional objects, and Lewis' logic VA [64], whose sphere semantics has been related to comparative possibility relations in [65] and can thus be simulated in GPL in a similar way.

Compared to logics such as NP^+ and VA, the main advantage of GPL is that we are able to provide a more intuitive semantics. Another advantage is that GPL can be implemented using SAT solvers in a relatively straightforward way, which should enable very efficient reasoning about default rules. Finally, embedding a logic of conditionals in GPL has the advantage that it allows us to combine conditionals with other types of epistemic knowledge. For example, we can use $\neg \mathbf{II}_1(\beta)$ to express that β is an abnormal situation, and, e.g., use $\mathbf{N}_1(\alpha) \rightarrow \neg \mathbf{II}_1(\beta)$ to encode that if α has been observed then β should be considered abnormal, which is more cautious than $\alpha \sim \neg \beta$.

6. Non-monotonic logic programming in GPL

The ability of GPL to model limitations on the knowledge of an agent makes it a natural framework for implementing various forms of non-monotonic reasoning. Section 5.4 already explained how to capture reasoning with exception-tainted rules. In the following, we show how the semantics of answer set programs can also be naturally captured in GPL. Section 6.2 then shows a close correspondence between GPL and the logic of minimal belief and negation as failure [66]. In particular, we obtain that the notion of minimality that is required in the latter logic is less demanding than the principle of minimal specificity in GPL.

6.1. Casting answer set programs in GPL

Consider the GPL formula $\mathbf{\Pi}_1(a) \rightarrow \mathbf{N}_1(b)$. Intuitively, this formula allows us to reason about the absence of information: as long as there is no reason to believe that a is false, we assume that b is necessarily true. Note, however, that $\mathbf{\Pi}_1(a) \rightarrow \mathbf{N}_1(b)$, which is equivalent to $\mathbf{N}_{\frac{1}{k}}(\neg a) \vee \mathbf{N}_1(b)$, has two minimally specific e-models: the e-models π_a^* and π_b defined as follows

$$\pi_a^*(\omega) = \begin{cases} \frac{k-1}{k} & \text{if } \omega \models a \\ 1 & \text{otherwise} \end{cases} \quad \pi_b(\omega) = \begin{cases} 0 & \text{if } \omega \models \neg b \\ 1 & \text{otherwise} \end{cases}$$

Note that π_b is Boolean in the sense that $\pi_b(\omega) \in \{0, 1\}$ for every $\omega \in \Omega$, whereas π_a^* is not, if $k \geq 2$. It turns out that in general, if we restrict our attention to minimally specific Boolean e-models of GPL formulas of this type with $k = 2$, we obtain a semantics for reasoning from the absence of information, which captures the stable model semantics of answer set programs. As we shall see the condition $k > 1$ is crucial. If $k = 1$ then $\mathbf{\Pi}_1(a) \rightarrow \mathbf{N}_1(b)$ is equivalent to $\mathbf{N}_1(\neg a) \vee \mathbf{N}_1(b)$, which does not allow for nonmonotonicity and only corresponds to a GPL (or more specifically MEL) translation of Kleene logic implication [67]. For the remainder of this section, we will focus on GPL_2 , although all the results readily translate to GPL_k for any $k \geq 2$.

Recall that an answer set program is a set of rules of the form:

$$a_1 \vee \dots \vee a_n \leftarrow b_1 \wedge \dots \wedge b_m \wedge \text{not } c_1 \wedge \dots \wedge \text{not } c_\ell \quad (27)$$

Intuitively, this rule encodes the idea that if we have no knowledge that any of c_1, \dots, c_ℓ are true, then whenever we can derive b_1, \dots, b_m we will assume that one of a_1, \dots, a_n must be true as well. As suggested in [68], we can view such a rule as a constraint on the possible epistemic states that a given agent may have. Let $\text{Lit} = \text{At} \cup \{\neg a \mid a \in \text{At}\}$ be the set of literals in the language. Let $M \subseteq \text{Lit}$ be such that $\{a, \neg a\} \not\subseteq M$ for every $a \in \text{At}$. Such a set $M \subseteq \text{Lit}$ can intuitively be viewed as a partial model: $a \in M$ means that a is known to be true, $\neg a \in M$ means that a is known to be false, and $a, \neg a \notin M$ means that the truth value of a is unknown. The reduct P^M of an answer set program P w.r.t. a partial model M is defined as follows:

$$P^M = \{a_1 \vee \dots \vee a_n \leftarrow b_1 \wedge \dots \wedge b_m \mid M \cap \{c_1, \dots, c_\ell\} = \emptyset, \\ (a_1 \vee \dots \vee a_n \leftarrow b_1 \wedge \dots \wedge b_m \wedge \text{not } c_1 \wedge \dots \wedge \text{not } c_\ell) \in P\}$$

Note that the reduct P^M is free of the negation-as-failure operator “not”. The set M is called a model of a rule $a_1 \vee \dots \vee a_n \leftarrow b_1 \wedge \dots \wedge b_m$ iff $\{b_1, \dots, b_m\} \not\subseteq M$ or $M \cap \{a_1, \dots, a_n\} \neq \emptyset$. Furthermore, M is called a model of a set of rules P^M (without negation-as-failure) if it is a model of every rule in P^M . Finally, M is called an answer set of an answer set program P iff it is the (unique) minimal model of P^M w.r.t. set inclusion.

Several equivalent methods have been identified to characterize the semantics of answer set programs [69]. Most of these characterizations are based on some kind of fixpoint construction. In the definition above, this is captured by the reduct, which fixes the interpretation of all literals under the scope of the “not” operator.

Using GPL, however, we can semantically characterize answer set programs purely in terms of minimally specific e-models. Specifically, given an answer set program P , we let K_P be the GPL knowledge base obtained by translating each rule of the form (27) into the following formula:

$$\mathbf{N}_1(b_1) \wedge \dots \wedge \mathbf{N}_1(b_m) \wedge \mathbf{\Pi}_1(\neg c_1) \wedge \dots \wedge \mathbf{\Pi}_1(\neg c_\ell) \rightarrow \mathbf{N}_1(a_1) \vee \dots \vee \mathbf{N}_1(a_n) \quad (28)$$

In other words, the body of a rule of the form (27) is satisfied if the agent knows each b_i with maximal certainty and moreover the agent considers $\neg c_j$ fully possible for each j . Note that $\mathbf{\Pi}_1(\neg c_j)$ is equivalent to $\neg \mathbf{N}_{\frac{1}{2}}(c_j)$ in GPL_2 .

The transformation in (28) is by itself not enough, as ASP is based on the idea of forward chaining and in particular does not allow contrapositive reasoning (e.g., from the rule $a \leftarrow b$ and the fact $\neg a$ we should not derive $\neg b$). To see how forward chaining could be enforced using GPL, first note that there are three ways in GPL_2 in which the formula (28) can be satisfied by a minimally specific e-model π of Θ_P :

1. one of the meta-literals $\mathbf{N}_1(b_i)$ is not satisfied by π ;
2. one of the meta-literals $\mathbf{\Pi}_1(\neg c_i)$ is not satisfied by π , i.e., $\mathbf{N}_{\frac{1}{2}}(c_i)$ is satisfied by π ;
3. one of the meta-literals $\mathbf{N}_1(a_i)$ is satisfied by π .

The first case intuitively corresponds to an answer set which does not include b_i , i.e., to a situation in which the rule (27) does not apply. The third case intuitively corresponds to an answer set in which a_i has been included to make the rule (27) satisfied, i.e., to a situation in which a_i has been derived using (non-deterministic) forward chaining. The second case, however, intuitively corresponds to a contrapositive inference, i.e., (27) has been satisfied by making c_i true. The latter inference is not allowed in ASP and the second case should thus be excluded. To this end, we take advantage of the fact that it is only in the second case that certainty degrees other than 0 or 1 are needed. Note that here we do not use degrees for modelling uncertainty, but intuitively for differentiating between literals that are assumed to be true and literals that can effectively be derived. In particular, it turns out that answer sets correspond to the minimally specific e-models of Θ_P in which only the certainty degrees 0 and 1 occur. Formally, the requirement that only these certainty degrees occur is encoded by the GPL formula Φ , defined as follows:

$$\Phi = \bigwedge_{a \in At} \mathbf{N}_1(a) \vee \mathbf{N}_1(\neg a) \vee (\mathbf{\Pi}_1(a) \wedge \mathbf{\Pi}_1(\neg a)) \quad (29)$$

The formula Φ expresses that for every atom a , the agent is either fully certain about the truth value of a (in which case $\mathbf{N}_1(a) \vee \mathbf{N}_1(\neg a)$ holds) or the agent is completely ignorant about a (in which case $\mathbf{\Pi}_1(a) \wedge \mathbf{\Pi}_1(\neg a)$ holds). It turns out that the answer sets of P correspond to the minimally specific e-models of Θ_P that satisfy Φ . In particular, we have the following correspondence.

Proposition 8. *Let P be an answer set program and let K_P be the corresponding GPL₂ knowledge base. It holds that P has a consistent answer set iff*

$$K_P \models_b \bigwedge_{a \in At} \mathbf{N}_1(a) \vee \mathbf{N}_1(\neg a) \vee (\mathbf{\Pi}_1(a) \wedge \mathbf{\Pi}_1(\neg a)) \quad (30)$$

Furthermore, it holds that l is included in at least one answer set of P iff

$$K_P \models_b \left(\bigwedge_{a \in At} \mathbf{N}_1(a) \vee \mathbf{N}_1(\neg a) \vee (\mathbf{\Pi}_1(a) \wedge \mathbf{\Pi}_1(\neg a)) \right) \wedge \mathbf{N}_1(l) \quad (31)$$

Finally, it holds that l is included in all answer sets of P iff

$$K_P \models_c \left(\bigwedge_{a \in At} \mathbf{N}_1(a) \vee \mathbf{N}_1(\neg a) \vee (\mathbf{\Pi}_1(a) \wedge \mathbf{\Pi}_1(\neg a)) \right) \rightarrow \mathbf{N}_1(l) \quad (32)$$

Proof. The proof is presented in Appendix E. □

Note that this result does not hold in GPL₁. Indeed, for $k = 1$, like for instance in [28], we have that $\bigwedge_{a \in At} \mathbf{N}_1(a) \vee \mathbf{N}_1(\neg a) \vee (\mathbf{\Pi}_1(a) \wedge \mathbf{\Pi}_1(\neg a))$ is a tautology. This explains why a similar characterization is not possible in autoepistemic logic, or other modal logics which rely on Boolean certainty degrees only. In contrast, equilibrium logic [70] does allow a similar characterization, by using a Boolean valuation in two worlds (called *here* and *there*) instead of intermediary certainty degrees. The advantage of GPL over equilibrium logic is that the epistemic interpretation of formulas is explicit, which make them easier to interpret intuitively (although this comes at the cost of a less concise syntax); we refer to [8] for a more detailed discussion on the relation between GPL and equilibrium logic. Note that while Fariñas et al. [71] have proposed a characterization of equilibrium logic in modal logic, this characterization does not highlight the intuitive meaning of equilibrium logic formulas, from an epistemic reasoning point of view. Recently, the same authors [72] have proposed an epistemic equilibrium logic. Again, however, the aim of this logic is not to provide an intuitive method for epistemic reasoning, but to generalize epistemic specifications [73, 74], a generalization of ASP which allows a new type of literal Kl in the body of rules, intuitively stating that l is true in all answer sets.

6.2. Logic of minimal belief and negation as failure

The characterization of ASP using GPL easily allows us to generalize the stable model semantics to a larger class of logic programs. For example, we could readily provide a semantics for disjunctions of ASP rules, we could allow negation as failure to appear in the

head of a rule, or use expressions of the form $\mathbf{N}_1(\alpha)$ where α can be an arbitrary propositional expressions instead of only a literal (see [68] for a more elaborate discussion on the latter point).

There are a number of existing logics which can similarly be used to provide a semantics for negation-as-failure in a more general setting. One of the simplest and earliest of these logics is Lifschitz' logic of minimal belief and negation as failure (MBNF) [66]. The language of this logic is the usual propositional language, extended with two modalities B and “not”. The semantics are defined w.r.t. triples of the form (ω, S^b, S^n) , where $\omega \in \Omega$ is a propositional interpretation and $S^b, S^n \subseteq \Omega$ are sets of propositional interpretations:

- For an atom a , $(\omega, S^b, S^n) \models_{MBNF} a$ iff $\omega \models a$;
- $(\omega, S^b, S^n) \models_{MBNF} \neg\psi$ iff $(\omega, S^b, S^n) \not\models_{MBNF} \psi$;
- $(\omega, S^b, S^n) \models_{MBNF} \psi \wedge \phi$ iff $(\omega, S^b, S^n) \models_{MBNF} \psi$ and $(\omega, S^b, S^n) \models_{MBNF} \phi$;
- $(\omega, S^b, S^n) \models_{MBNF} B\psi$ iff $(\omega', S^b, S^n) \models_{MBNF} \psi$ for every $\omega' \in S^b$ (i.e., $S^b \subseteq \llbracket \psi \rrbracket$ if ψ is a propositional formula);
- $(\omega, S^b, S^n) \models_{MBNF} \text{not } \psi$ iff there exists an $\omega' \in S^n$ such that $(\omega', S^b, S^n) \not\models_{MBNF} \psi$ (i.e., $S^n \not\subseteq \llbracket \psi \rrbracket$ if ψ is a propositional formula)

Intuitively, $B\psi$ is true if ψ is known to be true, i.e., we can think of S^b as the set of worlds that the agent considers possible. Intuitively $\text{not } \psi$ is true unless ψ is known to be true, where we instead consider S^n as the set of worlds that the agent considers possible. The satisfaction relation \models does not require any constraints on the relationship between S^b and S^n , although as we will see below, in models we will have that $S^b = S^n$. So, S^b is used to evaluate the “B” modalities and S^n is used to evaluate the “not” modalities. The use of two separate epistemic states can thus be thought of as a technical trick to encode a notion of minimality. As usual, $\psi \vee \phi$ is seen as an abbreviation of $\neg(\neg\psi \wedge \neg\phi)$ and $\psi \rightarrow \phi$ as an abbreviation of $\neg(\psi \wedge \neg\phi)$. Moreover, (ω, S^b, S^n) satisfies a set of MBNF formulas K iff it satisfies every formula in that set. A structure (ω, S) is called a model of a formula ψ iff

1. $(\omega, S, S) \models_{MBNF} \psi$; and
2. $(\omega', S', S) \not\models_{MBNF} \psi$ for any $S' \supset S$ and any propositional interpretation ω'

This second condition essentially plays a similar role to the notion of minimal specificity in GPL (and the notion of h-minimality in equilibrium logic). Note that the fact that modalities can be nested in MBNF does not really increase its expressive power, as e.g., $B(B(\psi))$ and $B(\psi)$ are equivalent (i.e., are satisfied by the same triples). Let us now consider the restriction of the language of MBNF to formulas without nested modalities, in which all atomic formulas occur within the scope of a modality. Let us refer to this fragment as MBNF_1 . This fragment is particularly interesting, because it can be used to define the semantics of answer set programming, in a way which is similar to the characterization in GPL from Proposition 8. In particular, consider an ASP rule of the following form

$$a_1 \vee \dots \vee a_n \leftarrow b_1 \wedge \dots \wedge b_m \wedge \text{not } c_1 \wedge \dots \wedge \text{not } c_\ell$$

The corresponding formula in MBNF_1 is given by

$$B(b_1) \wedge \dots \wedge B(b_m) \wedge \text{not}(c_1) \wedge \dots \wedge \text{not}(c_\ell) \rightarrow B(a_1) \vee \dots \vee B(a_n)$$

Lifschitz showed the following result.

Proposition 9. [66] *Let P be an answer set program and let K be the corresponding MBNF knowledge base. It holds that a set of literals M is a (consistent) answer set of P iff there exists a model (ω, S) of K such that $S = \llbracket M \rrbracket$.*

We will now show that MBNF_1 is related to, but different from GPL_2 . In particular, with each MBNF_1 knowledge base Θ , the corresponding GPL knowledge base $g(\Theta)$ is obtained by replacing each occurrence of $B(\alpha)$ by $\mathbf{N}_1(\alpha)$ and each occurrence of $\text{not}(\alpha)$ by $\mathbf{\Pi}_1(\neg\alpha)$. Note that for MBNF_1 formulas, we can identify models with sets $S \subseteq \Omega$ since whenever (ω, S) is also a model of an MBNF_1 formula, then (ω', S) is a model for any interpretation ω' . For the same reason, we will write (S^b, S^n) instead of (ω, S^b, S^n) when the choice of ω is irrelevant.

In the following, for any $S \subseteq \Omega$, we define the possibility distribution π_S as $\pi_S(\omega) = 1$ if $\omega \in S$ and $\pi_S(\omega) = 0$ otherwise.

