A Simple Doxastic Deontic Logic for Norm-Guided Decision Making

Thorsten Engesser[0000-0002-2129-3207]* and Agata Ciabattoni[0000-0001-6947-8772]

TU Wien
thorsten@logic.at, agata@logic.at

Abstract. Making decisions despite conflicting norms and incomplete or unreliable information is a fundamental challenge for autonomous systems. We introduce a simple doxastic deontic logic for this setting: a classically reducible fragment of Chellas' Minimal Deontic Logic, extended with explicit conditional norms and combined with multi-agent KD45, so that norms can depend on agents' beliefs about both facts and norms. On this logic we define the Doxastic Norm Compliance Optimization Problem, where an agent chooses a decision minimizing weighted norm violations. We distinguish subjective optimization (relative to the agent's beliefs) from objective optimization (relative to the actual facts). We give conditions under which (i) the two coincide and (ii) optimal decision-making can be reduced to weighted partial MaxSAT in polynomial time.

Keywords: Deontic logic · Doxastic logic · Norm compliance

1 Introduction

AI systems increasingly operate under legal, organizational, or ethical norms specifying obligatory and forbidden actions. However, conflicting norms and partial or unreliable beliefs can make it difficult to determine which norms apply.

Norm-guided decision making therefore raises two key challenges. First, full compliance with norms may be impossible. In such cases, an intelligent agent should select the action that minimizes the severity of norm violations (the lesser evil). Second, many norms are conditional, and their applicability may depend on the agent's beliefs about the world. Norm-guided decisions must therefore account not only for the norms themselves, but also for the agent's epistemic state.

Existing approaches address different aspects of these challenges but do not provide a unified framework for doxastic normative reasoning that is simple, expressive and suitable for automated decision making. Deontic logics remain the principal tool to reason about norms [10, 11]. Preference-based deontic logics [3, 25] elegantly handle conditional norms and contrary-to-duty (CTD) obligations[1] but struggle with factual detachment, i.e., inferring a concrete obligation from a conditional norm given the relevant facts, e.g. [8]. I/O logics [20], instead, naturally support factual detachment, and in their constrained variants [21], certain conflicts and CTD reasoning, but not

* Corresponding author.

[1] Obligations that come into force when other obligations are violated.

higher-order norms. Moreover, these formalisms do not work well in the presence of dilemmas (where an action is both obligatory and prohibited) or belief-dependent norms. Logic-programming approaches (e.g. [27, 13, 17]), including deontic extensions of Answer Set Programming (ASP) (via modal operators [1, 12] or atomic norms [6]), offer practical tools but have not been combined with doxastic reasoning.

Beliefs and norms have already been combined in agent-oriented frameworks: BOID integrates beliefs, obligations, intentions, and desires and studies their interaction for different agent types [5], while BIO combines beliefs, intentions, and obligations within a defeasible-logic framework [14]. These approaches focus on the interaction and resolution of conflicts among different mental attitudes. Our focus is instead on how beliefs determine the applicability of conditional norms and how an agent can optimize its decisions when full norm compliance is impossible. In this paper, we develop a logical framework tailored to this task, with a lightweight deontic component, rather than aiming for a comprehensive theory of agent deliberation. Our starting point is a restricted fragment of Chellas' Minimal Deontic Logic [7], in which obligations apply only to atomic formulas. We show that reasoning in this fragment reduces to classical propositional reasoning. We then enrich it with explicitly represented conditional norms and factual detachment, obtaining a Lightweight Deontic Logic (LDL) that captures key aspects of conditional normative reasoning, including higher-order norms for expressing obligations that govern when other obligations apply. To capture the role of beliefs in norm applicability, we introduce LDLB, which combines LDL with multi-agent doxastic logic [16, 15], allowing explicit norms to depend on agents' beliefs, while agents themselves may reason about both factual information and their beliefs about norms. Building on LDLB, we define the Doxastic Norm Compliance Optimization Problem (DNCOP), in which an acting agent must choose among alternative decisions while minimizing violations of weighted norms, adopting a sort of utilitarian perspective [22] that aggregates norm violations into a single quantitative measure of harm.

The resulting framework supports encoding CTDs and exceptions through optimization of weighted norms, while distinguishing optimization based on the agent's subjective beliefs from optimization based on the objective state of the world. We identify natural locality and truthfulness conditions under which these two perspectives coincide; under these conditions, computing optimal decisions reduces in polynomial time to weighted partial MaxSAT [2], allowing existing optimization technology to be used directly.

LDLB combines logical transparency with computational efficiency: norms remain explicitly represented, beliefs affect their applicability, and decisions are obtained through a well-defined optimization problem. Its modular design also allows integration with existing epistemic logics [16, 15, 18, 9].

