

Reasoning about Uncertainty and Explicit Ignorance in Generalized Possibilistic Logic

Didier Dubois¹ and Henri Prade² and Steven Schockaert³

Abstract. Generalized possibilistic logic (GPL) is a logic for reasoning about the revealed beliefs of another agent. It is a two-tier propositional logic, in which propositional formulas are encapsulated by modal operators that are interpreted in terms of uncertainty measures from possibility theory. Models of a GPL theory represent weighted epistemic states and are encoded as possibility distributions. One of the main features of GPL is that it allows us to explicitly reason about the ignorance of another agent. In this paper, we study two types of approaches for reasoning about ignorance in GPL, based on the idea of minimal specificity and on the notion of guaranteed possibility, respectively. We show how these approaches naturally lead to different flavours of the language of GPL and a number of decision problems, whose complexity ranges from the first to the third level of the polynomial hierarchy.

1 INTRODUCTION

Possibilistic logic [7] (PL) is a logic for reasoning with conjunctions of more or less certain 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). Inference is carried out by refutation using the following resolution principle

$$(\alpha \vee \beta, \lambda_1), (\neg\alpha \vee \gamma, \lambda_2) \vdash (\beta \vee \gamma, \min(\lambda_1, \lambda_2))$$

Inference in possibilistic logic thus remains close to inference in propositional logic, and in particular, 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 tentative ranking of propositional formulas. As such, it is closely related to the notion of epistemic entrenchment [11], as has been pointed out in [8]. This makes PL a natural vehicle for implementing strategies for belief revision [6] and managing inconsistency [3]. Along similar lines, there are close connections between PL and default reasoning in the sense of System P [14] and Rational Closure, which can be exploited to implement several forms of reasoning about rules with exceptions [4].

In many applications, a PL theory encodes the epistemic state of an agent. We then assume that all that the agent knows is contained in the knowledge base and its logical consequences (with the weights referring to the degree of epistemic entrenchment or the strength of belief). However, as a tool for epistemic reasoning, possibilistic logic is limited in at least two ways. First, given that a theory 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 the agent 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 we know is that one of these two situations holds, and in particular we know that the agent is not ignorant about the outcome of the coin flip. To express this in PL we would need to write a disjunction $(tails, 1) \vee (\neg tails, 1)$ which is not allowed in the language of PL. Second, PL does not allow us to encode information about the absence of knowledge, as it only relies on the assumption that the agent does not know α if α cannot be derived from the given theory. When reasoning about the beliefs of another agent, this assumption is not valid anymore, and we need to distinguish between situations where we know that the agent is ignorant about α and situations where we do not know whether the agent knows α .

Recently, we have proposed a generalized possibilistic logic (GPL) [10] which allows arbitrary propositional combinations of assertions of the form (α, λ) . This brings PL closer to modal logics for epistemic reasoning, and to emphasize this we will use a slightly different notation, writing e.g. $\mathbf{N}_1(tails) \vee \mathbf{N}_1(\neg tails)$ instead of $(tails, 1) \vee (\neg tails, 1)$. The aim of [10] was to show how GPL can naturally encode the semantics of answer set programming [12] and equilibrium logic [17]. Our aim in this paper is to develop GPL as a logic for reasoning about the revealed beliefs of another agent, with a GPL theory corresponding to a set of epistemic states, each of which is compatible with these revealed beliefs. In the classification of Aucher [1], it corresponds to the imperfect external point of view. This is similar to an auto-epistemic logic, except that in such a logic, an agent is supposed to reason about its own beliefs. One crucial difference, which has been pointed out in [13], is that when reasoning about one's own beliefs, it should not be possible to have $\mathbf{N}_1(\alpha) \vee \mathbf{N}_1(\beta)$ without either having $\mathbf{N}_1(\alpha)$ or $\mathbf{N}_1(\beta)$. GPL is closely related to modal logics for epistemic reasoning such as KD45 and S5. However, instead of using Kripke frames, the semantics of GPL are based on possibility distributions, which explicitly represent epistemic states. Apart from being a more intuitive way of capturing revealed beliefs, this has the advantage that (the strength of) belief can be naturally encoded as a graded notion. Another difference with traditional modal logics for epistemic reasoning is that we do not allow the modality \mathbf{N} to be nested and we do not allow objective formulas. In other words, we are not concerned with introspection nor with relating what the agent believes to what is objectively true.