Proposition 10. *Let Θ be an MBNF_1 knowledge base and let $g(\Theta)$ be the corresponding GPL_2 knowledge base. Let $S \subseteq \Omega$. Then $(S, S) \models_{\text{MBNF}} \Theta$ iff $\pi_S \models_{\text{GPL}} g(\Theta)$.*

Proof. For any propositional formula α , we have that $(S, S) \models_{\text{MBNF}} B(\alpha)$ iff α is true in every $\omega \in S$ iff $\pi_S \models_{\text{GPL}} \mathbf{N}_1(\alpha)$. Similarly, we have $(S, S) \models_{\text{MBNF}} \text{not}(\alpha)$ iff α is false in some $\omega \in S$ iff $\pi_S \models_{\text{GPL}} \mathbf{\Pi}_1(\neg\alpha)$. It follows that $(S, S) \models_{\text{MBNF}} \Theta$ iff $\pi_S \models_{\text{GPL}} g(\Theta)$. \square

Moreover, if π_S is a minimally specific e-model of $g(\Theta)$ then S is obviously a model of Θ . However, we do not have that every model S of an MBNF_1 formula Θ corresponds to a minimally specific e-model of $g(\Theta)$, as is illustrated by the following example.

Example 5. *We consider an example with only one atom a and we denote $\Omega = \{\omega_a, \omega_{\neg a}\}$, where ω_a is the interpretation which makes a true and $\omega_{\neg a}$ the interpretation which makes a false. Let $\psi = B(a) \vee \text{not}(a)$. Then $g(\psi) = \mathbf{N}_1(a) \vee \mathbf{\Pi}_1(\neg a)$ has only one minimally specific e-model π , defined by $\pi(\omega_a) = \pi(\omega_{\neg a}) = 1$. Accordingly, $S_1 = \{\omega_a, \omega_{\neg a}\}$ is a model of ψ ; indeed we have $(S_1, S_1) \models_{\text{MBNF}} \text{not}(a)$ and S_1 does not have any supersets since ω_a and $\omega_{\neg a}$ are the only interpretations. However, we show that $S_2 = \{\omega_a\}$ is also a model of ψ . First note that $(S_2, S_2) \models_{\text{MBNF}} B(a)$ hence we already have $(S_2, S_2) \models_{\text{MBNF}} \psi$. To show that S_2 is a model, it suffices to show that $(S_1, S_2) \not\models_{\text{MBNF}} \psi$ for the only superset S_1 , which is easy to verify. Indeed S_1 is not in $\llbracket a \rrbracket$, which violates $B(a)$, and $S_2 = \llbracket a \rrbracket$, which violates $\text{not}(a)$.*

This means that the notion of minimal specificity in GPL is more demanding than the notion of minimality imposed on models in MBNF. In [8], we obtained a similar result when comparing GPL to equilibrium logic [70]. This discrepancy especially seems to arise for theories which correspond to logic programs with negation-as-failure in the head. It is well-known that in the presence of negation-as-failure in the head, most semantics that cover

such cases lead to answer sets for which minimality no longer holds. While this has been advocated in e.g., [75] as a useful feature to encode particular constructs, such as inclusive disjunction, this behaviour remains somewhat counter-intuitive in a setting where minimal commitment is one of the main guiding principles.

The encoding of MBNF_1 in GPL is similar in spirit to the encoding of equilibrium logic that we proposed in [8]. In particular, the GPL encoding of an equilibrium logic theory is also such that (some) stable models correspond to those minimally specific e-models π in which $\pi(\omega) \in \{0, 1\}$ for each $\omega \in \Omega$. However, the encoding of MBNF_1 is considerably more intuitive. This is a consequence of the fact that the modalities in MBNF have a clear epistemic flavor, which does not seem to be the case for the connectives in equilibrium logic.

7. Computational complexity

In this section, we will consider the computational complexity of the main reasoning tasks for GPL. The modalities $\Delta_\lambda, \nabla_\lambda, \Delta_\lambda^X, \nabla_\lambda^X$ were introduced in Section 4 as abbreviations of formulas that only contain modalities of the form \mathbf{N}_λ . However, without these abbreviations, formulas may be exponentially longer, and as the computational complexity of reasoning tasks is expressed in function of the size of the formulas in the knowledge base, this has an impact on the complexity results. In other words, while these modalities have been introduced as abbreviations, they are not considered part of the language of GPL for the complexity results. We then also consider the variants GPL^Δ and GPL_R^Δ , in which these modalities are assumed to be included in the language:

GPL formulas are formulas in which all modalities are of the form \mathbf{N}_λ or $\mathbf{\Pi}_\lambda$.

GPL $^\Delta$ formulas are formulas in which also modalities of the form Δ_λ and ∇_λ are allowed.

GPL $_R^\Delta$ formulas are formulas in which moreover modalities of the form Δ_λ^X and ∇_λ^X are allowed.

Table 2 provides an overview of the complexity results that we will establish. In addition to the results from Table 2, we will also show that satisfiability checking in $\text{GPL}_{\succ}^{\text{core}}$ and GPL_{\succ} are NP-complete (and thus that entailment checking is coNP-complete), but the brave and cautious inference relations will not be considered in this case, as the notion of minimal specificity is not well-defined for these logics (e.g. $a \succ \perp$ does not have a minimally specific model).

Recall that a decision problem is in Σ_i^P ($i > 1$) if it can be solved in polynomial time on a non-deterministic Turing machine using a Σ_{i-1}^P -oracle, where $\Sigma_1^P = \text{NP}$. A decision problem is in Π_i^P if its complement is in Σ_i^P . A decision problem is in Θ_2^P if it can be solved in polynomial time on a deterministic Turing machine, by making a logarithmic number of calls to an NP-oracle.

7.1. Complexity of reasoning about GPL formulas

Proposition 11. *The problem of deciding whether a GPL formula is satisfiable is NP-complete (w.r.t. the size of the formula).*

Table 2: Overview of the complexity results.

	\models	\models_b	\models_c
GPL	coNP	Σ_2^P	Π_2^P
GPL $^\Delta$	Θ_2^P	Σ_2^P	Π_2^P
GPL $^\Delta_R$	Π_3^P	Σ_4^P	Π_4^P

Proof. Hardness follows straightforwardly from the NP-completeness of satisfiability in propositional logic. In particular, note that the propositional formula α is satisfiable iff the GPL formula $\mathbf{\Pi}_1(\alpha)$ is satisfiable.

We now propose an NP procedure for checking the satisfiability of an arbitrary GPL formula Φ . Each GPL formula Φ is equivalent to a disjunction of meta-terms, and it is sufficient that one of these terms is satisfiable. In polynomial time, we can guess such a term:

$$\mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n) \wedge \mathbf{\Pi}_{\mu_1}(\beta_1) \wedge \dots \wedge \mathbf{\Pi}_{\mu_m}(\beta_m) \quad (33)$$

We know that $\mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n)$ has a unique least specific e-model π if $\alpha_1 \wedge \dots \wedge \alpha_n$ is satisfiable. All that we need to check is whether this is the case, and whether $\Pi(\beta_i) \geq \mu_i$ for each i , with Π the possibility measure induced by π . In other words, there are two conditions which we need to check. First, the following formula needs to be consistent:

$$\alpha_1 \wedge \dots \wedge \alpha_n \quad (34)$$

Second, to check whether $\Pi(\beta_i) \geq \mu_i$ in the least specific model of $\mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n)$, we need to verify that β_i has a model ω such that $\omega \models \alpha_j$ holds for all j where $\lambda_j \geq \nu(\theta)$. In other words, we need to verify for each β_i that the following formula is consistent:

$$\bigwedge \{\alpha_i \mid \lambda_i \geq \nu(\theta)\} \wedge \beta_j \quad (35)$$

To check satisfiability in NP, when we guess the term (33), we can also guess an interpretation for each of these SAT instances and verify that they are indeed models of the corresponding propositional formulas. \square

Corollary 1. *The problem of deciding whether $\Phi \models \Psi$, with Φ and Ψ GPL formulas, is coNP-complete.*

Proof. This follows immediately from the observation that $\Phi \models \Psi$ holds iff $\Phi \wedge \neg\Psi$ is not satisfiable. \square

From the proof it is immediately clear that the same complexity results hold in the case of MEL, i.e., when only the certainty levels 0 and 1 are used. In other words, there is no penalty, in terms of computational complexity, for allowing more certainty levels.

The proof of Proposition 11 suggests a way to reason with GPL formulas using standard SAT solvers. In particular, let Φ be a GPL formula in which no negations occur at the meta-level. Since $\models_{GPL} \mathbf{N}_\lambda(\alpha \wedge \beta) \equiv \mathbf{N}_\lambda(\alpha) \wedge \mathbf{N}_\lambda(\beta)$, we can also assume w.l.o.g. that every meta-literal of the form $\mathbf{N}_\lambda(\alpha)$ is such that α is a disjunction of literals and that every meta-literal of the form $\mathbf{\Pi}_\lambda(\alpha)$ is such that α is a conjunction of literals. Let $f(\Phi)$ be the propositional formula which is obtained from Φ by replacing every meta-literal of the form $\mathbf{N}_\lambda(\alpha)$ by a fresh atom $a(\alpha, \lambda)$ and every meta-literal of the form $\mathbf{\Pi}_\lambda(\alpha)$ by a fresh atom $b(\alpha, \lambda)$. The SAT instance Θ corresponding with Φ contains the formula $f(\Phi)$ as well as the following formulas, involving fresh atomic formulas of the form $x^{(\beta, \mu)}$, for each meta-literal $\mathbf{\Pi}_\mu(\beta)$. Specifically, for each meta-literal of the form $\mathbf{\Pi}_\mu(\beta)$ we add:

$$b(\beta, \mu) \rightarrow a_1^{(\beta, \mu)} \wedge \dots \wedge a_n^{(\beta, \mu)} \wedge \neg b_1^{(\beta, \mu)} \wedge \dots \wedge \neg b_m^{(\beta, \mu)} \quad (36)$$

where we assume $\beta = a_1 \wedge \dots \wedge a_n \wedge \neg b_1 \wedge \dots \wedge \neg b_m$. Furthermore for each meta-literal of the form $\mathbf{N}_\lambda(\alpha)$ such that $\lambda \geq \nu(\mu)$ we add:

$$a(\alpha, \lambda) \wedge b(\beta, \mu) \rightarrow c_1^{(\beta, \mu)} \vee \dots \vee c_r^{(\beta, \mu)} \vee \neg d_1^{(\beta, \mu)} \vee \dots \vee \neg d_s^{(\beta, \mu)} \quad (37)$$

where we assume $\alpha = c_1 \vee \dots \vee c_r \vee \neg d_1 \vee \dots \vee \neg d_s$. Note that the formulas (36) and (37) are added to check the condition that the formula (35) has to be satisfiable for each meta-literal of the form $\mathbf{\Pi}_\mu(\beta)$ in the chosen meta-term. The other condition that we need to check is (34), which we can do by adding the following formulas for each meta-literal of the form $\mathbf{N}_\lambda(\alpha)$:

$$a(\alpha, \lambda) \rightarrow c_1 \vee \dots \vee c_r \vee \neg d_1 \vee \dots \vee \neg d_s \quad (38)$$

where we again assume that $\alpha = c_1 \vee \dots \vee c_r \vee \neg d_1 \vee \dots \vee \neg d_s$.

The following example illustrates the proposed reduction to SAT.

Example 6. Consider the following GPL_k formula for $k = 4$:

$$\Phi = \mathbf{N}_1(x \vee y) \wedge \mathbf{N}_{\frac{3}{4}}(\neg y) \wedge \mathbf{N}_{\frac{2}{4}}(\neg x \vee z) \wedge (\mathbf{\Pi}_{\frac{3}{4}}(\neg z) \vee \mathbf{\Pi}_{\frac{1}{4}}(x))$$

The resulting SAT instance Θ contains the following propositional formulas:

$$a(x \vee y, 1) \wedge a(\neg y, \frac{3}{4}) \wedge a(\neg x \vee z, \frac{2}{4}) \wedge (b(\neg z, \frac{3}{4}) \vee b(x, \frac{1}{4})) \quad (39)$$

$$b(\neg z, \frac{3}{4}) \rightarrow \neg z^{(\neg z, \frac{3}{4})} \quad (40)$$

$$b(x, \frac{1}{4}) \rightarrow x^{(x, \frac{1}{4})} \quad (41)$$

$$a(x \vee y, 1) \wedge b(\neg z, \frac{3}{4}) \rightarrow x^{(\neg z, \frac{3}{4})} \vee y^{(\neg z, \frac{3}{4})} \quad (42)$$

$$a(\neg y, \frac{3}{4}) \wedge b(\neg z, \frac{3}{4}) \rightarrow \neg y^{(\neg z, \frac{3}{4})} \quad (43)$$

$$a(\neg x \vee z, \frac{2}{4}) \wedge b(\neg z, \frac{3}{4}) \rightarrow \neg x^{(\neg z, \frac{3}{4})} \vee z^{(\neg z, \frac{3}{4})} \quad (44)$$

$$a(x \vee y, 1) \wedge b(x, \frac{1}{4}) \rightarrow x^{(x, \frac{1}{4})} \vee y^{(x, \frac{1}{4})} \quad (45)$$

$$a(x \vee y, 1) \rightarrow x \vee y \quad (46)$$

$$a(\neg y, \frac{3}{4}) \rightarrow \neg y \quad (47)$$

$$a(\neg x \vee z, \frac{2}{4}) \rightarrow \neg x \vee z \quad (48)$$

Recall that expressions such as $a(x \vee y, 1)$ and $z(\neg z, \frac{3}{4})$ are viewed as atomic formulas. From (39) we know that the atomic formulas $a(x \vee y, 1)$, $a(\neg y, \frac{3}{4})$ and $a(\neg x \vee z, \frac{2}{4})$ all need to be true, as well as $b(\neg z, \frac{3}{4})$ or $b(x, \frac{1}{4})$. However, from (40), together with (42)–(44) it follows that $b(\neg z, \frac{3}{4})$ cannot be satisfied in any model of Θ . This corresponds to the observation that the meta-term $\mathbf{N}_1(x \vee y) \wedge \mathbf{N}_{\frac{3}{4}}(\neg y) \wedge \mathbf{N}_{\frac{2}{4}}(\neg x \vee z) \wedge \mathbf{\Pi}_{\frac{3}{4}}(\neg z)$ is not satisfiable. On the other hand, the meta-term $\mathbf{N}_1(x \vee y) \wedge \mathbf{N}_{\frac{3}{4}}(\neg y) \wedge \mathbf{N}_{\frac{2}{4}}(\neg x \vee z) \wedge \mathbf{\Pi}_{\frac{1}{4}}(x)$ is satisfiable, and accordingly, it can readily be verified that Θ has a model in which the atom $b(x, \frac{1}{4})$ is true.

Proposition 12. *Let Φ be a GPL formula and let Θ be the associated SAT instance, constructed using the method explained above. It holds that Φ is satisfiable iff Θ is satisfiable.*

Proof. Suppose Φ has an e-model π . Then π satisfies some meta-term ϕ of the form (33). We define a partial interpretation ω as follows: $\omega \models a(\alpha, \lambda)$ iff the meta-term $\mathbf{N}_\lambda(\alpha)$ appears in ϕ and $\omega \models \neg a(\alpha, \lambda)$ otherwise; similarly $\omega \models b(\alpha, \lambda)$ iff the meta-term $\mathbf{\Pi}_\lambda(\alpha)$ appears in ϕ and $\omega \models \neg b(\alpha, \lambda)$ otherwise. Clearly, ω satisfies $f(\Phi)$. It remains to be shown that ω can be extended to an interpretation of all atomic formulas appearing in Θ , such that (36)–(38) are satisfied. However, the existence of such an extension follows directly from the fact that the formulas (35) and (34) are satisfiable if π is an e-model of the meta-term ϕ .

Conversely, if Θ has a model ω it is clear that there is some meta-term ϕ of the form (33) such that $\omega \models a(\alpha, \lambda)$ for every meta-literal $\mathbf{N}_\lambda(\alpha)$ appearing in ϕ and $\omega \models b(\alpha, \lambda)$ for every meta-literal $\mathbf{\Pi}_\lambda(\alpha)$ appearing in ϕ . As in the proof of Proposition 11 we find that ϕ has an e-model iff the formulas (35)–(34) are satisfied, and this follows straightforwardly from the fact that ω satisfies (36)–(38). \square

As already follows from the results in Section 6, reasoning about minimally specific e-models is more expensive than reasoning about what is true for all e-models of a GPL knowledge base. This stands in contrast to standard possibilistic logic, where both notions of entailment coincide.