2 Preliminaries: Chellas' Minimal Deontic Logic

Chellas' Minimal Deontic Logic (MDL) [7] is often presented as a response to the deontic paradoxes that have beset the first formalisation of deontic logic—Standard Deontic Logic [29] (i.e., the normal modal logic KD)—and as a system that tolerates deontic conflicts. MDL is a very weak system whose language $\mathcal{L}_{\text{MDL}}$ is defined as

$\varphi := p \in \mathit{Prop} \mid \neg\varphi \mid \varphi \wedge \varphi \mid O\varphi$, where (for the purposes of this paper) *Prop* is a finite set of (classical) atoms.

$O\varphi$ means "φ is obligatory". We define $F\varphi := O\neg\varphi$ ("φ is forbidden") and $P\varphi := \neg F\varphi$ ("φ is permitted"). Boolean connectives and constants are as usual.

Formulas from $\mathcal{L}_{\mathrm{MDL}}$ are evaluated over *neighborhood models* $\mathcal{M} = (W, N, V)$, where W is a set of *worlds*, $N : W \to 2^{2^W}$ is a *neighborhood function*, and $V : W \to 2^{\mathit{Prop}}$ assigns to each world a propositional *valuation*. As frame properties, Chellas' semantics assumes that $\emptyset \notin N(w)$ for all $w \in W$, and that $N(w)$ is closed under supersets. That is, whenever $U \in N(w)$ and $U \subseteq X \subseteq W$, then $X \in N(w)$. Truth in a world w of $\mathcal{M}$ is defined as follows:

$$
\begin{array}{lll}
\mathcal{M}, w \Vdash_{\mathrm{MDL}} p & \text{iff} & p \in V(w), \\
\mathcal{M}, w \Vdash_{\mathrm{MDL}} O\varphi & \text{iff} & \{w' \mid \mathcal{M}, w' \Vdash_{\mathrm{MDL}} \varphi\} \in N(w), \\
\mathcal{M}, w \Vdash_{\mathrm{MDL}} \neg\varphi & \text{iff} & \mathcal{M}, w \nVdash_{\mathrm{MDL}} \varphi, \\
\mathcal{M}, w \Vdash_{\mathrm{MDL}} \varphi \wedge \psi & \text{iff} & \mathcal{M}, w \Vdash_{\mathrm{MDL}} \varphi \text{ and } \mathcal{M}, w \Vdash_{\mathrm{MDL}} \psi.
\end{array}
$$

Remark 1. MDL obligations are monadic. A naive conditional obligation can be defined via material implication as $O(\varphi|\psi) := \psi \to O\varphi$. This yields detachment, since $O(\varphi|\psi) \wedge \psi \to O\varphi$, but it behaves inadequately when the obligation essentially depends on the condition. In particular, because $\to$ is true whenever its antecedent is false, $O(\varphi|\psi)$ becomes true in all states where ψ does not hold.

3 A Tiny Fragment of MDL

We identify a fragment of MDL reducible to classical propositional logic, where obligations and prohibitions apply only to atomic normative contents rather than arbitrary formulas. This restriction is inspired by the deontic theories of the Sanskrit philosophical school of Mīmāṃsā, where commands are understood as having a single goal [4], and aligns with the deontic extension DELX of ASP [6].

More precisely, we define the *tiny deontic language* $\mathcal{L}_{\mathrm{TDL}}$ as

$$\varphi := \omega \mid \neg\varphi \mid \varphi \wedge \varphi \qquad \text{and} \qquad \omega := p \in \mathit{Prop} \mid O\omega \mid F\omega$$

We call $\omega ::= p \mid O\omega \mid F\omega$ *simple normative atoms*, and denote their set by Ω^t.

Remark 2. Compared to the full language $\mathcal{L}_{\mathrm{MDL}}$, $\mathcal{L}_{\mathrm{TDL}}$ excludes obligations over Boolean combinations (and therefore also over the constants $\top$ or $\bot$). In particular, formulas such as $O(\varphi \wedge \psi)$ are not well formed; instead, we use $O\varphi \wedge O\psi$. Note that in Chellas' semantics, $O(\varphi \wedge \psi)$ implies $O\varphi \wedge O\psi$, but not conversely. Also obligations over disjunctions are likewise not admitted. An additional advantage of this restricted formulation is that it blocks problematic inferences such as Ross' paradox [26], according to which the obligation to mail a letter entails the obligation to mail or burn it.

3.1 Classical Logic as Alternative Semantics

We show that reasoning in $\mathcal{L}_{\mathrm{TDL}}$ reduces to reasoning in classical propositional logic. Intuitively, the only propositionally valid implications between simple normative atoms are identities of the form $\omega \to \omega$ (as $\bot \notin \mathcal{L}_{\mathrm{TDL}}$, the MDL axiom $\neg O\bot$ is not in $\mathcal{L}_{\mathrm{TDL}}$).