In this paper, we investigate a number of different approaches for reasoning about the ignorance of another agent in GPL. After presenting some basic notions from possibility theory in the next section, Section 3 recalls the syntax and semantics of GPL, shows the soundness and completeness of an axiomatization for GPL and

¹ Université Paul Sabatier, CNRS, IRIT, France, email: didier.dubois@irit.fr

² Université Paul Sabatier, CNRS, IRIT, France, email: henri.prade@irit.fr

³ Cardiff University, UK, email: s.schockaert@cs.cardiff.ac.uk

shows that entailment checking in GPL is coNP-complete, as in propositional logic. Then in Section 4, we analyse how a modality based on the guaranteed possibility measure can be used for reasoning about ignorance. Adding this modality to the language does not increase its expressivity, but allows us to encode some formulas more compactly. As a result the computational complexity goes up to Θ_2^P , and to Π_3^P if a context-sensitive version of the modality is used. Then, in Section 5 we analyze how ignorance can be modelled based on minimal specificity. This gives rise to cautious and brave non-monotonic consequence relations, and a complexity of respectively Π_2^P and Σ_2^P , as in many formalisms for non-monotonic reasoning, such as answer set programming [12], equilibrium logic [17] and autoepistemic logic [16].

2 PRELIMINARIES

Consider a variable X which has an unknown value from some finite universe \mathcal{U} . In possibility theory [19, 9], available knowledge about the value of X is encoded as a $\mathcal{U} \rightarrow [0, 1]$ mapping π , 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. A possibility distribution π induces a possibility measure Π and a necessity measure N , defined for $A \subseteq \mathcal{U}$ as

$$\Pi(A) = \max_{u \in A} \pi(u) \quad N(A) = 1 - \Pi(\mathcal{U} \setminus A)$$

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 induced by π are the guaranteed possibility measure Δ and the potential necessity measure ∇ , defined for $A \subseteq \mathcal{U}$ as

$$\Delta(A) = \min_{u \in A} \pi(u) \quad \nabla(A) = 1 - \Delta(\mathcal{U} \setminus A)$$

Intuitively, $\Delta(A)$ reflects the extent to which all values in A are considered possible, while $\nabla(A)$ reflects the extent to which some values outside A are impossible.

3 GENERALIZED POSSIBILISTIC LOGIC

Let \mathcal{L} be the language of all propositional formulas, obtained from a finite set of atomic propositions $At = \{a, b, c, \dots\}$ using the connectives \neg and \wedge . As usual, we also use the abbreviations $\alpha \vee \beta \equiv \neg(\neg\alpha \wedge \neg\beta)$ and $\alpha \rightarrow \beta \equiv \neg(\alpha \wedge \neg\beta)$. Let $\Lambda_k = \{0, \frac{1}{k}, \frac{2}{k}, \dots, 1\}$ with $k \in \mathbb{N} \setminus \{0\}$ be the set of certainty degrees under consideration, and let $\Lambda_k^+ = \Lambda_k \setminus \{0\}$. We define the language $\mathcal{L}_{\mathbb{N}}^k$ of GPL with $k + 1$ certainty levels as follows:

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

The semantics of GPL are defined in terms of normalized possibility distributions over propositional interpretations, where possibility degrees are limited to Λ_k . Let \mathcal{P}_k be the set of all such Λ_k -valued possibility distributions. A model of a GPL formula is any possibility distribution from \mathcal{P}_k which satisfies:

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

where N is the necessity measure induced by π . As usual, π is called a model of a set of GPL formulas K , written $\pi \models K$, if π is a model of each formula in K . We write $K \models \Phi$, for K a set of GPL formulas and Φ a GPL formula, iff every model of K is also a model of Φ .

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 (i.e. it is considered more plausible that α is true than that α is false). 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, given the finiteness of the set of considered certainty degrees, means $N(\alpha) \leq \lambda - \frac{1}{k}$ and thus $\Pi(\neg\alpha) \geq 1 - \lambda + \frac{1}{k}$. We shall use the notation $\nu(\lambda) = 1 - \lambda + \frac{1}{k}$. Then, $\nu(\lambda) \in \Lambda_k^+$ iff $\lambda \in \Lambda_k^+$, and $\nu(\nu(\lambda)) = \lambda, \forall \lambda \in \Lambda_k^+$. We will also use the following abbreviation:

$$\mathbf{\Pi}_\lambda(\alpha) \equiv \neg\mathbf{N}_{\nu(\lambda)}(\neg\alpha)$$

Then $\mathbf{\Pi}_1(\alpha)$ 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 ($\Pi(\alpha) \geq \lambda$).

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, such 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:*

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

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 cannot fully exclude Wales has won its first round game, but lacks evidence suggesting their victory. We can also express comparative uncertainty in GPL, e.g.:

$$\bigvee_{i=1}^k \mathbf{\Pi}_{\frac{i}{k}}(won_1(wal)) \wedge \neg\mathbf{\Pi}_{\frac{i}{k}}(won_1(eng))$$

This 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, i.e. $\exists \lambda > 0. \Pi(won_1(wal)) \geq \lambda > \Pi(won_1(eng))$.