Proposition 13. *Let Φ and Ψ be two GPL formulas. The problem of checking whether $\Phi \models_c \Psi$ is Π_2^P -complete (in the joint size of Φ and Ψ).*

Proof. hardness Consider the following QBF formula:

$$\psi = \forall x_1, \dots, x_n \cdot \exists y_1, \dots, y_m \cdot \phi(x_1, \dots, x_n, y_1, \dots, y_m)$$

In the following, we will abbreviate such formulas as $\forall X \exists Y \cdot \phi(X, Y)$ where $X = \{x_1, \dots, x_n\}$ and $Y = \{y_1, \dots, y_m\}$. We show that checking the validity of ψ can be

reduced to the problem of checking whether $\Phi \models_c \Psi$ for Φ and Ψ GPL formulas. Specifically, we choose Φ and Ψ as follows:

$$\begin{aligned}\Phi &= (\mathbf{N}_1(x_1) \vee \mathbf{N}_1(\neg x_1)) \wedge \dots \wedge (\mathbf{N}_1(x_n) \vee \mathbf{N}_1(\neg x_n)) \\ \Psi &= \mathbf{\Pi}_1(\phi(x_1, \dots, x_n, y_1, \dots, y_m))\end{aligned}$$

Let $\pi : \Omega \rightarrow [0, 1]$ be a minimally specific e-model of Ψ . Let us define $L \subseteq \{x_1, \dots, x_n, \neg x_1, \dots, \neg x_n\}$ as the set of literals that are known to be true in the epistemic state π :

$$L = \{x_i \mid \pi \models \mathbf{N}_1(x_i)\} \cup \{\neg x_i \mid \pi \models \mathbf{N}_1(\neg x_i)\}$$

It is clear, by definition of Ψ , that for every x_i either $\pi \models \mathbf{N}_1(x_i)$ or $\pi \models \mathbf{N}_1(\neg x_i)$, and as a result either $x_i \in L$ or $(\neg x_i) \in L$. It is also clear, by construction of L , that x_i and $\neg x_i$ cannot both be in L . In other words L defines a propositional interpretation over $\{x_1, \dots, x_n\}$. Conversely, each propositional interpretation over $\{x_1, \dots, x_n\}$ will correspond to some minimally specific e-model of Ψ .

Because π was assumed to be minimally specific, every interpretation ω which is consistent with L will be such that $\pi(\omega) = 1$. In particular, if there exists a model ω of $\phi(x_1, \dots, x_n, y_1, \dots, y_m)$ which is compatible with L , it will satisfy $\pi(\omega) = 1$ and thus $\pi \models \mathbf{\Pi}_1(\phi(x_1, \dots, x_n, y_1, \dots, y_m))$.

The QBF ψ is valid if and only if such a model ω exists for every choice of L . In other words, we have that ψ is valid iff $\Phi \models_c \Psi$.

membership Follows from the membership result in Proposition 15 below. □

To characterize the complexity of brave reasoning, note that $\Phi \models_c \Psi$ iff it is not the case that $\Phi \models_b \neg\Psi$. Hence we immediately get the following result.

Corollary 2. *Let Φ and Ψ be two GPL formulas. The problem of checking whether $\Phi \models_b \Psi$ is Σ_2^P -complete (in the joint size of Φ and Ψ).*

Given this complexity result, it is clear that no polynomial transformation to SAT will allow us to check $\Phi \models_c \Psi$ or $\Phi \models_b \Psi$, unless the polynomial hierarchy collapses. However, it is straightforward to reduce these entailment queries to QBFs, using a translation similar to the proposed SAT translation for checking the consistency of GPL formulas. In this way, we can use QBF solvers for reasoning about the minimally specific models of a GPL knowledge base.

7.2. Complexity of reasoning about GPL^Δ formulas

We now consider GPL^Δ formulas, i.e., formulas in which also meta-literals of the form $\Delta_\lambda(\alpha)$ and $\nabla_\lambda(\alpha)$ can occur.

Proposition 14. *The problem of deciding whether $\Phi \models \Psi$, for Φ and Ψ two GPL^Δ formulas, is Θ_2^P -complete (in the joint size of Φ and Ψ).*

Proof. hardness A standard Θ_2^P -complete problem is the following. Let ϕ_1, \dots, ϕ_n be propositional formulas. Decide whether the smallest i for which ϕ_i is unsatisfiable is an odd number.

Without loss of generality, we can assume that n is odd (as otherwise we could simply omit ϕ_n). Now consider the following GPL $^\Delta$ formula:

$$\begin{aligned} \Psi = & \neg \mathbf{\Pi}_1(\phi_1) \vee (\mathbf{\Pi}_1(\phi_1) \wedge \mathbf{\Pi}_1(\phi_2) \wedge \neg \mathbf{\Pi}_1(\phi_3)) \\ & \vee \dots \vee (\mathbf{\Pi}_1(\phi_1) \wedge \dots \wedge \mathbf{\Pi}_1(\phi_{n-1}) \wedge \neg \mathbf{\Pi}_1(\phi_n)) \end{aligned}$$

We show that $\mathbf{\Delta}_1(\top) \models \Psi$ iff the smallest i for which ϕ_i is unsatisfiable is odd. Clearly, $\mathbf{\Delta}_1(\top)$ has exactly one e-model, which is the possibility distribution π^* for which every world is fully possible, i.e., $\pi^*(\omega) = 1$ for every $\omega \in \Omega$. For a propositional formula ϕ , we then have $\mathbf{\Pi}^*(\phi) = 1$ iff $\llbracket \phi \rrbracket \neq \emptyset$. In other words, π^* will be an e-model of Ψ iff ϕ_1 is not satisfiable, or ϕ_1 and ϕ_2 are satisfiable but not ϕ_3 , etc.

membership It is well-known that the class Θ_2^P coincides with the class of decision problems which can be solved in polynomial time on a deterministic Turing machine by using a polynomial number of parallel queries to an NP-oracle, i.e., such that the result of one query to the NP-oracle cannot be used to formulate another query to the NP-oracle [76]. Surprisingly, allowing two rounds of parallel queries does not lead to an increased complexity ([77], Theorem 9). We will show that $\Phi \models \Psi$ can be decided in this way, thus proving membership in Θ_2^P .

Since $\Phi \models \Psi$ holds iff $\Phi \wedge \neg \Psi$ is unsatisfiable, it is sufficient to show that satisfiability checking of GPL $^\Delta$ formulas is in Θ_2^P . Let Ψ be a GPL $^\Delta$ formula. Without loss of generality, we can assume that no implications occur in Ψ and that all negations occur inside a modality, i.e., the meta-literals in Ψ are connected using conjunction and disjunction only.

Assume that the meta-literals occurring in Ψ are:

$$\begin{aligned} & \mathbf{N}_{\frac{1}{k}}(\alpha_1^1), \dots, \mathbf{N}_{\frac{1}{k}}(\alpha_{n_1}^1), \mathbf{N}_{\frac{2}{k}}(\alpha_1^2), \dots, \mathbf{N}_{\frac{2}{k}}(\alpha_{n_2}^2), \dots, \mathbf{N}_1(\alpha_1^k), \dots, \mathbf{N}_1(\alpha_{n_k}^k) \\ & \mathbf{\Pi}_{\frac{1}{k}}(\beta_1^1), \dots, \mathbf{\Pi}_{\frac{1}{k}}(\beta_{m_1}^1), \mathbf{\Pi}_{\frac{2}{k}}(\beta_1^2), \dots, \mathbf{\Pi}_{\frac{2}{k}}(\beta_{m_2}^2), \dots, \mathbf{\Pi}_1(\beta_1^k), \dots, \mathbf{\Pi}_1(\beta_{m_k}^k) \\ & \mathbf{\Delta}_{\frac{1}{k}}(\gamma_1^1), \dots, \mathbf{\Delta}_{\frac{1}{k}}(\gamma_{p_1}^1), \mathbf{\Delta}_{\frac{2}{k}}(\gamma_1^2), \dots, \mathbf{\Delta}_{\frac{2}{k}}(\gamma_{p_2}^2), \dots, \mathbf{\Delta}_1(\gamma_1^k), \dots, \mathbf{\Delta}_1(\gamma_{p_k}^k) \\ & \mathbf{\nabla}_{\frac{1}{k}}(\delta_1^1), \dots, \mathbf{\nabla}_{\frac{1}{k}}(\delta_{r_1}^1), \mathbf{\nabla}_{\frac{2}{k}}(\delta_1^2), \dots, \mathbf{\nabla}_{\frac{2}{k}}(\delta_{r_2}^2), \dots, \mathbf{\nabla}_1(\delta_1^k), \dots, \mathbf{\nabla}_1(\delta_{r_k}^k) \end{aligned}$$

Using a first round of parallel calls to an NP-oracle, we check $\gamma_i^u \models \alpha_j^v$ for all $1 \leq i \leq p_u$, $1 \leq j \leq n_v$, and $u + v \geq k + 1$. Note that the number of calls to the oracle is at most quadratic in the number of meta-literals appearing in Ψ .

Using the result of these oracle calls, we can decide the satisfiability of Ψ in NP, i.e., by making one additional call to the NP-oracle, as follows. Note that Ψ is equivalent to a disjunction of meta-terms. In polynomial time we may guess such a meta-term,

of the following form:

$$\Theta = \mathbf{N}_{v_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{v_n}(\alpha_n) \wedge \mathbf{\Pi}_{w_1}(\beta_1) \wedge \dots \wedge \mathbf{\Pi}_{w_m}(\beta_m) \\ \wedge \mathbf{\Delta}_{u_1}(\gamma_1) \wedge \dots \wedge \mathbf{\Delta}_{u_p}(\gamma_p) \wedge \mathbf{\nabla}_{z_1}(\delta_1) \wedge \dots \wedge \mathbf{\nabla}_{z_r}(\delta_r)$$

We will further refine the meta-literals of the form $\mathbf{\Pi}_{w_i}(\beta_i)$ and $\mathbf{\nabla}_{z_i}(\delta_i)$ in Θ . To refine a meta-literal of the form $\mathbf{\Pi}_{w_i}(\beta_i)$ we need to replace β_i by a more restrictive formula. To this end, for each β_i we guess a specific model $\omega_{\beta_i} \in \llbracket \beta_i \rrbracket$, and we define $\beta_i^* = \bigwedge_{\omega_{\beta_i} \models l} l$, i.e., β_i^* is chosen such that $\llbracket \beta_i^* \rrbracket = \{\omega_{\beta_i}\}$. It follows that $\models_{GPL} \mathbf{\Pi}_{w_i}(\beta_i^*) \equiv \mathbf{\Delta}_{w_i}(\beta_i^*)$.

To refine a meta-literal of the form $\mathbf{\nabla}_{z_i}(\delta_i)$, we need to replace δ_i with a less restrictive formula. In particular, we guess a world $\omega_{\delta_i} \notin \llbracket \delta_i \rrbracket$ and choose the formula δ_i^* such that $\llbracket \delta_i^* \rrbracket = \Omega \setminus \{\omega_{\delta_i}\}$. It then holds that $\models_{GPL} \mathbf{\nabla}_{z_i}(\delta_i^*) \equiv \mathbf{N}_{z_i}(\delta_i^*)$. Note that the size of the formulas β_i^* and δ_i^* is linear in the number of literals, hence we can indeed guess these formulas in polynomial time.

Clearly, the term Θ is satisfiable iff such refinements can be found that make the following term Θ^* satisfiable

$$\Theta^* = \mathbf{N}_{v_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{v_n}(\alpha_n) \wedge \mathbf{\Delta}_{w_1}(\beta_1^*) \wedge \dots \wedge \mathbf{\Delta}_{w_m}(\beta_m^*) \\ \wedge \mathbf{\Delta}_{u_1}(\gamma_1) \wedge \dots \wedge \mathbf{\Delta}_{u_p}(\gamma_p) \wedge \mathbf{N}_{z_1}(\delta_1^*) \wedge \dots \wedge \mathbf{N}_{z_r}(\delta_r^*)$$

Let π_* be the most specific possibility distribution satisfying

$$\mathbf{\Delta}_{w_1}(\beta_1^*) \wedge \dots \wedge \mathbf{\Delta}_{w_m}(\beta_m^*) \wedge \mathbf{\Delta}_{u_1}(\gamma_1) \wedge \dots \wedge \mathbf{\Delta}_{u_p}(\gamma_p)$$

and let π^* be the least specific possibility distribution satisfying

$$\mathbf{N}_{v_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{v_n}(\alpha_n) \wedge \mathbf{N}_{z_1}(\delta_1^*) \wedge \dots \wedge \mathbf{N}_{z_r}(\delta_r^*)$$

It is clear that every possibility distribution π from \mathcal{P}_k satisfying $\pi_*(\omega) \leq \pi(\omega) \leq \pi^*(\omega)$ is an e-model of Θ^* . To complete the proof, we need to show that it can be checked in polynomial time whether such a possibility distribution π exists. This is the case exactly when $\alpha_1 \wedge \dots \wedge \alpha_n \wedge \delta_1^* \wedge \dots \wedge \delta_r^*$ is satisfiable, and the following entailment relations are valid:

- $\beta_i^* \models \alpha_j$ for every i, j such that $w_i \geq \nu(v_j)$
- $\beta_i^* \models \delta_j^*$ for every i, j such that $w_i \geq \nu(z_j)$
- $\gamma_i \models \alpha_j$ for every i, j such that $u_i \geq \nu(v_j)$
- $\gamma_i \models \delta_j^*$ for every i, j such that $u_i \geq \nu(z_j)$

To verify satisfiability in NP (given the result of the first round of calls to the oracle), we can guess the term Θ^* and at the same time guess a model of $\alpha_1 \wedge \dots \wedge \alpha_n \wedge \delta_1^* \wedge \dots \wedge \delta_r^*$. We can use that model to verify the satisfiability of $\alpha_1 \wedge \dots \wedge \alpha_n \wedge \delta_1^* \wedge \dots \wedge \delta_r^*$. Moreover, the entailment relations of the form $\gamma_i \models \alpha_j$ can be verified by looking up the result

of the first round of calls to the NP oracle. To check $\beta_i^* \models \alpha_j$, since β^* only has one model w_{β_i} it suffices to check whether α_j is true in that model, which can clearly be done in polynomial time. Similarly, $\beta_i^* \models \delta_j^*$ can be decided in this way. In fact, in the latter case it suffices to check that $w_{\beta_i} \neq w_{\delta_j}$. Finally, to check $\gamma_i \models \delta_j^*$, we can equivalently check $\neg\delta_j^* \models \neg\gamma_i$. Since $\neg\delta_j^*$ has a unique model w_{δ_j} it suffices to check that γ_i is false in this model. \square

Thus we find that allowing the Δ modality causes a jump in complexity. It should be noted, however, that this increased complexity is the result of how \mathbf{N} and Δ interact. In particular, the NP-completeness result from Proposition 11 could straightforwardly be adapted to a logic where only Δ modalities are allowed, by taking advantage of the duality expressed in (15).

Proposition 15. *Let Φ and Ψ be two GPL^Δ formulas. The problem of checking whether $\Phi \models_c \Psi$ is Π_2^P -complete (in the joint size of Φ and Ψ).*

Proof. hardness Follows immediately from Proposition 13.

membership We now present a Σ_2^P algorithm for checking that Ψ is false in at least one minimally specific e-model of Φ .

The GPL formula Φ is equivalent to a disjunction of meta-terms. In polynomial time, we can guess such a meta-term. Moreover, as in the proof of Proposition 14, without loss of generality, we can assume that the only meta-literals which occur in this meta-term are of the form $\mathbf{N}_{\lambda_1}(\alpha_1)$ and $\mathbf{\Delta}_{\mu_1}(\beta_1)$, by refining any meta-literals of the form $\mathbf{\Pi}_{\lambda_1}(\alpha_1)$ and $\mathbf{\nabla}_{\mu_1}(\beta_1)$. Assume that we guess a meta-term of the following form:

$$\mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n) \wedge \mathbf{\Delta}_{\mu_1}(\beta_1) \wedge \dots \wedge \mathbf{\Delta}_{\mu_m}(\beta_m) \quad (49)$$

Using an NP-oracle we can check that this formula is consistent, verifying that $\alpha_1 \wedge \dots \wedge \alpha_n$ is consistent, and that $\beta_i \models \alpha_j$ whenever $\mu_i \geq \nu(\lambda_j)$. If the meta-term is consistent, it has a unique least specific e-model π_1 , which is the least specific e-model of $\mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n)$, noting that the latter corresponds to a standard possibilistic logic base.

Using the NP-oracle we can check that β is false in π_1 . In particular:

- a meta-literal $\mathbf{N}_{\delta_1}(\gamma_1)$ is satisfied by π_1 iff $\{\alpha_i \mid \lambda_i \geq \delta_1\} \models \gamma_1$;
- a meta-literal $\mathbf{\Pi}_{\delta_1}(\gamma_1)$ is falsified by π_1 iff $\{\alpha_i \mid \lambda_i \geq \nu(\delta_1)\} \models \neg\gamma_1$;
- a meta-literal $\mathbf{\Delta}_{\delta_1}(\gamma_1)$ is satisfied by π_1 iff $\gamma_1 \models \{\alpha_i \mid \lambda_i \geq \nu(\delta_1)\}$;
- a meta-literal $\mathbf{\nabla}_{\delta_1}(\gamma_1)$ is falsified by π_1 iff $\neg\gamma_1 \models \{\alpha_i \mid \lambda_i \geq \delta_1\}$.