For the proof we provide a canonical model ($\mathcal{M}^{\mathrm{TDL}}$), with one world for each propositional valuation over Ω^t.[2]

Definition 1. *$\mathcal{M}^{\mathrm{TDL}}$ is (W^t, N^t, V^t) with $W^t = 2^{\Omega^t}$, and for all $w \in W^t$, $V^t(w) = w \cap Prop$, and $N^t(w) = closure(w^O \cup w^F)$, where*

- $w^O = \{\{w' \in W^t \mid \omega \in w'\} \mid O\omega \in w\}$,
- $w^F = \{\{w' \in W^t \mid \omega \notin w'\} \mid F\omega \in w\}$, *and*
- $closure(X) = \{U \subseteq W^t \mid$ *there exists an* $x \in X$ *such that* $x \subseteq U\}$.

Lemma 2. *$\mathcal{M}^{\mathrm{TDL}}$ is a Chellas-style neighborhood model (by construction).*

Lemma 3. *Let $w \in W^t$ and $\omega \in \Omega^t$. Then $\mathcal{M}^{\mathrm{TDL}}, w \Vdash_{\mathrm{MDL}} \omega$ iff $\omega \in w$.*

Proof (sketch). By structural induction on ω. The case $\omega = p$ follows from $V^t(w) = w \cap Prop$. For $O\omega$, we have to show that $\{w' \mid \omega \in w'\} \in N^t(w)$ iff $O\omega \in w$. The backward direction follows from the definition of w^O. For the forward direction, assume $\{w' \mid \omega \in w'\} \in N^t(w)$. By the definition of $closure$, some $x \in w^O \cup w^F$ is contained in $\{w' \mid \omega \in w'\}$. The world $\emptyset$ excludes $x \in w^F$; $\{\psi\}$ excludes $x = \{w' \mid \psi \in w'\}$ for $\psi \neq \omega$. Thus $x = \{w' \mid \omega \in w'\} \in w^O$. Hence $O\omega \in w$. The case for $F\omega$ is analogous.

Writing $\Vdash_{\mathrm{CL}}$ for classical propositional truth over Ω^t, Lemma 3 extends by structural induction to all Boolean combinations over Ω^t, that is, formulas of $\mathcal{L}_{\mathrm{TDL}}$:

Theorem 4. *Let $w \subseteq \Omega^t$ and $\varphi \in \mathcal{L}_{\mathrm{TDL}}$. Then $\mathcal{M}^{\mathrm{TDL}}, w \Vdash_{\mathrm{MDL}} \varphi$ iff $w \Vdash_{\mathrm{CL}} \varphi$.*

Each finite input involves only finitely many simple normative atoms, so satisfiability and entailment in $\mathcal{L}_{\mathrm{TDL}}$ reduce to classical propositional reasoning over these atoms.

4 Lightweight Deontic Logic

We now extend the previous fragment to a deontic logic that better captures the conditional nature of normative statements: Lightweight Deontic Logic (LDL). Encoding conditional norms via material implication is inadequate (see Remark 1). Instead, LDL assumes *explicit conditional norms* over atomic propositions, specified directly by the system designer. By factual detachment, these conditional norms give rise to unconditional obligations and prohibitions corresponding to the tiny fragment of MDL, yielding the *implicit unconditional norms* of the system. We define $\mathcal{L}_{\mathrm{LDL}}$ by:

$$\varphi ::= \omega \mid \neg\varphi \mid \varphi \wedge \varphi,$$
$$\omega ::= p \in Prop \mid O\omega \mid F\omega \mid \mathbf{O}(\omega|\varphi) \mid \mathbf{F}(\omega|\varphi).$$

We read $\mathbf{O}(\omega|\varphi)$ as "ω is explicitly obligatory given φ" and $\mathbf{F}(\omega|\varphi)$ as "ω is explicitly forbidden given φ". We call a formula of the form $\omega ::= p \mid O\omega \mid F\omega \mid \mathbf{O}(\omega|\varphi) \mid \mathbf{F}(\omega|\varphi)$ a *normative atom* and denote by Ω the set of all such atoms. Note that Ω contains higher-order norms, for example $\mathbf{O}(F\,smoking|F\,drinking)$.

[2] A version with full proofs is available at https://arxiv.org/abs/2610.00668.

4.1 Propositional Semantics

Although $\mathcal{L}_{\text{LDL}}$ generalizes $\mathcal{L}_{\text{TDL}}$ ($\varphi \in \mathcal{L}_{\text{LDL}}$ for each $\varphi \in \mathcal{L}_{\text{TDL}}$), as shown below, it can still be characterized by the semantics of classical logic.