The following axiomatization of GPL has been introduced in [10]:

- (PL) The Hilbert axioms of classical logic
- (K) $\mathbf{N}_\lambda(\alpha \rightarrow \beta) \rightarrow (\mathbf{N}_\lambda(\alpha) \rightarrow \mathbf{N}_\lambda(\beta))$
- (N) $\mathbf{N}_1(\top)$
- (D) $\mathbf{N}_\lambda(\alpha) \rightarrow \mathbf{\Pi}_1(\alpha)$
- (W) $\mathbf{N}_{\lambda_1}(\alpha) \rightarrow \mathbf{N}_{\lambda_2}(\alpha)$, if $\lambda_1 \geq \lambda_2$

with modus ponens as the only inference rule. Note in particular that when λ is fixed we get a fragment of the modal logic KD. The case where $k = 1$ coincides with the Meta-Epistemic Logic (MEL) that was introduced in [2]. This simpler logic, a fragment of KD with no nested modalities nor objective formulas, can express full certainty and full ignorance only and its semantics is in

terms of non-empty subsets of interpretations. Note that in MEL, we have $\Pi_1(\alpha) \equiv \neg \mathbf{N}_1(\neg\alpha)$ whereas in general we only have $\Pi_1(\alpha) \equiv \neg \mathbf{N}_{\frac{1}{k}}(\neg\alpha)$.

Proposition 1 (Soundness and completeness). *Let K be a GPL theory and Φ a GPL formula. It holds that $K \models \Phi$ iff Φ can be derived from K using modus ponens and the axioms **(PL)**, **(K)**, **(N)**, **(D)**, **(W)**.*

Proof. Soundness of the axioms can readily be verified. Completeness was shown in [10] in a constructive way. Here, we present a more elegant proof, which relies on the completeness of the Hilbert axioms for propositional logic and extends the one given for the case $k = 1$ in [2]. Let $\mathcal{L}_{CNF} \subseteq \mathcal{L}$ be a finite set of all formulas over At which are in conjunctive-normal form. Since At is finite, it follows that \mathcal{L}_{CNF} is finite as well (if we disallow the same literal to appear more than once in a clause and we disallow the same clause to appear more than once in the formula). Without loss of generality, we can assume that every sub-formula of the form $\mathbf{N}_\lambda(\alpha)$ occurring in K and Φ is such that $\alpha \in \mathcal{L}_{CNF}$ (since **(K)**, **(W)** and **(N)** imply that $\mathbf{N}_\lambda(\alpha) \equiv \mathbf{N}_\lambda(\beta)$ if $\alpha \equiv \beta$).

We can see formulas in GPL as propositional formulas which are built from the set of atoms $F = \{\mathbf{N}_\lambda(\alpha) \mid \alpha \in \mathcal{L}_{CNF}, \lambda \in \Lambda_k^+\}$. Let Ω^* be the set of all propositional interpretations over F . Let K^* be the propositional knowledge base over F containing all formulas in K , as well as:

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

where $\text{cnf}(\alpha \rightarrow \beta)$ is an arbitrary (but fixed) formula from \mathcal{L}_{CNF} which is equivalent to $\alpha \rightarrow \beta$, and similar for $\text{cnf}(\neg\alpha)$. We then clearly have that Φ can be derived from K in GPL iff Φ can be derived from K^* in propositional logic.

To finish the proof, we show that there exists a bijection between the set of propositional models of K^* (seen as a theory in propositional logic) and the set of possibilistic models of K (seen as a GPL theory). First note that with every model I of K^* , we can consistently associate, due to axiom **(W)**, a set-function $g_I : 2^{\Omega} \rightarrow \Lambda_k^+$ defined for $\alpha \in \mathcal{L}_{CNF}$ 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 $\mathbf{N}_\lambda(\perp) \rightarrow \mathbf{N}_{\frac{1}{k}}(\perp)$ ($\equiv \neg \mathbf{N}_{\frac{1}{k}}(\perp) \rightarrow \neg \mathbf{N}_\lambda(\perp)$) for every $\lambda \in \Lambda_k^+$ (as an instantiation of **(W)**) we know that $\{\lambda \mid I \models \mathbf{N}_\lambda(\perp)\} = \emptyset$.
- 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 **(K)**, using **(N)**, **(W)**). 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 $\alpha, \beta \in \mathcal{L}_{CNF}$. 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 the instantiation of **(K)** on the tautology $\beta \rightarrow (\alpha \rightarrow (\alpha \wedge \beta))$ we find $I \models \mathbf{N}_\lambda(\text{cnf}(\alpha \rightarrow$

$(\alpha \wedge \beta))$). Using another instantiation of **(K)** we find from $I \models \mathbf{N}_\lambda(\text{cnf}(\alpha \rightarrow (\alpha \wedge \beta)))$ and $I \models \mathbf{N}_\lambda(\alpha)$ that $I \models \mathbf{N}_\lambda(\text{cnf}(\alpha \wedge \beta))$. 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 $g_I(\llbracket \alpha \wedge \beta \rrbracket) \geq \min(g_I(\llbracket \alpha \rrbracket), g_I(\llbracket \beta \rrbracket))$.