What remains to be verified is that there does not exist another consistent meta-term which is an implicant of Φ and which has a least specific e-model π_2 that is strictly less specific than π_1 .

The check this, we define a GPL^Δ knowledge base Φ' as follows. Without loss of generality, we can assume that Φ is in negation normal form and in particular that there are no occurrences of negation or implication outside the modalities. Starting from Φ , for each meta-literal $\mathbf{N}_\delta(\gamma)$ that occurs, we test whether

$$\mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n) \models \mathbf{N}_\delta(\gamma) \quad (50)$$

If this is not the case, no e-model of $\mathbf{N}_\delta(\gamma)$ can be less specific than π_1 ; we then replace $\mathbf{N}_\delta(\gamma)$ by \perp . Furthermore, for each meta-literal $\nabla_\delta(\gamma)$ which occurs, we test whether $\neg\gamma \wedge \neg \bigwedge \{\alpha_i \mid \lambda_i \geq \delta\}$ is consistent. If not, we find that no e-model of $\nabla_\delta(\gamma)$ can be less specific than π_1 , and we replace $\nabla_\delta(\gamma)$ by \perp . If, on the other hand, $\neg\gamma \wedge \neg \bigwedge \{\alpha_i \mid \lambda_i \geq \delta\}$ is consistent, then we replace the meta-literal $\nabla_\delta(\gamma)$ by $\nabla_\delta(\gamma \vee \bigwedge \{\alpha_i \mid \lambda_i \geq \delta\})$, since the minimally specific e-models of the latter GPL formula are exactly those minimally specific e-models of $\nabla_\delta(\gamma)$ that are less specific than π_1 . Note that the resulting knowledge base Φ' is consistent, and in particular that (49) is an implicant of Φ' .

By replacing a meta-literal $\mathbf{N}_\delta(\gamma)$ or $\nabla_\delta(\gamma)$, we potentially reduce the set of e-models of the knowledge base. However, by construction, none of these e-models can be less specific than π_1 . Moreover, each minimally specific e-model of Φ' is either equal to π_1 or strictly less specific than π_1 . Therefore, we finally test whether $\Phi' \models \mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n)$. If this is the case, then none of the e-models of Φ' , and by extension of Φ , can be less specific than π_1 . On the other hand, if this is not the case, then Φ' has an e-model which is not a refinement of π (since any refinement of π_1 is also an e-model of $\mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n)$). By construction, Φ' then has an e-model which is strictly less specific than π_1 , which means that the guess in (49) did not induce a minimally specific e-model of Φ . □

Corollary 3. *Let Φ and Ψ be two GPL^Δ formulas. The problem of checking whether $\Phi \models_b \Psi$ is Σ_2^P -complete (in the joint size of Φ and Ψ).*

7.3. Complexity of reasoning about GPL_R^Δ formulas

Next we consider the complexity of reasoning in the presence of the restricted guaranteed possibility modality.

Proposition 16. *The problem of deciding whether $\Phi \models \Psi$, for Φ and Ψ two GPL_R^Δ formulas, is Π_3^P -complete (in the joint size of Φ and Ψ).*

Proof. We will prove that checking the satisfiability of a GPL_R^Δ formula is Σ_3^P -complete, from which the stated result readily follows.

Hardness Let $X \cup Y \cup Z$ be a partition of the set of atomic formulas. We can show that checking the validity of the QBF $\Psi = \exists X \forall Y \exists Z . \phi(X, Y, Z)$ is equivalent to checking whether the following formula is satisfiable:

$$\bigwedge_{x \in X} (\mathbf{N}_1(x) \vee \mathbf{N}_1(\neg x)) \wedge \mathbf{\Delta}_1^Y(\top) \wedge \mathbf{N}_1(\phi(X, Y, Z)) \quad (51)$$

Indeed, this formula is equivalent to a disjunction of meta-terms, each of which corresponds to an interpretation of the variables in X . Let $X_1 \cup X_2$ be a partition of X and let us consider the corresponding implicant of (51):

$$\left(\bigwedge_{x \in X_1} \mathbf{N}_1(x) \right) \wedge \left(\bigwedge_{x \in X_2} \mathbf{N}_1(\neg x) \right) \wedge \Delta_1^Y(\top) \wedge \mathbf{N}_1(\phi(X, Y, Z))$$

This is equivalent to

$$\Delta_1^Y(\top) \wedge \mathbf{N}_1 \left(\phi(X, Y, Z) \wedge \left(\bigwedge_{x \in X_1} x \right) \wedge \left(\bigwedge_{x \in X_2} \neg x \right) \right)$$

The latter formula is satisfiable iff every truth assignment of the variables in Y can be extended to a model of $\phi(X, Y, Z) \wedge (\bigwedge_{x \in X_1} x) \wedge (\bigwedge_{x \in X_2} \neg x)$. Clearly this means that (51) is satisfiable iff the QBF Ψ is valid. This means that satisfiability checking in GPL_R^Δ is Σ_3^P -hard, from which it follows that entailment checking is Π_3^P -hard.

Membership We provide a Σ_3^P procedure for verifying that a GPL_R^Δ formula Ψ is satisfiable. Similarly as in the proof of Proposition 14, we can guess an implicant of Ψ of the following form:

$$\Theta = \mathbf{N}_{v_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{v_n}(\alpha_n) \wedge \Delta_{w_1}^{X_1}(\beta_1) \wedge \dots \wedge \Delta_{w_m}^{X_m}(\beta_m)$$

where X_1, \dots, X_m are sets of atomic formulas. We give a Σ_2^P procedure for checking that Θ is not satisfiable. First verify whether $\alpha_1 \wedge \dots \wedge \alpha_n$ is satisfiable, using an NP oracle. If this is the case, select a β_i , guess a model ω of β_i , and verify using the NP-oracle that $\text{conj}_\omega^X \wedge \{\alpha_j \mid v_j \geq \nu(w_i)\}$ is inconsistent. It follows that checking the satisfiability of Θ is in Π_2^P , and can thus be done in constant time using a Σ_2^P -oracle. \square

Proposition 17. *The problem of deciding whether $\Phi \models_c \Psi$, for Φ and Ψ two GPL_R^Δ formulas, is Π_4^P -complete (in the joint size of Φ and Ψ).*

Proof. Hardness Let $X \cup Y \cup Z \cup U$ be a partition of the set of atomic formulas. We show that checking the validity of the QBF $\psi = \forall X \exists Y \forall Z \exists U . \phi(X, Y, Z, U)$ is equivalent to checking whether $\Phi \models_c \Psi$ where:

$$\begin{aligned} \Phi &= \left(\bigwedge_{x \in X} \mathbf{N}_1(x) \vee \mathbf{N}_1(\neg x) \right) \wedge \mathbf{N}_1(\phi(X, Y, Z, U)) \\ &\quad \wedge \left(\left(\bigwedge_{y \in Y} \mathbf{N}_1(y \vee a) \vee \mathbf{N}_1(\neg y \vee a) \right) \wedge \Delta_1^{Z \cup \{a\}}(\neg a) \right) \vee \mathbf{N}_1(a) \\ \Psi &= \neg \mathbf{N}_1(a) \end{aligned}$$

Indeed, first note that for every subset $X_0 \subset X$ (i.e., for every interpretation of X), Φ has a minimally specific e-model in which $\bigwedge_{x \in X_0} \mathbf{N}_1(x) \wedge \bigwedge_{x \notin X_0} \mathbf{N}_1(\neg x)$ is true. To see why this is the case, note that the least specific e-model π_1 of $\bigwedge_{x \in X_0} \mathbf{N}_1(x) \wedge \bigwedge_{x \notin X_0} \mathbf{N}_1(\neg x) \wedge \mathbf{N}_1(a)$ is an e-model of Φ and for any e-model π_2 of Φ which is strictly less specific than π_1 , it must be the case that $\pi_2 \models \bigwedge_{x \in X_0} \mathbf{N}_1(x) \wedge \bigwedge_{x \notin X_0} \mathbf{N}_1(\neg x)$. We thus find that $\Phi \models_c \Psi$ iff for every $X_0 \subset X$, it holds that $f(X_0) \models_c \Psi$, where:

$$\begin{aligned} f(X_0) = & \bigwedge_{x \in X_0} \mathbf{N}_1(x) \wedge \bigwedge_{x \notin X_0} \mathbf{N}_1(\neg x) \wedge \mathbf{N}_1(\phi(X, Y, Z, U)) \\ & \wedge \left(\left(\bigwedge_{y \in Y} \mathbf{N}_1(y \vee a) \vee \mathbf{N}_1(\neg y \vee a) \right) \wedge \Delta_1^{Z \cup \{a\}}(\neg a) \right) \vee \mathbf{N}_1(a) \end{aligned}$$

Note that $\models_{GPL} \Phi \equiv \bigvee_{X_0 \subseteq X} f(X_0)$. Now let us define $g(X_0)$ and $f(X_0, Y_0)$ for $X_0 \subseteq X$ and $Y_0 \subseteq Y$ as follows:

$$\begin{aligned} g(X_0) = & \bigwedge_{x \in X_0} \mathbf{N}_1(x) \wedge \bigwedge_{x \notin X_0} \mathbf{N}_1(\neg x) \wedge \mathbf{N}_1(\phi(X, Y, Z, U)) \wedge \mathbf{N}_1(a) \\ f(X_0, Y_0) = & \bigwedge_{x \in X_0} \mathbf{N}_1(x) \wedge \bigwedge_{x \notin X_0} \mathbf{N}_1(\neg x) \wedge \mathbf{N}_1(\phi(X, Y, Z, U)) \\ & \wedge \bigwedge_{y \in Y_0} \mathbf{N}_1(y \vee a) \wedge \bigwedge_{y \notin Y_0} \mathbf{N}_1(\neg y \vee a) \wedge \Delta_1^{Z \cup \{a\}}(\neg a) \end{aligned}$$

Note that $f(X_0) = g(X_0) \vee \bigvee_{Y_0 \subseteq Y} f(X_0, Y_0)$. If there exists a $Y_0 \subseteq Y$ such that $f(X_0, Y_0)$ is consistent, then clearly the least specific e-model of $f(X_0, Y_0)$ will be strictly less specific than any e-model of $g(X_0)$. Furthermore note that $f(X_0, Y_0) \models_c \neg \mathbf{N}_1(a)$ while $g(X_0) \models_c \mathbf{N}_1(a)$. In other words, we have $f(X_0) \models_c \neg \mathbf{N}_1(a)$ iff there exists a $Y_0 \subseteq Y$ such that $f(X_0, Y_0)$ is consistent. The latter condition will be satisfied iff for every $Z_0 \subseteq Z$ it holds that $\bigwedge_{x \in X_0} x \wedge \bigwedge_{x \notin X_0} \neg x \wedge \bigwedge_{y \in Y_0} y \wedge \bigwedge_{y \notin Y_0} \neg y \wedge \bigwedge_{z \in Z_0} z \wedge \bigwedge_{z \notin Z_0} \neg z \wedge \phi(X, Y, Z, U)$ is consistent. In other words, iff for every $Z_0 \subseteq Z$ there exists a $U_0 \subseteq U$ such that $X_0 \cup Y_0 \cup Z_0 \cup U_0$ defines a model of $\phi(X, Y, Z, U)$.

In summary we have that $\Phi \models_c \Psi$ iff for every X_0 it holds that $f(X_0) \models_c \Psi$, iff for every X_0 there exists a Y_0 such that $f(X_0, Y_0)$ is consistent, iff for every X_0 there exists a Y_0 such that for every Z_0 there exists a U_0 such that $X_0 \cup Y_0 \cup Z_0 \cup U_0$ defines a model of $\phi(X, Y, Z, U)$, iff the QBF ψ is valid.

Membership We give a Σ_4^P procedure for checking $\Phi \not\models_c \Psi$, from which the membership result immediately follows. As in the proof of Proposition 16, we can guess an implicant of Φ of the following form:

$$\Phi_0 = \mathbf{N}_{v_1}(\gamma_1) \wedge \dots \wedge \mathbf{N}_{v_n}(\gamma_n) \wedge \Delta_{w_1}^{X_1}(\delta_1) \wedge \dots \wedge \Delta_{w_m}^{X_m}(\delta_m)$$

where X_1, \dots, X_m are sets of atomic formulas. Using a Σ_2^P oracle, we can verify that Φ_0 is consistent, as in the proof of Proposition 16. Since the unique least specific e-model of Φ_0 is also the least specific e-model of $\mathbf{N}_{v_1}(\gamma_1) \wedge \dots \wedge \mathbf{N}_{v_n}(\gamma_n)$, using a Σ_2^P oracle, we can check in polynomial time that $\Phi_0 \not\models_c \Psi$. Indeed:

- The satisfaction of meta-literals of the form $\mathbf{N}_\lambda(\epsilon)$, $\mathbf{\Pi}_\lambda(\epsilon)$, $\mathbf{\Delta}_\lambda(\epsilon)$ and $\mathbf{\nabla}_\lambda(\epsilon)$ occurring in Ψ can be verified as in the proof of Proposition 15.
- To check whether $\mathbf{N}_{v_1}(\gamma_1) \wedge \dots \wedge \mathbf{N}_{v_n}(\gamma_n) \models_c \mathbf{\Delta}_\lambda^X(\epsilon)$, it suffices to check the validity of the following QBF:

$$\forall X . (\exists (At \setminus X) . \epsilon) \rightarrow (\exists (At \setminus X) . \bigwedge \{\gamma_i \mid v_i \geq \nu(\lambda)\})$$

This can be accomplished in constant time using a Σ_2^P oracle.

- To check whether $\mathbf{N}_{v_1}(\gamma_1) \wedge \dots \wedge \mathbf{N}_{v_n}(\gamma_n) \models_c \mathbf{\nabla}_\lambda^X(\epsilon)$, it suffices to check that $\mathbf{N}_{v_1}(\gamma_1) \wedge \dots \wedge \mathbf{N}_{v_n}(\gamma_n) \not\models_c \mathbf{\Delta}_\lambda^X(\neg\epsilon)$, since $\mathbf{N}_{v_1}(\gamma_1) \wedge \dots \wedge \mathbf{N}_{v_n}(\gamma_n)$ has a unique least specific e-model.

Finally, we give a Σ_3^P procedure for showing that the least specific e-model of Φ_0 is not a minimally specific e-model of Φ . In particular, we guess an implicant of Φ of the form:

$$\Phi_1 = \mathbf{N}_{u_1}(\epsilon_1) \wedge \dots \wedge \mathbf{N}_{u_s}(\epsilon_s) \wedge \mathbf{\Delta}_{z_1}^{Y_1}(\zeta_1) \wedge \dots \wedge \mathbf{\Delta}_{z_t}^{Y_t}(\zeta_t)$$

We can then verify using a Σ_2^P oracle that Φ_1 is consistent. Using an NP oracle, we can furthermore verify that $\mathbf{N}_{v_1}(\gamma_1) \wedge \dots \wedge \mathbf{N}_{v_n}(\gamma_n) \models \mathbf{N}_{u_1}(\epsilon_1) \wedge \dots \wedge \mathbf{N}_{u_s}(\epsilon_s)$ while $\mathbf{N}_{u_1}(\epsilon_1) \wedge \dots \wedge \mathbf{N}_{u_s}(\epsilon_s) \not\models \mathbf{N}_{v_1}(\gamma_1) \wedge \dots \wedge \mathbf{N}_{v_n}(\gamma_n)$, from which it follows that the least specific e-model of Φ_1 is strictly less specific than the least specific e-model of Φ_0 . \square

Corollary 4. *The problem of deciding whether $\Phi \models_b \Psi$, for Φ and Ψ two GPL_R^Δ formulas, is Σ_4^P -complete (in the joint size of Φ and Ψ).*

7.4. Complexity of reasoning in GPL_\succ and $\text{GPL}_\succ^{\text{core}}$

To characterize the complexity of satisfiability checking in $\text{GPL}_\succ^{\text{core}}$, we can take advantage of a straightforward reduction to GPL_k . First note that when only finitely many certainty degrees are considered, \succ can be introduced as an abbreviation in GPL_k :

$$\alpha \succ \beta = \bigvee_{i=1}^k (\mathbf{N}_{\frac{i}{k}}(\alpha) \wedge \neg \mathbf{N}_{\frac{i}{k}}(\beta)) \quad (52)$$

For finite knowledge bases, we never really need infinitely many certainty degrees, although the required number can depend on the size of the considered formulas. This is made precise in the following proposition.

Proposition 18. *Let $\Phi = \{\alpha_1 \succ \beta_1, \dots, \alpha_n \succ \beta_n, \gamma_{n+1} \sim \delta_{n+1}, \dots, \gamma_m \sim \delta_m\}$. If $k \geq n$, it holds that Φ is satisfiable in GPL_k iff Φ is satisfiable in $\text{GPL}_\succ^{\text{core}}$.*

Proof. Since any e-model of Φ in GPL_k is also an e-model of Φ in $\text{GPL}_{\succ}^{\text{core}}$, it is clear that satisfiability in GPL_k entails satisfiability in $\text{GPL}_{\succ}^{\text{core}}$.

Conversely, let π be an e-model of Ψ in $\text{GPL}_{\succ}^{\text{core}}$. In particular, among all such e-models, let π be such that the number of certainty levels in $\Lambda = \{\pi(\omega) \mid \omega \in \Omega\}$ is minimal. Let $\Lambda' = \{1 - N(\alpha_i) \mid 1 \leq i \leq n\} \cup \{1 - N(\beta_i) \mid 1 \leq i \leq n\} \subseteq \Lambda$. It holds that $\Lambda = \Lambda'$. Indeed, if this were not the case, we could define a possibility distribution π' as follows:

$$\pi'(\omega) = \begin{cases} \max\{\lambda \mid \lambda < \pi(\omega), \lambda \in \Lambda'\} & \text{if } \pi(\omega) \notin \Lambda' \text{ and } \pi(\omega) > \min \Lambda' \\ \min \Lambda & \text{if } \pi(\omega) < \min \Lambda' \\ \pi(\omega) & \text{otherwise} \end{cases} \quad (53)$$

It is straightforward to verify that the necessity measure induced by π' still satisfies all constraints. This shows that it is possible to choose a possibility distribution π' which only takes values from Λ' . Since we moreover clearly have that $\Lambda' \subseteq \Lambda$, and π was assumed to minimize $|\Lambda|$, we find $\Lambda' = \Lambda$.

We now show that $|\Lambda| \leq n + 1$. In particular, we show that if $\lambda \in \Lambda \setminus \{\min \Lambda, \max \Lambda\}$ it holds that there are at least two different formulas χ_1, χ_2 among $\{\alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n\}$ such that $N(\chi_1) = N(\chi_2) = 1 - \lambda$. Suppose this were not the case, and that e.g., α_i is the only formula for which $N(\alpha_i) = 1 - \lambda$. Define π' as follows:

$$\pi'(\omega) = \begin{cases} \max\{\mu : \mu \in \Lambda', \mu < \lambda\} & \text{if } \pi(\omega) = 1 - \lambda \\ \pi(\omega) & \text{otherwise} \end{cases}$$

Then it is clear that the necessity measure induced by π' still satisfies all constraints, while π' uses strictly fewer certainty levels than π , a contradiction. The case where β_i is the only formula with necessity $1 - \lambda$ is entirely analogous. Finally, since only the relative ordering of the certainty levels matters, it is always possible to choose π such that $\Lambda = \{0, \frac{1}{k}, \dots, 1\}$, $k \geq n$. In other words, there exists an e-model $\pi \in \mathcal{P}_k$ of Φ . \square

In general, to verify whether $\Phi \models_{\succ} \Psi$ holds, we can rewrite $\Phi \wedge \neg \Psi$ such that it is free of negations, by using the fact that $\neg(\alpha \succ \beta)$ is equivalent to $(\alpha \sim \beta) \vee (\beta \succ \alpha)$, and similarly $\neg(\alpha \sim \beta)$ is equivalent to $(\alpha \succ \beta) \vee (\beta \succ \alpha)$. Let Θ be the resulting formula. Then a suitable lower bound for k , ensuring that $\Phi \models_{\succ} \Psi$ iff $\Phi \models_{\text{GPL}}^k \Psi$, can be found as follows:

$$\begin{aligned} \text{bound}(\alpha \succ \beta) &= 1 \\ \text{bound}(\alpha \sim \beta) &= 0 \\ \text{bound}(\Theta_1 \wedge \Theta_2) &= \text{bound}(\Theta_1) + \text{bound}(\Theta_2) \\ \text{bound}(\Theta_1 \vee \Theta_2) &= \max(\text{bound}(\Theta_1), \text{bound}(\Theta_2)) \end{aligned}$$

Since satisfiability checking in $\text{GPL}_{\succ}^{\text{core}}$ can thus be reduced to checking the satisfiability of a GPL formula (whose size is polynomial in the size of the initial formula), it follows that this problem is in NP.

Proposition 19. *The complexity of deciding whether a GPL_{\succ}^{core} formula is satisfiable is NP-complete (w.r.t. the size of the formula).*

Proof. To see why satisfiability checking in GPL_{\succ}^{core} is NP-hard, note that the propositional formula α is satisfiable iff $\alpha \succ \perp$ is satisfiable. NP-membership directly follows from Proposition 18. \square

Proposition 20. *The complexity of deciding whether a GPL_{\succ} formula is satisfiable is NP-complete (w.r.t. the size of the formula).*

Proof. NP-hardness trivially follows from Proposition 19. We now propose an NP procedure to check the satisfiability of a GPL_{\succ} formula Φ . First, if Φ is satisfiable, in polynomial time we can guess a satisfiable term of the following form:

$$\bigwedge_{i=1}^n \mathbf{N}_{\lambda_i}(\alpha_i) \wedge \bigwedge_{i=n+1}^m \neg \mathbf{N}_{\lambda_i}(\alpha_i) \wedge \bigwedge_{i=m+1}^p (\alpha_i \succ \beta_i) \wedge \bigwedge_{i=p+1}^q (\alpha_i \sim \beta_i)$$

From Lemma 3 in Appendix C, we know that this term is satisfiable iff the following GPL_{\succ}^{core} formula is satisfiable.

$$\begin{aligned} & \bigwedge \{ \alpha_i \succ \alpha_j \mid 1 \leq i \leq n, n+1 \leq j \leq m, \lambda_i \geq \lambda_j \} \wedge \bigwedge \{ \alpha_i \sim \top \mid 1 \leq i \leq n, \lambda_i = 1 \} \\ & \wedge \bigwedge \{ \top \succ \alpha_i \mid n+1 \leq i \leq m \} \wedge \bigwedge \{ \alpha_i \succ \perp \mid 1 \leq i \leq n \} \\ & \wedge \bigwedge \{ \alpha_i \succ \beta_i \mid m+1 \leq i \leq p \} \wedge \bigwedge \{ \alpha_i \sim \beta_i \mid p+1 \leq i \leq q \} \end{aligned}$$

As in the proof of Proposition 19, we find that the satisfiability of this latter formula can be checked using an NP procedure. \square

8. Concluding remarks

We have introduced generalized possibilistic logic (GPL) as a general framework for reasoning about the revealed beliefs of an external agent. At the syntactic level, formulas in GPL are propositional combinations of meta-literals of the form $\mathbf{N}_{\lambda}(\alpha)$, expressing that it is known that an external agent believes α with certainty (at least) λ . Meta-literals of the form $\mathbf{\Pi}_{\lambda}(\alpha)$, $\mathbf{\Delta}_{\lambda}(\alpha)$ and $\mathbf{\nabla}_{\lambda}(\alpha)$ have also been introduced as abbreviations in the language. At the semantic level, the four considered types of meta-literals correspond to lower bounds on the four main uncertainty measures from possibility theory, i.e., the necessity, possibility, guaranteed possibility and potential necessity measures. We have moreover introduced a refinement of $\mathbf{\Delta}_{\lambda}(\alpha)$ and $\mathbf{\nabla}_{\lambda}(\alpha)$ to express context-dependent information about the ignorance of the agent in a more compact way.

After presenting an axiomatization of GPL and proving its soundness and completeness, we have studied two different ways to reason about the ignorance of an external agent, based on the principle of minimal specificity and based on guaranteed possibility respectively. Subsequently, we discussed the ability of GPL to model comparative uncertainty.

Among others, we axiomatized a logic, which can be embedded in GPL, to reason about arbitrary propositional combinations of statements of the form “ α is (strictly) more certain than β ”. As a special case, we obtain that GPL can be used to reason about propositional combinations of defaults, in the sense of System P. Next, we showed that the ability of GPL to model ignorance makes it a natural vehicle for expressing the semantics of non-monotonic logic programming formalisms. In particular, we showed how disjunctive answer set programs naturally correspond to a type of GPL theories, with answer sets corresponding to minimally specific e-models which are Boolean, in the sense that all interpretations are either possible to degree 1 or to degree 0. We then compared GPL with a fragment of Lifschitz’ logic of minimal belief and negation as failure (MBNF) which generalizes disjunctive answer set programming. While there is a close relationship between theories in this fragment of MBNF and the corresponding GPL theories, we have found that the notion of minimality demanded of MBNF models is less strict than the notion of minimal specificity, which is similar to an observation we made in [8] about equilibrium logic. While the less demanding notion of minimality in MBNF and equilibrium logic may have technical advantages, in particular for modelling inclusive disjunction [75], this finding casts doubt on the appropriateness of logics such as MBNF and equilibrium logic for epistemic reasoning.

In terms of computational complexity, we found natural decision problems at the first, second, third and fourth level of the polynomial hierarchy, where the third and fourth level are only reached when the refined modalities $\Delta_\lambda(\alpha)$ and $\nabla_\lambda(\alpha)$ are allowed. This confirms that the latter modalities allow us to compactly express knowledge that would otherwise require exponentially long formulas (unless the polynomial hierarchy collapses).

The ability of GPL to model both negation-as-failure and conditionals in an intuitive way demonstrates its versatility as a general logic for reasoning about the beliefs of an agent from an outsider point of view (as opposed to introspective reasoning). Among others, this makes GPL a natural choice for the formal study of access control mechanisms that need to maintain the confidentiality of some pieces of knowledge. For example, [78] discusses a number of settings where an information system needs to be able to determine whether answering a given query would allow the user to derive information that is supposed to remain secret, based on possibly incomplete knowledge of what that user already knows. The use of GPL is also natural in game theoretic settings, where agents need to reason based on their incomplete knowledge about the goals of other agents, e.g., as part of a negotiation process [79].

There are several ways in which GPL can be further extended. For example, a framework for multi-agent epistemic reasoning could be obtained by encapsulating GPL formulas similarly to how GPL encapsulates propositional formulas. Let us write $\mathbf{N}_{(\lambda,A)}(\alpha)$ to denote that agent A knows α with certainty λ . A formula such as $\mathbf{N}_{(\lambda,A)}(\mathbf{N}_{(\mu,B)}(\beta))$ expresses that (I know) that A knows with certainty λ that B knows β with certainty μ . At the semantic level, in the two-agent case, e-models would be of the form $(\pi_A, \pi_B, \pi_{AB}, \pi_{BA})$, where π_A and π_B are possibility distributions over propositional interpretations (encoding what objective formulas A and B know), and π_{AB} and π_{BA} are possibility distributions over possibility distributions over propositional interpretations (encoding resp. what A knows about what B knows, and what B knows about what A knows). We then have e.g.

$(\pi_A, \pi_B, \pi_{AB}, \pi_{BA}) \models \mathbf{N}_{(\lambda,A)}(\alpha \wedge \mathbf{N}_{(\mu,B)}(\beta))$, expressing that (I know that) A knows α and A knows that B knows β , iff $\pi_A \models \mathbf{N}_\lambda(\alpha)$ and $N_{AB}\{\tau \mid \tau \models \mathbf{N}_\mu(\beta)\} \geq \lambda$, where N_{AB} is the necessity measure induced by π_{AB} . Note that this approach does not allow us to consider chains of arbitrary length, e.g., formulas such as $\mathbf{N}_{(\lambda_1,A)}(\mathbf{N}_{(\mu,B)}(\mathbf{N}_{(\lambda_2,A)}(\alpha)))$ would require e-models of the form $(\pi_A, \pi_B, \pi_{AB}, \pi_{BA}, \pi_{ABA}, \pi_{BAB})$. In practice, this would not be a restriction, as we only need to consider those chains that appear in the given GPL knowledge base. Among others, it would be interesting to see how the interplay between minimal specificity and guaranteed possibility would allow us to model limits on agents' knowledge, and how such models would compare against multi-agent extensions of only knowing [80, 81].

Acknowledgments

We thank the anonymous reviewers for their valuable suggestions for improving this paper. Steven Schockaert has been supported by a grant from the Leverhulme Trust (RPG-2014-164) and by ERC Starting Grant 637277.

Appendix A. Proof of Proposition 1

Proof. The soundness of the axioms **(PL)**, **(K)**, **(N)**, **(D)** and **(W)** w.r.t. the semantics of GPL can readily be verified. Here we show that these axioms are also complete.

Let $F = \{\mathbf{N}_\lambda(\alpha) \mid \alpha \in \mathcal{L}, \lambda \in \Lambda_k^+\}$. Let Ω^* be the set of all propositional interpretations over the set of atomic formulas F . Given a GPL knowledge base K , let K^* be the propositional knowledge base over F , defined as

$$\begin{aligned} K^* = & K \cup \{\mathbf{N}_\lambda(\alpha \rightarrow \beta) \rightarrow (\mathbf{N}_\lambda(\alpha) \rightarrow \mathbf{N}_\lambda(\beta)) \mid \alpha, \beta \in \mathcal{L} \text{ and } \lambda \in \Lambda_k^+\} \\ & \cup \{\mathbf{N}_1(\top)\} \cup \{\mathbf{N}_\lambda(\alpha) \rightarrow \neg \mathbf{N}_{\frac{1}{k}}(\neg \alpha) \mid \alpha \in \mathcal{L}\} \\ & \cup \{\mathbf{N}_{\lambda_1}(\alpha) \rightarrow \mathbf{N}_{\lambda_2}(\alpha) \mid \alpha \in \mathcal{L}, \lambda_1 \geq \lambda_2\} \end{aligned}$$

where $\alpha \rightarrow \beta$ is an arbitrary (but fixed) formula from \mathcal{L} which is equivalent to $\alpha \rightarrow \beta$, and similarly for $\neg \alpha$. We then have that Φ can be derived from K using the axioms **(PL)**, **(K)**, **(N)**, **(D)**, **(W)** and modus ponens iff Φ can be derived from K^* in propositional logic.

To finish the proof, note that with every model I of K^* , we can associate a set-function $g_I : 2^\Omega \rightarrow \Lambda$ defined for $\alpha \in \mathcal{L}$ as

$$g_I(\llbracket \alpha \rrbracket) = \max\{\lambda \mid I \models \mathbf{N}_\lambda(\alpha)\}$$

where we define $g_I(\llbracket \alpha \rrbracket) = 0$ if $\{\lambda \mid I \models \mathbf{N}_\lambda(\alpha)\} = \emptyset$. From the fact that K^* contains every instantiation of the axioms **(K)**, **(N)**, **(D)** and **(W)**, we can derive the following properties for the function g_I :

- We have $g_I(\Omega) = 1$ thanks to the fact that $\mathbf{N}_1(\top) \in K^*$.
- We have $g_I(\emptyset) = 0$. Indeed, since K^* contains $\mathbf{N}_1(\top)$ and $\mathbf{N}_1(\top) \rightarrow \neg \mathbf{N}_{\frac{1}{k}}(\perp)$ (as an instantiation of **(D)**) and $\mathbf{N}_\lambda(\perp) \rightarrow \mathbf{N}_{\frac{1}{k}}(\perp)$ for every $\lambda \in \Lambda_k^+$ (as an instantiation of **(W)**) we know that $I \models \neg \mathbf{N}_\lambda(\perp)$ for every $\lambda \in \Lambda_k^+$. It follows that $\{\lambda \mid I \models \mathbf{N}_\lambda(\perp)\} = \emptyset$ and thus $g_I(\emptyset) = g_I(\llbracket \perp \rrbracket) = 0$.

- We have that g_I is monotone w.r.t. set inclusion. Indeed, if $\llbracket \alpha \rrbracket \subseteq \llbracket \beta \rrbracket$ then $\alpha \models \beta$ holds, which means that K^* will entail $\mathbf{N}_\lambda(\alpha) \rightarrow \mathbf{N}_\lambda(\beta)$ for every $\lambda \in \Lambda_k^+$ (as an instantiation of (\mathbf{K})). It follows that $\{\lambda \mid I \models \mathbf{N}_\lambda(\alpha)\} \subseteq \{\lambda \mid I \models \mathbf{N}_\lambda(\beta)\}$ and $g_I(\llbracket \alpha \rrbracket) \leq g_I(\llbracket \beta \rrbracket)$.
- We have that $g_I(\llbracket \alpha \wedge \beta \rrbracket) = \min(g_I(\llbracket \alpha \rrbracket), g_I(\llbracket \beta \rrbracket))$ for every $\alpha, \beta \in \mathcal{L}$. Indeed from the monotonicity of g_I we already have $g_I(\llbracket \alpha \wedge \beta \rrbracket) \leq \min(g_I(\llbracket \alpha \rrbracket), g_I(\llbracket \beta \rrbracket))$. Conversely, assume $I \models \mathbf{N}_\lambda(\alpha)$ and $I \models \mathbf{N}_\lambda(\beta)$. Using $\mathbf{N}_\lambda(\beta)$ and the instantiation of (\mathbf{K}) on the tautology $\beta \rightarrow (\alpha \rightarrow (\alpha \wedge \beta))$ we find $I \models \mathbf{N}_\lambda(\alpha \rightarrow (\alpha \wedge \beta))$. Using another instantiation of (\mathbf{K}) we find from $I \models \mathbf{N}_\lambda(\alpha \rightarrow (\alpha \wedge \beta))$ and $I \models \mathbf{N}_\lambda(\alpha)$ that $I \models \mathbf{N}_\lambda(\alpha \wedge \beta)$. It follows that $\{\lambda \mid I \models \mathbf{N}_\lambda(\alpha)\} \cap \{\lambda \mid I \models \mathbf{N}_\lambda(\beta)\} \subseteq \{\lambda \mid I \models \mathbf{N}_\lambda(\alpha \wedge \beta)\}$ and $g_I(\llbracket \alpha \wedge \beta \rrbracket) \geq \min(g_I(\llbracket \alpha \rrbracket), g_I(\llbracket \beta \rrbracket))$.