Every set-function satisfying these four criteria is a necessity measure [9], and this necessity measure uniquely identifies a normalized possibility distribution π such that $\pi(\omega) = 1 - g_I(\Omega \setminus \{\omega\})$, which by construction will be a model of K . Conversely, it is easy to see that every model π of K corresponds to a 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

The following formulas can be proven as theorems in GPL:

$$\begin{aligned} & \mathbf{N}_{\lambda_1}(\alpha) \wedge \mathbf{N}_{\lambda_2}(\alpha \rightarrow \beta) \rightarrow \mathbf{N}_{\min(\lambda_1, \lambda_2)}(\beta) \\ & \mathbf{\Pi}_{\lambda_1}(\alpha) \wedge \mathbf{N}_{\lambda_2}(\alpha \rightarrow \beta) \rightarrow \mathbf{\Pi}_{\lambda_1}(\beta), \text{ if } \lambda_2 \geq \nu(\lambda_1) \end{aligned}$$

The following GPL theorems are the counterpart of well-known properties of necessity and possibility measures:

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

Finally, we show that inference in GPL is not harder than inference in classical logic. Let us define a *meta-literal* as an expression of the form $\mathbf{N}_\lambda(\alpha)$ or $\mathbf{\Pi}_\lambda(\alpha)$. A *meta-clause* is an expression of the form $\Phi_1 \vee \dots \vee \Phi_n$ with each Φ_i a meta-literal. Similarly, a *meta-term* is an expression of the form $\Phi_1 \wedge \dots \wedge \Phi_n$ with each Φ_i a meta-literal.

Proposition 2. *The problem of deciding whether $\Phi \models \Psi$, with Φ and Ψ GPL formulas, is coNP-complete⁴.*

Proof. Hardness: follows straightforwardly from the coNP-completeness of entailment checking in propositional logic.

Membership: we show that checking the satisfiability of a GPL formula is in NP. 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:

$$\mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n) \wedge \mathbf{\Pi}_{\mu_1}(\beta_1) \wedge \dots \wedge \mathbf{\Pi}_{\mu_m}(\beta_m) \quad (1)$$

We know from PL that $\mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n)$ has a unique least specific 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, for each j and $\theta \in \Lambda_k^+$ such that $\mu_j \geq \nu(\theta)$ the following formula needs to be consistent:

$$\bigwedge \{\alpha_i \mid \lambda_i \geq \theta\} \wedge \beta_j \quad (2)$$

When guessing (1), we can also guess an interpretation for each of the SAT instances (2). We can then verify in polynomial time that they are indeed models of the corresponding propositional formulas. \square

As to the possible kinds of conclusions that can be inferred from a GPL base K regarding a propositional formula α , with only two certainty levels, one can already distinguish the following cases [2]:

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

⁴ As usual, complexity results are stated w.r.t. the size of the formulas involved (i.e. the number of occurrences of atomic propositions).

- $K \models \mathbf{N}_1(\neg\alpha)$ means that we know that the agent knows that α is false.
- $K \models \mathbf{N}_1(\alpha) \vee \mathbf{N}_2(\neg\alpha)$, $K \not\models \mathbf{N}_1(\alpha)$ 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 .

These five cases can be contrasted with 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. The underlying reason is that theories in classical logic correspond to a single epistemic state, and in this sense, we can think of classical logic as a mechanism for reasoning about the beliefs of a given agent, or indeed about one's own beliefs. In contrast, a GPL theory corresponds to a set of epistemic states, and is useful for reasoning about what is known about the beliefs of another agent.

4 GPL WITH 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. $conj_\omega = \bigwedge_{\omega \models a} a \wedge \bigwedge_{\omega \not\models a} \neg a$. Then we define:

$$\Delta_\lambda(\alpha) = \bigwedge_{\omega \in \llbracket \alpha \rrbracket} \mathbf{\Pi}_\lambda(conj_\omega) \quad \nabla_\lambda(\alpha) = \neg \Delta_{\nu(\lambda)}(\neg\alpha)$$