It is well-known [12] that every set-function which satisfies these four criteria is a necessity measure, and this necessity measure uniquely identifies a normalized possibility distribution π , which by construction will be an e-model of K . Conversely, it is easy to see that every e-model π of K corresponds to a unique propositional model I of K^* , defined as $I \models \mathbf{N}_\lambda(\alpha)$ iff $N(\alpha) \geq \lambda$ for N the necessity measure induced by π . \square

Appendix B. Proof of Proposition 4

Before we present the proof of the main result, which will apply to arbitrary propositional combinations of comparative certainty statements, we show that the proposed axioms are sufficient for detecting inconsistencies in sets of statements of the form $\alpha \succ \beta$.

Lemma 1. *Let $\Theta = \{\alpha_1 \succ \beta_1, \dots, \alpha_n \succ \beta_n\}$. It holds that Θ has an e-model iff $\Theta \not\models_{\succ}^{\text{core}} \perp$.*

Proof. The soundness of the axioms follows easily from well-known properties of necessity measures. We thus focus on showing that \perp can be derived if Θ is not satisfiable.

For each α_i there exist formulas $\alpha_i^1, \dots, \alpha_i^{m_i}$ such that $\alpha_i = \alpha_i^1 \wedge \dots \wedge \alpha_i^{m_i}$ and such that for each formula α_i^j it holds that $\llbracket \alpha_i^j \rrbracket = \Omega \setminus \{\omega_i^j\}$ for some propositional interpretation ω_i^j . Similarly, for each β_i there are formulas $\beta_i^1, \dots, \beta_i^{n_i}$ such that $\beta_i = \beta_i^1 \wedge \dots \wedge \beta_i^{n_i}$ and for each k it holds that $\llbracket \beta_i^k \rrbracket = \Omega \setminus \{\omega_i^k\}$ for some propositional interpretation ω_i^k .

From $\alpha_i \succ \beta_i$ we can derive using (\mathbf{Ax}_3) that $\alpha_i^j \succ \beta_i$ for every $j \in \{1, \dots, m_i\}$. Furthermore, using (25) we can derive $A_i^j = (\alpha_i^j \succ \beta_i^1) \vee \dots \vee (\alpha_i^j \succ \beta_i^{n_i})$. Conversely, from A_i^j we can derive $\alpha_i^j \succ \beta_i$ using (\mathbf{Ax}_3) , and from $\{\alpha_i^1 \succ \beta_i, \dots, \alpha_i^{m_i} \succ \beta_i\}$ we can derive $\alpha_i \succ \beta_i$ using (\mathbf{Ax}_2) and (\mathbf{Ax}_3) . It follows that $\alpha_i \succ \beta_i$ is equivalent to $\{A_i^1, \dots, A_i^{m_i}\}$.

Let ϕ be a mapping from $\{1, \dots, n\} \times \{1, \dots, m_i\}$ to $\{1, \dots, n_i\}$, allowing us to choose for each formula A_i^j a disjunct $\alpha_i^j \succ \beta_i^{\phi(i,j)}$. Clearly Θ is satisfiable iff there exists a mapping ϕ such that $\Theta_\phi = \{\alpha_i^j \succ \beta_i^{\phi(i,j)} \mid i \in \{1, \dots, n\}, j \in \{1, \dots, m_i\}\}$ is satisfiable. Accordingly, for Θ to be unsatisfiable, it suffices to show that for each such mapping ϕ , \perp can be derived from the formulas in Θ_ϕ .

Each formula $\alpha_i^j \succ \beta_i^{\phi(i,j)}$ corresponds to a constraint of the form $N(\alpha_i^j) > N(\beta_i^{\phi(i,j)})$, which by construction corresponds to the constraint $\pi(\omega_i^j) < \pi(\omega_i^{\phi(i,j)})$ on the associated possibility distribution. Clearly a set of such constraints can be satisfied unless there is

a cycle of the form $\pi(\omega^1) < \pi(\omega^2), \pi(\omega^2) < \pi(\omega^3), \dots, \pi(\omega^r) < \pi(\omega^1)$. In such a case, Θ_ϕ contains formulas of the form $\chi_1 \succ \chi_2, \chi_2 \succ \chi_3, \dots, \chi_r \succ \chi_1$ (up to syntactic variations of the arguments χ_l which we can ignore because of **(Ax₃)**). By repeatedly applying (20) we can then derive $\chi_1 \succ \chi_1$, which allows us to derive \perp using **(Ax₄)**. Thus we have shown that Θ_ϕ is unsatisfiable iff \perp can be derived. \square

In the next lemma, we additionally consider formulas of the form $\gamma \sim \delta$.

Lemma 2. *Let $\Theta = \{\alpha_1 \succ \beta_1, \dots, \alpha_n \succ \beta_n, \gamma_{n+1} \sim \delta_{n+1}, \dots, \gamma_m \sim \delta_m\}$. It holds that Θ has an e-model iff $\Theta \not\vdash_{\succ}^{core} \perp$.*

Proof. We show that Θ has an e-model if $\Theta \not\vdash_{\succ}^{core} \perp$; the other direction follows from the soundness of the axioms.

From Lemma 1 we know that $\{\alpha_1 \succ \beta_1, \dots, \alpha_n \succ \beta_n\}$ is satisfiable, given that we assumed that no inconsistency can be derived. Let π be a possibility distribution that satisfies $\{\alpha_1 \succ \beta_1, \dots, \alpha_n \succ \beta_n\}$. From the fact that \succ is a strict partial order, the fact that \sim is an equivalence relation and **(Ax₅)**, it follows that we can partition the set of formulas $X = \{\alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \gamma_{n+1}, \dots, \gamma_m, \delta_{n+1}, \dots, \delta_m\}$ as $X = X_1 \cup \dots \cup X_s$ where for $\chi \in X_r$ and $\chi' \in X_t$, with $r < t$, it holds that Θ contains a conjunct $\alpha \succ \beta$ where $\models \alpha \equiv \chi$ and $\models \beta \equiv \chi'$. Let $L_r = \{1 - N(\chi) \mid \chi \in X_r\}$. Because of how we choose π , it holds that $\max L_r < \min L_t$ for $r < t$. We now define the possibility distribution π' as follows:

$$\pi'(\omega) = \begin{cases} \min L_1 & \text{if } \pi(\omega) \leq \max L_1 \\ \min L_s & \text{if } \pi(\omega) > \max L_{s-1} \\ \min L_r & \text{if } 1 < r < s \text{ and } \max L_{r-1} < \pi(\omega) \leq \max L_r \end{cases}$$

Let N' be the necessity measure induced by π' . It is straightforward to verify that $N'(\chi) = N'(\chi')$ for $\chi, \chi' \in X_r$ and that $N'(\chi) > N'(\chi')$ if $\chi \in X_r$ and $\chi' \in X_t$ with $r < t$. In other words, it holds that $\pi' \models_{\succ} \Theta$. \square

Proposition 4 now follows easily.

Proof. The soundness of the axioms can be verified straightforwardly. To see why the completeness result holds, note that when $\Phi \models_{\succ} \Psi$ holds, we have that $\Phi \wedge \neg\Psi$ is unsatisfiable. Let $\Phi_1 \vee \dots \vee \Phi_n$ be a formula in DNF which is equivalent to $\Phi \wedge \neg\Psi$, where each disjunct Φ_i is a conjunction $\phi_1^i \wedge \dots \wedge \phi_n^i$ of meta-literals of the form $\alpha \succ \beta$ and $\alpha \sim \beta$. From Lemma 2 it immediately follows that each such a disjunct Φ_i is unsatisfiable iff $\phi_1^i \wedge \dots \wedge \phi_n^i \vdash_{\succ}^{core} \perp$. Since $\Phi_1 \vee \dots \vee \Phi_n$ is inconsistent, none of the disjuncts Φ_i are satisfiable, from which we can thus conclude $\Phi_1 \vee \dots \vee \Phi_n \vdash_{\succ}^{core} \perp$ and thus $\Phi \vdash_{\succ}^{core} \Psi$. \square

Appendix C. Proof of Proposition 5

Lemma 3. Let $K = \{\mathbf{N}_{\lambda_i}(\alpha_i) \mid 1 \leq i \leq n\} \cup \{\neg \mathbf{N}_{\lambda_i}(\alpha_i) \mid n+1 \leq i \leq m\} \cup \{\alpha_i \succ \beta_i \mid m+1 \leq i \leq p\} \cup \{\alpha_i \sim \beta_i \mid p+1 \leq i \leq q\}$. Let the set of GPL_{\succ}^{core} formulas L be given by:

$$\begin{aligned} L = & \{\alpha_i \succ \alpha_j \mid 1 \leq i \leq n, n+1 \leq j \leq m, \lambda_i \geq \lambda_j\} \cup \{\alpha_i \sim \top \mid 1 \leq i \leq n, \lambda_i = 1\} \\ & \cup \{\top \succ \alpha_i \mid n+1 \leq i \leq m\} \cup \{\alpha_i \succ \perp \mid 1 \leq i \leq n\} \\ & \cup \{\alpha_i \succ \beta_i \mid m+1 \leq i \leq p\} \cup \{\alpha_i \sim \beta_i \mid p+1 \leq i \leq q\} \end{aligned}$$

It holds that K has an e-model iff L has an e-model.

Proof. It is straightforward to verify that all the considered axioms are sound, hence when K has an e-model it must be the case that L has an e-model as well. Conversely, suppose that L has an e-model π . Note that by definition of e-model, π is then normalised. Clearly all formulas of the form $\alpha \succ \beta$ and $\alpha \sim \beta$ in K are satisfied by π , as these formulas are also included in L . Furthermore, for every formula of the form $\mathbf{N}_1(\alpha)$ in K , L will contain the formula $\alpha \sim \top$, and thus $N(\alpha) = N(\top) = 1$ for N the necessity measure induced by π . Hence all formulas of the form $\mathbf{N}_1(\alpha)$ from K are satisfied by π .

- Assume that some formula $\mathbf{N}_{\lambda_l}(\alpha_l)$, with $\lambda_l < 1$, is not satisfied by π and let $c = N(\alpha_l)$; note that we then have $c < \lambda_l$. Furthermore note that $c > 0$ since L contains the formula $\alpha \succ \perp$. Let d be the smallest element from the set $\{\lambda_{n+1}, \dots, \lambda_m, 1\}$ which is strictly greater than λ_l ; since $\lambda_l < 1$ such an element d must indeed exist. We define the normalized possibility distribution π' for $\omega \in \Omega$ as follows:

$$\pi'(\omega) = \begin{cases} \pi(\omega) & \text{if } \pi(\omega) > 1 - c \text{ or } \pi(\omega) \leq 1 - d \\ 1 - \lambda_l - \frac{(1-c-\pi(\omega))}{d-c}(d - \lambda_l) & \text{otherwise} \end{cases}$$

The transformation from π to π' is illustrated in Figure 1(a). First note that from $c > 0$ and the fact that π is normalised, it follows that π' is normalised. Furthermore, since $c < \lambda_l < d$ we have that the transformation from π to π' is order-preserving, i.e. we have $\pi(\omega_1) < \pi(\omega_2)$ iff $\pi'(\omega_1) < \pi'(\omega_2)$. It follows that π' satisfies all formulas of the form $\alpha \succ \beta$ and $\alpha \sim \beta$ in K , given that π satisfies these formulas, as the satisfaction of such formulas only relies on the ordering of the possibility degrees. It clearly also holds that $\pi' \models \mathbf{N}_{\lambda_l}(\alpha_l)$. Indeed, since $N(\alpha_l) = c$ we know that $\pi(\omega) \leq 1 - c$ for every model of $\neg \alpha_l$. By definition of π' this means that $\pi(\omega) \leq 1 - \lambda_l$ for each such ω , and thus $\pi' \models \mathbf{N}_{\lambda_l}(\alpha_l)$. Furthermore, since $\pi'(\omega) \leq \pi(\omega)$ for every $\omega \in \Omega$, it holds that π' satisfies all the formulas of the form $\mathbf{N}_{\lambda}(\alpha)$ that were already satisfied by π .

We now show that the same holds for formulas of the form $\neg \mathbf{N}_{\lambda}(\alpha)$. Suppose $\pi \models_{\succ} \neg \mathbf{N}_{\lambda_p}(\alpha_p)$, with $p \in \{n+1, \dots, m\}$.

- If $\lambda_p > \lambda_l$ then $\lambda_p \geq d$. If it were the case that $\pi' \models_{\succ} \mathbf{N}_{\lambda_p}(\alpha_p)$, then for every model ω of $\neg \alpha_p$ we would have $\pi'(\omega) \leq 1 - \lambda_p \leq 1 - d$. However, by construction of π' this would mean $\pi'(\omega) = \pi(\omega)$ for each model of $\neg \alpha_p$ and thus $N(\alpha_p) = N'(\alpha_p)$, a contradiction.

- If $\lambda_p \leq \lambda_l$ then $\alpha_l \succ \alpha_p$ is in L . Hence we have that $c = N(\alpha_l) > N(\alpha_p)$. It follows that there is some model ω^* of $\neg\alpha_p$ such that $\pi(\omega^*) > 1 - \min(c, \lambda_p)$. By definition of π' we then have $\pi'(\omega^*) = \pi(\omega)$ and thus $\pi' \models_{\succ} \neg\mathbf{N}_{\min(c, \lambda_p)}(\alpha_p)$ and a fortiori $\pi' \models_{\succ} \neg\mathbf{N}_{\lambda_p}(\alpha_p)$.
- Now consider the case where some formula $\neg\mathbf{N}_{\lambda_l}(\alpha_l)$ is not satisfied by π , and let us write $c = N(\alpha_l)$; note that we then have $c \geq \lambda_l > 0$. Furthermore note that we have $c < 1$ since L contains the formula $\top \succ \alpha$. Let d be the largest element from the set $\{\lambda_1, \dots, \lambda_n, 0\}$ which is strictly smaller than λ_l ; since $\lambda_l > 0$, such an element d must exist. Let us write $e = \frac{d + \lambda_l}{2}$. We define the normalized possibility distribution π' for $\omega \in \Omega$ as follows:

$$\pi'(\omega) = \begin{cases} \pi(\omega) & \text{if } \pi(\omega) < 1 - c \text{ or } \pi(\omega) \geq 1 - d \\ 1 - e + \frac{(1 - c - \pi(\omega))}{d - c}(e - d) & \text{otherwise} \end{cases}$$

This transformation from π to π' is illustrated in Figure 1(b). Since $\pi'(\omega) \geq \pi(\omega)$ and π is normalised, we have that π' is normalised as well. Furthermore, since $d < e < \lambda_l \leq c$, we have that the transformation from π to π' is order-preserving, and thus that π' satisfies all formulas of the form $\alpha \succ \beta$ and $\alpha \sim \beta$. We also have that $\pi' \models \neg\mathbf{N}_{\lambda_l}(\alpha_l)$. Indeed, since $c < 1$ there must exist model ω^* of $\neg\alpha_l$ such that $\pi(\omega^*) = 1 - c$. By construction, it holds that $\pi'(\omega^*) = 1 - e$, from which it follows that $N'(\alpha_l) \leq e < \lambda_l$, with N' the necessity measure induced by π' . Furthermore, since $\pi'(\omega) \geq \pi(\omega)$ for every $\omega \in \Omega$, it holds that π' satisfies all the formulas $\neg\mathbf{N}_{\lambda_p}(\alpha_p)$ that were already satisfied by π .

We now show that the same holds for formulas of the form $\mathbf{N}_{\lambda_p}(\alpha_p)$. Suppose $\pi \models_{\succ} \mathbf{N}_{\lambda_p}(\alpha_p)$, with $p \in \{1, \dots, n\}$.

- If $\lambda_l > \lambda_p$ then $\lambda_p \leq d$. Suppose $\pi' \not\models_{\succ} \mathbf{N}_{\lambda_p}(\alpha_p)$. Then there exists a model ω^* of $\neg\alpha_p$ such that $\pi'(\omega^*) > 1 - \lambda_p$. However, since $1 - \lambda_p \geq 1 - d$, we have $\pi'(\omega^*) = \pi(\omega^*)$, which would mean $\pi \not\models_{\succ} \mathbf{N}_{\lambda_p}(\alpha_p)$, a contradiction.
- If $\lambda_l \leq \lambda_p$ then $\alpha_p \succ \alpha_l$ is in L , and thus $N(\alpha_p) > N(\alpha_l) = c$. It follows that $\pi(\omega) < 1 - c$ for every model of $\neg\alpha_p$. Thus we have $\pi(\omega) = \pi'(\omega)$ for every model of $\neg\alpha_p$, and in particular $N(\alpha_p) = N'(\alpha_p)$, for N' the necessity measure induced by π' .