Some useful theorems using the modalities Δ and ∇ include:

$$\begin{aligned} \Delta_{\lambda_1}(\alpha \wedge \beta) \wedge \Delta_{\lambda_2}(\neg\alpha \wedge \gamma) &\rightarrow \Delta_{\min(\lambda_1, \lambda_2)}(\beta \wedge \gamma) \\ \nabla_{\lambda_1}(\alpha \wedge \beta) \wedge \Delta_{\lambda_2}(\neg\alpha \wedge \gamma) &\rightarrow \nabla_{\lambda_1}(\beta \wedge \gamma), \text{ if } \lambda_2 \geq \nu(\lambda_1) \end{aligned}$$

$$\begin{aligned} \Delta_\lambda(\alpha) \wedge \Delta_\lambda(\beta) &\equiv \Delta_\lambda(\alpha \vee \beta) & \nabla_\lambda(\alpha) \wedge \nabla_\lambda(\beta) &\leftarrow \nabla_\lambda(\alpha \vee \beta) \\ \Delta_\lambda(\alpha) \vee \Delta_\lambda(\beta) &\rightarrow \Delta_\lambda(\alpha \wedge \beta) & \nabla_\lambda(\alpha) \vee \nabla_\lambda(\beta) &\equiv \nabla_\lambda(\alpha \wedge \beta) \end{aligned}$$

Using the modality Δ , for any possibility distribution π over Ω , we can easily define a GPL theory which has π as its only model. 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 theory Φ_π as:

$$\Phi_\pi = \bigwedge_{i=1}^k \mathbf{N}_{\nu(\frac{i}{k})}(\alpha_i) \wedge \Delta_{\frac{i}{k}}(\alpha_i). \quad (3)$$

In this equation, the degree of possibility of each $\omega \in \llbracket \alpha_i \rrbracket$ is defined by inequalities from above and from below. Indeed, $\Delta_{\frac{i}{k}}(\alpha_i)$ 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 model of Φ_π . 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 theory. Conceptually, the construction of Φ_π relates to the notion of ‘‘only knowing’’ from Levesque [15]. 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 (3), $\alpha_{k-j+1} = \dots \alpha_k = \beta$ and we obtain $\Phi_\pi \equiv \Delta_1(\beta) \wedge \mathbf{N}_{\frac{j}{k}}(\beta) \wedge \Delta_{\frac{k-j}{k}}(\top)$.

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

Since the modality Δ was introduced as an abbreviation, allowing this modality has no impact on the expressivity 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 theory Φ_S . This is reflected in the following complexity result. We will call formulas in which the modalities Δ and ∇ occur GPL $^\Delta$ formulas. It is not hard to see that $\mathbf{N}_\lambda(\alpha) \wedge \Delta_\mu(\beta)$, with $\mu \geq \nu(\lambda)$ is consistent iff α is consistent and $\beta \models \alpha$. Hence, checking the satisfiability of an arbitrary GPL $^\Delta$ formula cannot be in NP (unless the polynomial hierarchy collapses), which means that entailment checking cannot be in coNP. Recall that 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 for one query to the NP-oracle cannot be used to formulate another query for the NP-oracle [18]). Moreover, allowing two rounds of parallel queries does not lead to an increased complexity ([5], Theorem 9). We will show that $\Phi \models \Psi$ can be decided in this way.

Proposition 3. *The problem of deciding whether $\Phi \models \Psi$, for Φ and Ψ two GPL $^\Delta$ formulas, is Θ_2^P -complete.*

Proof. Hardness: A standard Θ_2^P -complete problem is the following. Let $\alpha_1, \dots, \alpha_n$ be propositional formulas. Decide whether the smallest i for which α_i is unsatisfiable is an odd number. Without loss of generality, we can assume that n is odd. Now consider the following GPL $^\Delta$ formula:

$$\begin{aligned} \Psi \equiv & \neg \mathbf{\Pi}_1(\alpha_1) \vee (\mathbf{\Pi}_1(\alpha_1) \wedge \mathbf{\Pi}_1(\alpha_2) \wedge \neg \mathbf{\Pi}_1(\alpha_3)) \\ & \vee \dots \vee (\mathbf{\Pi}_1(\alpha_1) \wedge \dots \wedge \mathbf{\Pi}_1(\alpha_{n-1}) \wedge \neg \mathbf{\Pi}_1(\alpha_n)) \end{aligned}$$

It is easy to show that $\Delta_1(\top) \models \Psi$ iff the smallest i for which α_i is unsatisfiable is odd.

Membership: Since $\Phi \models \Psi$ 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:

$$\begin{aligned} & \mathbf{N}_{\frac{1}{k}}(\alpha_1^1), \dots, \mathbf{N}_{\frac{1}{k}}(\alpha_{n_1}^1), \mathbf{N}_{\frac{2}{k}}(\alpha_1^2), \dots, \mathbf{N}_1(\alpha_1^k), \dots, \mathbf{N}_1(\alpha_{n_k}^k) \\ & \mathbf{\Pi}_{\frac{1}{k}}(\beta_1^1), \dots, \mathbf{\Pi}_{\frac{1}{k}}(\beta_{m_1}^1), \mathbf{\Pi}_{\frac{2}{k}}(\beta_1^2), \dots, \mathbf{\Pi}_1(\beta_1^k), \dots, \mathbf{\Pi}_1(\beta_{m_k}^k) \\ & \Delta_{\frac{1}{k}}(\gamma_1^1), \dots, \Delta_{\frac{1}{k}}(\gamma_{p_1}^1), \Delta_{\frac{2}{k}}(\gamma_1^2), \dots, \Delta_1(\gamma_1^k), \dots, \Delta_1(\gamma_{p_k}^k) \\ & \nabla_{\frac{1}{k}}(\delta_1^1), \dots, \nabla_{\frac{1}{k}}(\delta_{r_1}^1), \nabla_{\frac{2}{k}}(\delta_1^2), \dots, \nabla_1(\delta_1^k), \dots, \nabla_1(\delta_{r_k}^k) \end{aligned}$$

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:

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

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 $\mathbf{\Pi}_{w_i}(\beta_i^*) \equiv \mathbf{\Delta}_{w_i}(\beta_i^*)$. 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 $\mathbf{\nabla}_{z_i}(\delta_i^*) \equiv \mathbf{N}_{z_i}(\delta_i^*)$.

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

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

This latter formula is satisfiable 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 $u_i \geq \nu(z_j)$

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. Given that β_i^* and $\neg\delta_j^*$ have a unique model, it is clear that the remaining conditions can be verified in polynomial time. \square

The modality $\mathbf{\Delta}$ allows us to express limits 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 $\mathbf{N}_1(\text{won}_3(\text{wal})) \wedge \mathbf{\Delta}_1(\text{won}_3(\text{wal}))$, 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, we propose the following variant of the $\mathbf{\Delta}$ modality:

$$\mathbf{\Delta}_\lambda^X(\alpha) = \bigwedge_{\omega \in \llbracket \alpha \rrbracket} \mathbf{\Pi}_\lambda(\text{conj}_\omega^X)$$

where $X \subseteq \text{At}$ is a set of atoms and conj_ω^X is the restriction of conj_ω to those literals about atoms in X , i.e. $\text{conj}_\omega^X = \bigwedge \{x \mid x \in X, \omega \models x\} \wedge \bigwedge \{\neg x \mid x \in X, \omega \models \neg x\}$. Note that $\mathbf{\Delta}_\lambda(\alpha) \equiv \mathbf{\Delta}_\lambda^{\text{At}}(\alpha)$. For example, in the scenario from Example 2, instead of asserting $\mathbf{\Delta}_1(\text{won}_3(\text{wal}))$, we can assert $\mathbf{\Delta}_1^X(\text{won}_3(\text{wal}))$, with $X = \{\text{plays}_3(x, y) \mid x, y \in T\} \cup \{\text{won}_3(x) \mid x \in T\}$ the set of all atoms about round 3 of the championship.

The extra flexibility of the $\mathbf{\Delta}^X$ modality comes at the price of a higher computational complexity. 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 . We will call formulas in which modalities of the form $\mathbf{\Delta}^X$ occur GPL $_{\mathbf{\Delta}}^{\mathbf{\Delta}}$ formulas.

Proposition 4. The problem of deciding whether $\Phi \models \Psi$, for Φ and Ψ two GPL $_{\mathbf{\Delta}}^{\mathbf{\Delta}}$ formulas, is Π_3^P -complete.

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

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

This means that satisfiability checking in GPL $_{\mathbf{\Delta}}^{\mathbf{\Delta}}$ 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 $_{\mathbf{\Delta}}^{\mathbf{\Delta}}$ formula Ψ is satisfiable. Similarly as in the proof of Proposition 3, we can guess an implicant of Ψ of the following form:

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

where X_1, \dots, X_m are sets of atoms. We give a Σ_2^P procedure for checking that Θ is not satisfiable: select a β_i , guess a model ω of β_i , and verify 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

5 IGNORANCE AS MINIMAL SPECIFICITY

A possibility distribution π_1 in a universe U is called less specific than a possibility distribution π_2 in U , written $\pi_1 \preceq \pi_2$ if $\pi_1(u) \geq \pi_2(u)$ for all u in U . The relation \preceq defines a partial order on the set of models of a GPL theory 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 $K \models_{br} \Phi$ iff Φ is satisfied by a minimally specific model of K .
- We say that Φ is a cautious consequence of K , written $K \models_{cau} \Phi$ iff Φ is satisfied by all minimally specific models of K .

In standard possibilistic logic, every theory K has a least specific model π_K . Moreover, it holds that $K \models (\alpha, \lambda)$ iff $N_K(\alpha) \geq \lambda$, for N_K the necessity measure induced by π_K . In other words, in standard possibilistic logic, the entailment relations \models , \models_{br} and \models_{cau} coincide. In GPL, this is no longer the case.