By iterating this construction until all formulas of the form $\mathbf{N}_{\lambda_l}(\alpha_l)$ and $\neg\mathbf{N}_{\lambda_l}(\alpha_l)$ are satisfied, we obtain an e-model of K . \square

Noting that all formulas in the set L can be derived from K using \vdash_{\succ} , the completeness of the GPL_{\succ} axioms follows easily from the previous lemma, together with the completeness of the $\text{GPL}_{\succ}^{\text{core}}$ axioms from Section 5.2.

Proof. As it is clear that $\Phi \vdash_{\succ} \Psi$ implies $\Phi \models_{\succ} \Psi$, we focus on the completeness result. If $\Phi \not\models_{\succ} \Psi$ then $\Phi \wedge \neg\Psi$ is unsatisfiable. Let $\Phi_1 \vee \dots \vee \Phi_n$ be a formula in DNF which is




Figure C.1: Transformations used in the proof of Lemma 3.

equivalent to $\Phi \wedge \neg\Psi$, where each disjunct Φ_i is a conjunction $\phi_1^i \wedge \dots \wedge \phi_n^i$ of meta-literals of the form $\mathbf{N}_\lambda(\alpha)$, $\neg\mathbf{N}_\lambda(\alpha)$, $\alpha \succ \beta$ and $\alpha \sim \beta$. From Lemma 3 it immediately follows that each such a disjunct Φ_i is unsatisfiable iff $\phi_1 \wedge \dots \wedge \phi_n \vdash_{\succ} \perp$. Since $\Phi_1 \vee \dots \vee \Phi_n$ is inconsistent, none of the disjuncts Φ_i are satisfiable, from which we can thus conclude $\Phi_1 \vee \dots \vee \Phi_n \vdash_{\succ} \perp$ and thus $\Phi \vdash_{\succ} \Psi$. \square

Appendix D. Proof of Proposition 7

We first show the soundness of the axioms.

Proof. As already discussed, the soundness of the axioms **(RE)**, **(LLE)**, **(RW)**, **(OR)**, **(CM)**, **(CUT)** and **(INC)** follows from the soundness of **(Ax₁)** and **(RI₁)**–**(RI₃)**. To show that **(WRM)** is sound w.r.t. the possibilistic semantics it is sufficient to show that for any possibility measure Π it holds that $\Pi(\alpha \wedge \gamma) > \Pi(\alpha \wedge \neg\gamma)$ and $\Pi(\alpha \wedge \neg\beta) \leq \Pi(\alpha \wedge \beta)$ together imply $\Pi(\alpha \wedge \beta \wedge \gamma) > \Pi(\alpha \wedge \beta \wedge \neg\gamma)$. To see that this is the case⁶, note that $\Pi(\alpha \wedge \gamma) > \Pi(\alpha \wedge \neg\gamma)$ means that either $\Pi(\alpha \wedge \beta \wedge \gamma) > \Pi(\alpha \wedge \neg\gamma)$ or $\Pi(\alpha \wedge \neg\beta \wedge \gamma) > \Pi(\alpha \wedge \neg\gamma)$. In the former case, we readily obtain $\Pi(\alpha \wedge \beta \wedge \gamma) > \Pi(\alpha \wedge \beta \wedge \neg\gamma)$. In the latter case, we also need to have $\Pi(\alpha \wedge \beta \wedge \gamma) > \Pi(\alpha \wedge \beta \wedge \neg\gamma)$ since otherwise we find $\Pi(\alpha \wedge \neg\beta) \leq \Pi(\alpha \wedge \beta) = \Pi(\alpha \wedge \beta \wedge \neg\gamma)$, and in particular $\Pi(\alpha \wedge \neg\beta \wedge \gamma) \leq \Pi(\alpha \wedge \neg\beta) \leq \Pi(\alpha \wedge \beta \wedge \neg\gamma) \leq \Pi(\alpha \wedge \neg\gamma)$, a contradiction. \square

To show the completeness result, we will use two lemmas.

Lemma 4. $c(\alpha_1 \sim \beta_1) \wedge \dots \wedge c(\alpha_n \sim \beta_n)$ is satisfiable iff $c(\alpha_1 \sim \beta_1) \wedge \dots \wedge c(\alpha_n \sim \beta_n) \not\vdash_c \perp$.

⁶The proof appears in [83] but is given again as this paper may be hard to find.

Proof. We show that $c(\alpha_1 \sim \beta_1) \wedge \dots \wedge c(\alpha_n \sim \beta_n) \vdash_c \perp$ if $c(\alpha_1 \sim \beta_1) \wedge \dots \wedge c(\alpha_n \sim \beta_n)$ is unsatisfiable. The other direction trivially follows from the soundness of the axioms.

Without loss of generality we can assume that $\Psi = c(\alpha_1 \sim \beta_1) \wedge \dots \wedge c(\alpha_{n-1} \sim \beta_{n-1})$ is satisfiable, but that every e-model π of Ψ is such that $\Pi(\alpha_n \wedge \beta_n) \leq \Pi(\alpha_n \wedge \neg \beta_n)$, with Π the possibility measure induced by π (i.e., from an inconsistent conjunction of conditionals with a consistent antecedent, we can always select a non-empty, maximal consistent subset of conditionals). This is only possible if every e-model π of Ψ is such that $\Pi(\alpha_n \wedge \beta_n) < \Pi(\alpha_n \wedge \neg \beta_n)$. Indeed, suppose there was an e-model π of Ψ such that $\Pi(\alpha_n \wedge \beta_n) = \Pi(\alpha_n \wedge \neg \beta_n) < 1$ and define π' as follows ($0 < \varepsilon < 1 - \Pi(\alpha_n \wedge \beta_n)$):

$$\pi'(\omega) = \begin{cases} \pi(\omega) + \varepsilon & \text{if } \pi(\omega) = \Pi(\alpha_n \wedge \beta_n) \text{ and } \omega \models \alpha_n \wedge \beta_n \\ \pi(\omega) & \text{otherwise} \end{cases}$$

It is easy to see that if $\varepsilon < \min_{i=1}^{n-1} (\Pi(\alpha_i \wedge \beta_i) - \Pi(\alpha_i \wedge \neg \beta_i))$ it holds that π' is an e-model of $c(\alpha_1 \sim \beta_1) \wedge \dots \wedge c(\alpha_n \sim \beta_n)$. If $\Pi(\alpha_n \wedge \beta_n) = \Pi(\alpha_n \wedge \neg \beta_n) = 1$ we instead define π' as follows ($0 < \varepsilon < \Pi(\alpha_n \wedge \neg \beta_n)$):

$$\pi'(\omega) = \begin{cases} \pi(\omega) - \varepsilon & \text{if } \pi(\omega) = \Pi(\alpha_n \wedge \neg \beta_n) \text{ and } \omega \models \alpha_n \wedge \neg \beta_n \\ \pi(\omega) & \text{otherwise} \end{cases}$$

Thus we can assume that for every e-model π of Ψ , we have $\pi \models_c c(\alpha_n \sim \neg \beta_n)$. Given the completeness result from [59] for consistent sets of conditionals, it follows that $c(\alpha_n \sim \neg \beta_n)$ can be derived from $c(\alpha_1 \sim \beta_1) \wedge \dots \wedge c(\alpha_{n-1} \sim \beta_{n-1})$. Finally, using **(INC)** and the axioms of classical logic, we can derive \perp from $c(\alpha_n \sim \neg \beta_n)$ and $c(\alpha_n \sim \beta_n)$. \square

Lemma 5. *Let $\{\gamma_1 \sim \delta_1, \dots, \gamma_m \sim \delta_m\}$ be a rationally closed set of defaults and let $\Phi = c(\gamma_1 \sim \delta_1) \wedge \dots \wedge c(\gamma_m \sim \delta_m) \wedge \neg c(\gamma_{m+1} \sim \delta_{m+1}) \wedge \dots \wedge \neg c(\gamma_r \sim \delta_r)$. It holds that Φ is satisfiable iff $\Phi \not\vdash \perp$.*

Proof. Assume that $\Phi \not\vdash \perp$; we show that Φ has an e-model. Note that the other direction follows trivially from the soundness of the axioms.

Since $\Phi \not\vdash \perp$, it follows from Lemma 4 that the set of conditionals $\{\gamma_1 \sim \delta_1, \dots, \gamma_m \sim \delta_m\}$ is consistent. Given the correspondence between consistent sets of conditionals and possibility theory shown in [59], this means that $c(\gamma_1 \sim \delta_1) \wedge \dots \wedge c(\gamma_m \sim \delta_m)$ is satisfiable, and in particular that it has an e-model π such that the conditionals satisfied by π are exactly those in $\{\gamma_1 \sim \delta_1, \dots, \gamma_m \sim \delta_m\}$, since we assumed that this set is rationally closed. Since $\Phi \not\vdash \perp$ it holds that none of the defaults $\gamma_{m+1} \sim \delta_{m+1}, \dots, \gamma_r \sim \delta_r$ is included in this latter set, and thus that π is an e-model of Φ . \square

We now show the completeness result.

Proof. As it is clear that $\Phi \vdash_c \Psi$ implies $\Phi \models_c \Psi$, we focus on the completeness result. If $\Phi \models_c \Psi$ then $\Phi \wedge \neg \Psi$ is unsatisfiable. Let $\Phi_1 \vee \dots \vee \Phi_n$ be a formula in DNF which is equivalent to $\Phi \wedge \neg \Psi$, where each disjunct Φ_i is a conjunction $\phi_1^i \wedge \dots \wedge \phi_n^i$ of meta-literals

of the form $c(\gamma \sim \delta)$ or $\neg c(\gamma \sim \delta)$. Moreover, thanks to axiom **(WRM)** we can assume that the set of meta-literals of the form $c(\gamma \sim \delta)$ correspond to a rationally closed set of defaults. From Lemma 5 it then follows that each such a disjunct Φ_i is satisfiable iff $\phi_1 \wedge \dots \wedge \phi_n \vdash_c \perp$. Since $\Phi_1 \vee \dots \vee \Phi_n$ is inconsistent, none of the disjuncts Φ_i are satisfiable, from which we can thus conclude $\Phi_1 \vee \dots \vee \Phi_n \vdash_c \perp$ and thus $\Phi \vdash_c \Psi$. \square

Appendix E. Proof of Proposition 8

To prove the three results, we show that for every minimally specific e-model π of K_P which satisfies $\bigwedge_{a \in At} \mathbf{N}_1(a) \vee \mathbf{N}_1(\neg a) \vee (\mathbf{\Pi}_1(a) \wedge \mathbf{\Pi}_1(\neg a))$, it holds that the set M defined as follows is an answer set of P :

$$M = \{l \in Lit \mid \pi \models \mathbf{N}_1(l)\} \quad (\text{E.1})$$

and that all answer sets are of this form, i.e., that for every answer set M of P it holds that the possibility distribution π_M defined as follows is a minimally specific e-model of K_P which satisfies $\bigwedge_{a \in At} \mathbf{N}_1(a) \vee \mathbf{N}_1(\neg a) \vee (\mathbf{\Pi}_1(a) \wedge \mathbf{\Pi}_1(\neg a))$:

$$\pi_M(\omega) = \begin{cases} 1 & \text{if } \omega \models l \text{ for every literal } l \in M \\ 0 & \text{otherwise} \end{cases}$$

- Let M be a consistent answer set of P . We show that π_M is a minimally specific e-model of K_P which satisfies $\bigwedge_{a \in \mathcal{H}_P} \mathbf{N}_1(a) \vee \mathbf{N}_1(\neg a) \vee (\mathbf{\Pi}_1(a) \wedge \mathbf{\Pi}_1(\neg a))$. The latter trivially follows from the fact that $\pi_M(\omega) \in \{0, 1\}$ for each propositional interpretation ω . It remains to be shown that π_M is a minimally specific e-model of K_P . Note that because M is a consistent answer set, it holds that π_M is normalized.

To see why π_M is an e-model, consider a rule from P of the form (27) and assume that π_M satisfies $\mathbf{N}_1(b_1) \wedge \dots \wedge \mathbf{N}_1(b_m) \wedge \mathbf{\Pi}_1(\neg c_1) \wedge \dots \wedge \mathbf{\Pi}_1(\neg c_\ell)$. Since $\pi_M \models \mathbf{N}_1(b_i)$, we have $\pi_M(\omega) = 0$ for all worlds ω in which b_i is false. By construction this means that $b_i \in M$. Similarly, since $\pi_M \models \mathbf{\Pi}_1(\neg c_i)$ there is at least one world ω in which c_i is false, which by construction means that $c_i \notin M$.

It follows that the rule $a_1 \vee \dots \vee a_n \leftarrow b_1 \wedge \dots \wedge b_m$ is in the reduct P^M . Since M is a model of this rule (given the assumption that M is an answer set), one of a_1, \dots, a_n is in M . By construction this means that $\pi_M \models \mathbf{N}_1(a_i)$ for some $1 \leq i \leq n$.

It remains to be shown that π_M is minimally specific. Suppose π^* is an e-model of K_P which is strictly less specific than π_M . Then there exists a world ω such that $\pi_M(\omega) = 0$ and $\pi^*(\omega) > 0$. This means that there is a literal $l^* \in M$ such that $\pi \models \mathbf{N}_1(l^*)$ and $\pi^* \not\models \mathbf{N}_1(l^*)$. Let $M^* = \{l \mid \pi^* \models \mathbf{N}_1(l)\}$. It is clear that $M^* \subset M$. It is not hard to see that M^* is a model of P^M , which is a contradiction since M is an answer set and thus by definition the unique minimal model of P^M .

- Let π be a minimally specific e-model of K_P , which satisfies $\bigwedge_{a \in At} \mathbf{N}_1(a) \vee \mathbf{N}_1(\neg a) \vee (\mathbf{\Pi}_1(a) \wedge \mathbf{\Pi}_1(\neg a))$. Let M be defined as in (E.1). From the fact that π is normalized, it immediately follows that M is consistent.

First we show that M is a model of P^M . Consider a rule of the form (27), for which $\{c_1, \dots, c_\ell\} \cap M = \emptyset$, i.e., such that $a_1 \vee \dots \vee a_n \leftarrow b_1 \wedge \dots \wedge b_m$ is in the reduct P^M . Since $c_i \notin M$ we know that $\pi \not\models \mathbf{N}_1(c_i)$. Because π by assumption satisfies $\mathbf{N}_1(c_i) \vee \mathbf{N}_1(\neg c_i) \vee (\mathbf{\Pi}_1(c_i) \wedge \mathbf{\Pi}_1(\neg c_i))$ it follows that $\pi \models \mathbf{N}_1(\neg c_i)$ or $\pi \models \mathbf{\Pi}_1(\neg c_i)$. Using axiom **(D)** we find that $\pi \models \mathbf{\Pi}_1(\neg c_i)$ for each c_i . We then have that π is an e-model of $\mathbf{N}_1(a_1) \vee \dots \vee \mathbf{N}_1(a_n) \vee \neg \mathbf{N}_1(b_1) \vee \dots \vee \neg \mathbf{N}_1(b_m)$, from which we immediately find that M is a model of $a_1 \vee \dots \vee a_n \leftarrow b_1 \wedge \dots \wedge b_m$.

Now suppose that M is not an answer set. Then the minimal model M^* of P^M is such that $M^* \subset M$. The corresponding possibility distribution π^* is defined as

$$\pi^*(\omega) = \begin{cases} 1 & \text{if } \omega \models l \text{ for every literal } l \in M^* \\ 0 & \text{otherwise} \end{cases}$$

When $\pi^*(\omega) = 0$, we have that $\omega \models \neg l$ for some $l \in M^*$. Then we also have $l \in M$ which means $\pi \models \mathbf{N}_1(l)$ and in particular $\pi(\omega) = 0$. It follows that π^* is less specific than π , and since there is a literal $l_0 \in M \setminus M^*$ such that $\pi \models \mathbf{N}_1(l_0)$ and $\pi^* \not\models \mathbf{N}_1(l_0)$, it follows that π^* is strictly less specific than π .

Now we define a third possibility distribution π^+ as follows:

$$\pi^+(\omega) = \begin{cases} \frac{1}{2} & \text{if } \pi(\omega) = 0 \text{ and } \pi^*(\omega) = 1 \\ \pi^*(\omega) & \text{otherwise} \end{cases}$$

Clearly, we have that π^+ is strictly less specific than π . We show that π^+ is an e-model of K_P , contradicting our assumption that π were a minimally specific e-model of K_P , from which it then follows that M must be an answer set.