Example 3. The formula $\mathbf{N}_1(a) \vee \mathbf{N}_1(b)$, for two atomic variables a and b , has two minimally specific models π_a and π_b defined as:

$$\pi_a(\omega) = \begin{cases} 0 & \text{if } \omega \models \neg a \\ 1 & \text{otherwise} \end{cases} \quad \pi_b(\omega) = \begin{cases} 0 & \text{if } \omega \models \neg b \\ 1 & \text{otherwise} \end{cases} \quad (4)$$

This already shows that \models and \models_{br} do not coincide, as e.g. $\mathbf{N}_1(a) \vee \mathbf{N}_1(b) \models_{br} \mathbf{N}_1(a)$ while clearly $\mathbf{N}_1(a) \vee \mathbf{N}_1(b) \not\models \mathbf{N}_1(a)$. To see why \models and \models_{cau} do not coincide, note that $\mathbf{N}_1(a) \vee \mathbf{N}_1(b) \models_{cau} \mathbf{\Pi}_1(c) \wedge \mathbf{\Pi}_1(\neg c)$ while $\mathbf{N}_1(a) \vee \mathbf{N}_1(b) \not\models \mathbf{\Pi}_1(c) \wedge \mathbf{\Pi}_1(\neg c)$.

Reasoning about what is true in all minimally specific models, as opposed to all models, is similar to making a 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 the following GPL theory

$$K = (\mathbf{N}_1(\alpha) \wedge \mathbf{\Pi}_1(\beta) \wedge \mathbf{\Pi}_1(\neg\beta)) \vee (\mathbf{N}_1(\beta) \wedge \mathbf{\Pi}_1(\alpha) \wedge \mathbf{\Pi}_1(\neg\alpha))$$

This corresponds to a scenario in which we know that the agent either knows α and is ignorant about β , or knows β and is ignorant about α . It holds that $K \not\models_{cau} \mathbf{N}_1(\alpha) \vee \mathbf{N}_1(\neg\alpha)$, i.e. we cannot conclude that the agent knows about α , and $K \not\models_{cau} \mathbf{\Pi}_1(\alpha) \wedge \mathbf{\Pi}_1(\neg\alpha)$, i.e. we cannot conclude that the agent is ignorant about α either.

Reasoning about minimally specific models is more expensive than reasoning about what is true for all models of a GPL theory (unless the GPL_R^{Δ} formulas are allowed in the latter case).

Proposition 5. Let Φ and Ψ be two GPL formulas. The problem of checking whether $\Phi \models_{cau} \Psi$ is Π_2^P -complete.

Proof. **Hardness:** it is easy to show that checking the validity of the QBF formula $\forall X \exists Y. \phi(X, Y)$ is equivalent to checking whether $\Phi \models_{cau} \Psi$ for Φ and Ψ defined as follows:

$$\Phi \equiv \bigwedge_{x \in X} (\mathbf{N}_1(x) \vee \mathbf{N}_1(\neg x)) \quad \Psi \equiv \mathbf{\Pi}_1(\phi(X, Y))$$

Membership: We now present a Σ_2^P algorithm for checking that β is false in at least one minimally specific model of α .

The GPL formula Φ is equivalent to a disjunction of meta-terms. In polynomial time, we can guess such a term:

$$\mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n) \wedge \mathbf{\Pi}_{\mu_1}(\beta_1) \wedge \dots \wedge \mathbf{\Pi}_{\mu_m}(\beta_m) \quad (5)$$

Using an NP-oracle we can check that this formula is consistent. We then know that it has a unique least specific model π_1 , and we can check using our oracle that Ψ is false in this model. Indeed, a meta-literal $\mathbf{N}_{\theta}(\gamma)$ occurring in Ψ is true in π_1 if it can be derived from $\mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n)$. Similarly, a meta-literal $\mathbf{\Pi}_{\theta}(\gamma)$ occurring in Ψ is true in π_1 unless $\mathbf{N}_{\nu(\theta)}(\neg\gamma)$ can be derived from $\mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n)$.

What remains to be verified is that there does not exist another consistent meta-term which has a model π_2 that is strictly less specific than π_1 . To check this, we define a GPL theory K' as follows. Without loss of generality, we can assume that K does not contain any occurrences of \rightarrow and that all negations in K occur inside a modality. For each meta-term $\mathbf{N}_{\nu}(\gamma)$ occurring in K , we test whether

$$\mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n) \models \mathbf{N}_{\mu}(\gamma) \quad (6)$$

If this is not the case, we replace $\mathbf{N}_{\nu}(\gamma)$ by \perp . Note that the resulting theory K' is consistent, since none of the meta-literals occurring in (5) will have been replaced. By replacing a meta-literal $\mathbf{N}_{\nu}(\gamma)$ by \perp , we potentially reduce the set of models of the knowledge base. However, because (6) by construction does not hold, none of these models can be less specific than π_1 .

Finally, we test whether $K' \models \mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n)$. If this is the case, then none of the models of K' , and by extension of K , can be less specific than π_1 . If this is not the case, the guess in (5) did not induce a minimally specific possibility distribution of K . Indeed, because (6) holds for all the meta-literals $\mathbf{N}_{\mu}(\gamma)$ in K' , any model of K' which is not a model of $\mathbf{N}_{\lambda_1}(\alpha_1) \wedge \dots \wedge \mathbf{N}_{\lambda_n}(\alpha_n)$ has to be strictly less specific than π_1 . \square

We can also show that checking $\Phi \models_{cau} \Psi$ is Π_2^P -complete for GPL^{Δ} formulas, and Π_4^P -complete for GPL_R^{Δ} formulas. We omit the proofs due to space constraints. To characterize the complexity of brave reasoning, note that $\Phi \models_{cau} \Psi$ iff it is not the case that $\Phi \models_{br} \neg\Psi$. Hence we immediately get the following result.

Corollary 1. Let Φ and Ψ be two GPL formulas. The problem of checking whether $\Phi \models_{br} \Psi$ is Σ_2^P -complete.

6 CONCLUSION

Generalized possibilistic logic allows us to reason about the epistemic state of another agent. A particular feature of this type of reasoning is that we can draw conclusions about what that agent does not know. This requires that we encode information about the limits of the agent's knowledge. We have discussed several ways in which this can be accomplished in possibilistic logic, based on guaranteed possibility measures and on the principle of minimal specificity. We have shown that this leads to a range of decision problems, which are complete for the classes coNP , Θ_2^P , Π_2^P , Σ_2^P and Π_3^P .

REFERENCES

- [1] G. Aucher, 'An internal version of epistemic logic', *Studia Logica*, **94**(1), 1–22, (2010).
- [2] M. Banerjee and D. Dubois, 'A simple logic for reasoning about incomplete knowledge', *International Journal of Approximate Reasoning*, **55**(2), 639–653, (2014).
- [3] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, and H. Prade, 'Inconsistency management and prioritized syntax-based entailment', in *Proceedings of the International Joint Conference on Artificial Intelligence*, volume 93, pp. 640–645, (1993).
- [4] S. Benferhat, D. Dubois, and H. Prade, 'Practical handling of exception-tainted rules and independence information in possibilistic logic', *Applied Intelligence*, **9**(2), 101–127, (1998).
- [5] S.R. Buss and 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*, pp. 224–233, (1988).
- [6] D. Dubois, J. Lang, and H. Prade, 'Automated reasoning using possibilistic logic: semantics, belief revision, and variable certainty weights', *IEEE Transactions on Knowledge and Data Engineering*, **6**(1), 64–71, (1994).
- [7] D. Dubois, J. Lang, and H. Prade, 'Possibilistic logic', in *Handbook of Logic in Artificial Intelligence and Logic Programming*, ed., D. Nute D. Gabbay, C. Hogger J. Robinson, volume 3, 439–513, Oxford University Press, (1994).
- [8] D. Dubois and H. Prade, 'Epistemic entrenchment and possibilistic logic', *Artificial Intelligence*, **50**(2), 223–239, (1991).
- [9] D. Dubois and H. Prade, 'Possibility theory: qualitative and quantitative aspects', in *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, eds., D.M. Gabbay and P. Smets, volume 1, 169–226, Kluwer Academic, (1998).
- [10] D. Dubois, H. Prade, and S. Schockaert, 'Stable models in generalized possibilistic logic', in *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning*, pp. 519–529, (2012).
- [11] P. Gärdenfors, *Knowledge in Flux*, MIT Press, 1988.
- [12] M. Gelfond and V. Lifschitz, 'The stable model semantics for logic programming', in *Proceedings of the Fifth International Conference and Symposium on Logic Programming*, pp. 1081–1086, (1988).
- [13] J. Halpern and Y. Moses, 'Towards a theory of knowledge and ignorance: Preliminary report', in *Logics and Models of Concurrent Systems*, ed., Krzysztof R. Apt, volume 13, 459–476, Springer Berlin Heidelberg, (1985).
- [14] S. Kraus, D. Lehmann, and M. Magidor, 'Nonmonotonic reasoning, preferential models and cumulative logics', *Artificial Intelligence*, **44**(1-2), 167–207, (1990).
- [15] H. J. Levesque, 'All I know: A study in autoepistemic logic', *Artificial Intelligence*, **42**, 263–309, (1990).
- [16] R. C. Moore, 'Semantical considerations on nonmonotonic logic', *Artificial Intelligence*, **25**(1), 75–94, (1985).
- [17] D. Pearce, 'Equilibrium logic', *Annals of Mathematics and Artificial Intelligence*, **47**, 3–41, (2006).
- [18] K. Wagner, 'Bounded query classes', *SIAM Journal on Computing*, **19**(5), 833–846, (1990).
- [19] L.A. Zadeh, 'Fuzzy sets as a basis for a theory of possibility', *Fuzzy Sets and Systems*, **1**, 3–28, (1978).