Consider a rule of the form (27). If $\{c_1, \dots, c_\ell\} \cap M \neq \emptyset$, by construction $\pi \models \mathbf{N}_1(c_i)$ must hold for some c_i . This means $\pi^+ \models \mathbf{N}_{\frac{1}{2}}(c_i)$ and in particular $\pi^+ \models \neg \mathbf{\Pi}_1(\neg c_i)$. This means that π^+ satisfies the corresponding GPL formula of the form (28). On the other hand, if $\{c_1, \dots, c_\ell\} \cap M = \emptyset$, then M^* satisfies the rule $a_1 \vee \dots \vee a_n \leftarrow b_1 \wedge \dots \wedge b_m$, since this rule is in the reduct P^M . Moreover, by construction we have $l \in M^*$ iff $\pi^* \models \mathbf{N}_1(l)$ iff $\pi^+ \models \mathbf{N}_1(l)$. It follows that π^+ satisfies $\mathbf{N}_1(a_1) \vee \dots \vee \mathbf{N}_1(a_n) \vee \neg \mathbf{N}_1(b_1) \vee \dots \vee \neg \mathbf{N}_1(b_m)$.

References

- [1] D. Dubois, J. Lang, H. Prade, Possibilistic logic, in: D. N. D. Gabbay, C. Hogger J. Robinson (Ed.), Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3, Oxford University Press, 1994, pp. 439–513.
- [2] D. Dubois, H. Prade, Epistemic entrenchment and possibilistic logic, Artificial Intelligence 50 (2) (1991) 223 – 239.
- [3] J. Y. Halpern, R. Fagin, Y. Moses, M. Vardi, Reasoning About Knowledge, MIT Press (Revised paperback edition), 2003.
- [4] M. Gelfond, V. Lifschitz, The stable model semantics for logic programming, in: Proceedings of the Fifth International Conference and Symposium on Logic Programming, 1988, pp. 1081–1086.

- [5] S. Kraus, D. Lehmann, M. Magidor, Nonmonotonic reasoning, preferential models and cumulative logics, *Artificial Intelligence* 44 (1-2) (1990) 167–207.
- [6] H. J. Levesque, All I know: A study in autoepistemic logic, *Artificial Intelligence* 42 (1990) 263 – 309.
- [7] F. Touazi, C. Cayrol, D. Dubois, Possibilistic reasoning with partially ordered beliefs, *Journal of Applied Logic* 13 (4) (2015) 770–798.
- [8] D. Dubois, H. Prade, S. Schockaert, Stable models in generalized possibilistic logic, in: *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning*, 2012, pp. 519–529.
- [9] D. Dubois, H. Prade, S. Schockaert, Reasoning about uncertainty and explicit ignorance in generalized possibilistic logic, in: *Proceedings of the 21st European Conference on Artificial Intelligence*, 2014, pp. 261–266.
- [10] L. A. Zadeh, Fuzzy sets as a basis for a theory of possibility, *Fuzzy Sets and Systems* 1 (1978) 3–28.
- [11] D. Dubois, H. Prade, *Possibility Theory*, Plenum Press, New York (NY), 1988.
- [12] D. Dubois, H. Prade, Possibility theory: qualitative and quantitative aspects, in: D. Gabbay, P. Smets (Eds.), *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, Vol. 1, Kluwer Academic, 1998, pp. 169–226.
- [13] G. Shackle, *Expectation in Economics*, Cambridge University Press, 1949.
- [14] D. Dubois, H. Prade, S. Sandri, On possibility/probability transformations, in: R. Lowen, M. Roubens (Eds.), *Fuzzy Logic*, Vol. 12 of *Theory and Decision Library*, Springer, 1993, pp. 103–112.
- [15] D. Dubois, S. Moral, H. Prade, A semantics for possibility theory based on likelihoods, *Journal of Mathematical Analysis and Applications* 205 (2) (1997) 359 – 380.
- [16] D. Dubois, H. Prade, Fuzzy sets, probability and measurement, *European Journal of Operational Research* 40 (2) (1989) 135–154.
- [17] D. Dubois, H. Prade, Qualitative and semi-quantitative modeling of uncertain knowledge - A discussion, in: *Computational Models of Rationality*, Essays dedicated to Gabriele Kern-Isberner on the occasion of her 60th birthday, College Publications, 2016, pp. 280–296.
- [18] W. Spohn, Ordinal conditional functions: A dynamic theory of epistemic states, in: W. Harper, B. Skyrms (Eds.), *Causation in Decision, Belief Change, and Statistics*, Springer Netherlands, 1988, pp. 105–134.
- [19] J. Lang, D. Dubois, H. Prade, A logic of graded possibility and certainty coping with partial inconsistency, in: *Proceedings of the Seventh Conference on Uncertainty in Artificial Intelligence*, 1991, pp. 188–196.
- [20] P. Gärdenfors, *Knowledge in Flux*, MIT Press, 1988.
- [21] D. Dubois, J. Lang, H. Prade, Automated reasoning using possibilistic logic: semantics, belief revision, and variable certainty weights, *IEEE Transactions on Knowledge and Data Engineering* 6 (1) (1994) 64–71.
- [22] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, H. Prade, Inconsistency management and prioritized syntax-based entailment, in: *Proceedings of the International Joint Conference on Artificial Intelligence*, Vol. 93, 1993, pp. 640–645.
- [23] S. Benferhat, D. Dubois, H. Prade, Practical handling of exception-tainted rules and independence information in possibilistic logic, *Applied Intelligence* 9 (2) (1998) 101–127.
- [24] M. Richardson, P. Domingos, Markov logic networks, *Machine Learning* 62 (2006) 107–136.
- [25] F. Dupin de Saint-Cyr, J. Lang, T. Schiex, Penalty logic and its link with Dempster-Shafer theory, in: *Proceedings of the 10th International Conference on Uncertainty in Artificial Intelligence*, 1994, pp. 204–211.
- [26] D. Dubois, H. Prade, Practical methods for constructing possibility distributions, *Int. J. Intell. Syst.* 31 (3) (2016) 215–239.
- [27] O. Kuzelka, J. Davis, S. Schockaert, Encoding Markov logic networks in possibilistic logic, in: *Proc. of the 31st Conf. on Uncertainty in Artificial Intelligence*, 2015, pp. 454–463.
- [28] M. Banerjee, D. Dubois, A simple logic for reasoning about incomplete knowledge, *International Journal of Approximate Reasoning* 55 (2) (2014) 639 – 653.

- [29] R. C. Moore, Semantical considerations on nonmonotonic logic, *Artificial Intelligence* 25 (1) (1985) 75–94.
- [30] J. Halpern, Y. Moses, Towards a theory of knowledge and ignorance: Preliminary report, in: K. R. Apt (Ed.), *Logics and Models of Concurrent Systems*, Vol. 13, Springer Berlin Heidelberg, 1985, pp. 459–476.
- [31] D. Dubois, H. Prade, Resolution principles in possibilistic logic, *International Journal of Approximate Reasoning* 4 (1) (1990) 1–21.
- [32] L. A. Zadeh, A theory of approximate reasoning, in: J. E. Hayes, D. Mitchie, L. I. Mikulich (Eds.), *Machine Intelligence*, Vol. 9, Ellis Horwood, 1979, pp. 149–194.
- [33] L. A. Zadeh, PRUF — a meaning representation language for natural languages, *International Journal of Man-Machine Studies* 10 (1978) 395–460.
- [34] H. Prade, Reasoning with fuzzy default values, in: *Proc. 15th IEEE Inter. Symp. Multiple-Valued Logic*, Kingston, On., Canada, IEEE, 1985, pp. 191–197.
- [35] L. A. Zadeh, Fuzzy sets and information granularity, in: R. Ragade, M. Gupta, R. Yager (Eds.), *Advances in Fuzzy Sets Theory and Applications*, North Holland, Amsterdam, 1979, pp. 3–18.
- [36] L. Fariñas del Cerro, H. Prade, Rough sets, twofold fuzzy sets and modal logic. fuzziness in indiscernibility and partial information, in: *The Mathematics of Fuzzy Systems*, Vol. 88 of *Interdisciplinary Systems Research Series*, Verlag TÜV Rheinland, 1986, pp. 103–120.
- [37] L. Fariñas del Cerro, A. Herzig, A modal analysis of possibility theory, in: *Proceedings of the International Workshop on Fundamentals of Artificial Intelligence Research*, 1991, pp. 11–18.
- [38] D. Lewis, Counterfactuals and comparative possibility, *Journal of Philosophical Logic* 2 (4) (1973) 418–446.
- [39] D. Dubois, Belief structures, possibility theory and decomposable measures on finite sets, *Computers and AI* 5 (1986) 403–416.
- [40] C. Boutilier, Modal logics for qualitative possibility theory, *International Journal of Approximate Reasoning* 10 (2) (1994) 173–201.
- [41] D. Dubois, H. Prade, C. Testemale, In search of a modal system for possibility theory, in: *Proceedings of the European Conference on Artificial Intelligence*, 1988, pp. 501–506.
- [42] G. J. Klir, D. Harmanec, On modal logic interpretation of possibility theory, *Int. J. of Uncertainty, Fuzziness and Knowledge-based Systems* 2 (2) (1994) 237–245.
- [43] M. Banerjee, D. Dubois, L. Godo, Possibilistic vs. relational semantics for logics of incomplete information, in: *Proceedings of the 15th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems*, 2014, pp. 335–344.
- [44] P. Hájek, *Metamathematics of Fuzzy Logic*, Trends in Logic, Springer, 2001.
- [45] T. Flaminio, L. Godo, E. Marchioni, On the logical formalization of possibilistic counterparts of states over n -valued Łukasiewicz events, *Journal of Logic and Computation* 21 (3) (2011) 429–446.
- [46] T. Flaminio, L. Godo, E. Marchioni, Reasoning about uncertainty of fuzzy events: an overview, in: P. Cintula, C. Fermueller, L. Godo, P. Hájek (Eds.), *Understanding Vagueness*, College Publications, 2011, pp. 367–400.
- [47] M. Baaz, Infinite-valued gödel logic with 0-1-projections and relativisations, in: P. Hájek (Ed.), *Gödel96: Logical Foundations of Mathematics, Computer Science, and Physics*, Vol. 6 of *Lecture Notes in Logic*, Springer-Verlag, Berlin, 1996, pp. 23–33.
- [48] P. Cintula, C. Noguera, *Modal Logics of Uncertainty with Two-Layer Syntax: A General Completeness Theorem*, 2014, pp. 124–136.
- [49] C. Liao, B. I. Lin, Proof methods for reasoning about possibility and necessity, *Int. J. Approx. Reasoning* 9 (4) (1993) 327–364.
- [50] D. Dubois, P. Hájek, H. Prade, Knowledge-driven versus data-driven logics, *J. Logic, Language and Information* 9 (2000) 65–89.
- [51] J. Y. Halpern, Defining relative likelihood in partially-ordered preferential structures, *Journal of Artificial Intelligence Research* 7 (1997) 1–24.
- [52] D. Dubois, E. Lorini, H. Prade, Nonmonotonic desires - A possibility theory viewpoint, in: R. Booth,

- G. Casini, S. Klarman, G. Richard, I. J. Varzinczak (Eds.), Proceedings of the International Workshop on Defeasible and Ampliative Reasoning, DARE@ECAI 2014, co-located with the 21st European Conference on Artificial Intelligence (ECAI 2014), Prague, Czech Republic, August 19, 2014., Vol. 1212 of CEUR Workshop Proceedings, CEUR-WS.org, 2014.
 URL <http://ceur-ws.org/Vol-1212/DARE-14-paper-1.pdf>
- [53] R. Reiter, A logic for default reasoning, *Artificial Intelligence* 13 (1-2) (1980) 81 – 132.
 - [54] R. C. Moore, Semantical considerations on nonmonotonic logic, *Artificial Intelligence* 25 (1) (1985) 75 – 94.
 - [55] J. Pearl, Probabilistic reasoning in intelligent systems: networks of plausible inference, Morgan Kaufmann, 1988.
 - [56] J. Pearl, System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning, in: Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge, 1990, pp. 121–135.
 - [57] H. Geffner, J. Pearl, Conditional entailment: Bridging two approaches to default reasoning, *Artificial Intelligence* 53 (2) (1992) 209–244.
 - [58] S. Benferhat, D. Dubois, H. Prade, Practical handling of exception-tainted rules and independence information in possibilistic logic, *Applied intelligence* 9 (2) (1998) 101–127.
 - [59] S. Benferhat, D. Dubois, H. Prade, Nonmonotonic reasoning, conditional objects and possibility theory, *Artificial Intelligence* 92 (1-2) (1997) 259–276.
 - [60] G. Schurz, Probabilistic semantics for delgrande’s conditional logic and a counterexample to his default logic, *Artificial Intelligence* 102 (1) (1998) 81 – 95.
 - [61] D. Lehmann, M. Magidor, What does a conditional knowledge base entail?, *Artificial Intelligence* 55 (1) (1992) 1–60.
 - [62] S. Schockaert, H. Prade, Interpolative reasoning with default rules, in: Proceedings of the International Joint Conference on Artificial Intelligence, 2013, pp. 1090–1096.
 - [63] D. Dubois, H. Prade, Conditional objects as nonmonotonic consequence relationships, *IEEE Transactions on Systems, Man and Cybernetics* 24 (12) (1994) 1724–1740.
 - [64] D. Lewis, Counterfactuals, Blackwell Publishing, 1973.
 - [65] L. Fariñas del Cerro, A. Herzig, J. Lang, From ordering-based nonmonotonic reasoning to conditional logics, *Artificial Intelligence* 66 (2) (1994) 375 – 393.
 - [66] V. Lifschitz, Minimal belief and negation as failure, *Artificial Intelligence* 70 (1-2) (1994) 53 – 72.
 - [67] D. Ciucci, D. Dubois, A modal theorem-preserving translation of a class of three-valued logics of incomplete information, *Journal of Applied Non-Classical Logics* 23 (4) (2013) 321–352.
 - [68] K. Bauters, S. Schockaert, M. De Cock, D. Vermeir, Characterizing and extending answer set semantics using possibility theory, *Theory and Practice of Logic Programming* (2015) 79–116.
 - [69] V. Lifschitz, Twelve definitions of a stable model, in: Proceedings of the 24th International Conference on Logic Programming, 2008, pp. 37–51.
 - [70] D. Pearce, Equilibrium logic, *Annals of Mathematics and Artificial Intelligence* 47 (2006) 3–41.
 - [71] L. Fariñas del Cerro, A. Herzig, E. I. Su, Capturing equilibrium models in modal logic, *Journal of Applied Logic* 12 (2) (2014) 192–207.
 - [72] L. Fariñas del Cerro, A. Herzig, E. I. Su, Epistemic equilibrium logic, in: Proceedings of the 24th International Joint Conference on Artificial Intelligence, 2015, pp. 2964–2970.
 - [73] M. Gelfond, Strong introspection, in: Proceedings of the 9th National Conference on Artificial Intelligence, 1991, pp. 386–391.
 - [74] M. Gelfond, New semantics for epistemic specifications, in: Proceedings of the International Conference on Logic Programming and Nonmonotonic Reasoning, 2011, pp. 260–265.
 - [75] K. Inoue, C. Sakama, Negation as failure in the head, *The Journal of Logic Programming* 35 (1) (1998) 39 – 78.
 - [76] K. Wagner, Bounded query classes, *SIAM Journal on Computing* 19 (5) (1990) 833–846.
 - [77] S. Buss, L. Hay, On truth-table reducibility to SAT and the difference hierarchy over NP, in: Proceedings of the Third Annual Conference on Structure in Complexity Theory, 1988, pp. 224–233.

- [78] J. Biskup, Selected results and related issues of confidentiality-preserving controlled interaction execution, in: Proc. Int. Sym. on Foundations of Information and Knowledge Systems, 2016, pp. 211–234.
- [79] S. D. Clercq, S. Schockaert, A. Nowé, M. D. Cock, Multilateral negotiation in boolean games with incomplete information using generalized possibilistic logic, in: Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, 2015, pp. 2890–2896.
- [80] J. Y. Halpern, G. Lakemeyer, Multi-agent only knowing, *Journal of Logic and Computation* 11 (1) (2001) 41–70.
- [81] V. Belle, G. Lakemeyer, Semantical considerations on multiagent only knowing, *Artificial Intelligence* 223 (2015) 1–26.
- [82] P. Cintula, C. Noguera, Modal logics of uncertainty with two-layer syntax: A general completeness theorem, in: U. Kohlenbach, P. Barceló, R. J. G. B. de Queiroz (Eds.), *Logic, Language, Information, and Computation - 21st International Workshop, WoLLIC 2014, Valparaíso, Chile, September 1-4, 2014. Proceedings*, Vol. 8652 of *Lecture Notes in Computer Science*, Springer, 2014, pp. 124–136.
- [83] S. Benferhat, D. Dubois, H. Prade, Representing default rules in possibilistic logic, in: *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR'92)*. Cambridge, MA, October 25-29, 1992., Morgan Kaufmann, 1992, pp. 673–684.