The introduced semantics makes a conceptual distinction between explicit and implicit normative atoms. We say that a given explicit conditional obligation $\mathbf{O}(\omega|\varphi)$ or prohibition $\mathbf{F}(\omega|\varphi)$ is *in effect* if its condition φ is satisfied. This induces the corresponding implicit obligation $O\omega$ or prohibition $F\omega$ to be true (which concerns the actual state rather than hypothetical situations).

Definition 5. *An* LDL model *is a propositional valuation $M \subseteq \Omega$ satisfying for all $\omega \in \Omega$ and $\varphi \in \mathcal{L}_{\text{LDL}}$ the* LDL-consistency *conditions:* [3]

- $\mathbf{O}(\omega|\varphi) \in M \Longrightarrow M \nVdash_{\text{CL}} \varphi$ *or* $O\omega \in M$, *and*
- $\mathbf{F}(\omega|\varphi) \in M \Longrightarrow M \nVdash_{\text{CL}} \varphi$ *or* $F\omega \in M$.

Definition 6. *Let $\varphi \in \mathcal{L}_{\text{LDL}}$ and let M be an LDL model. We write $M \Vdash_{\text{LDL}} \varphi$ if and only if $M \Vdash_{\text{CL}} \varphi$.*

Intuitively, the consistency requirement allows us to infer (unconditional) implicit obligations from (conditional) explicit obligations via factual detachment, without strengthening of the antecedent for explicit obligations, as shown by the example below.

Remark 3. Like in MDL, having a dilemma (i.e. both $O\omega$ and $F\omega$) does not lead to a contradiction.

Example 7. Consider $\varphi = p \wedge \mathbf{O}(r|p)$ and $M = \{p, \mathbf{O}(r|p), Or\}$. Clearly, M is an LDL model and $M \Vdash_{\text{CL}} \varphi \wedge \neg\mathbf{O}(r|p \wedge q)$. Therefore, $M \Vdash_{\text{LDL}} \varphi \wedge \neg\mathbf{O}(r|p \wedge q)$.

Remark 4. LDL is a conservative extension of the tiny fragment of Chellas' MDL.

4.2 Axiomatization of LDL

We define a Hilbert-style proof system for LDL. Its axiom schemata are:

(Prop) Every instance of a classical propositional tautology over Ω.
(O-Det) $\mathbf{O}(\omega|\varphi) \wedge \varphi \rightarrow O\omega$, for all $\omega \in \Omega$ and $\varphi \in \mathcal{L}_{\text{LDL}}$.
(F-Det) $\mathbf{F}(\omega|\varphi) \wedge \varphi \rightarrow F\omega$, for all $\omega \in \Omega$ and $\varphi \in \mathcal{L}_{\text{LDL}}$.

The only inference rule is modus ponens: **(MP)** From φ and $\varphi \rightarrow \psi$, infer ψ.

For $\Gamma \subseteq \mathcal{L}_{\text{LDL}}$ and $\varphi \in \mathcal{L}_{\text{LDL}}$, we write $\Gamma \vdash_{\text{LDL}} \varphi$ iff there exists a finite sequence of formulas $\varphi_1, \ldots, \varphi_n \in \mathcal{L}_{\text{LDL}}$ with $\varphi_n = \varphi$ such that each φ_i is either a member of Γ, an instance of one of the above axiom schemata, or obtained from earlier formulas in the sequence by (MP).

To prove soundness and completeness, we first provide a lemma connecting semantic consequence in LDL-theories to that of classical propositional logic. For $\Gamma \subseteq \mathcal{L}_{\text{LDL}}$ and $\varphi \in \mathcal{L}_{\text{LDL}}$, we write $\Gamma \models_{\text{LDL}} \varphi$ if every LDL model M satisfying $M \Vdash_{\text{LDL}} \Gamma$ also satisfies $M \Vdash_{\text{LDL}} \varphi$. Similarly, we write $\Gamma \models_{\text{CL}} \varphi$ if $M \Vdash_{\text{CL}} \varphi$ for each $M \Vdash_{\text{CL}} \Gamma$. We abbreviate $\emptyset \models_{\text{LDL}} \varphi$ by $\models_{\text{LDL}} \varphi$ and $\emptyset \models_{\text{CL}} \varphi$ by $\models_{\text{CL}} \varphi$.

[3] We can equivalently write M as a Boolean function $\{\omega \mapsto \top \mid \omega \in M\} \cup \{\omega \mapsto \bot \mid \omega \notin M\}$. For $\omega \in \Omega$, we write $M \Vdash_{\text{CL}} \omega$ iff $\omega \in M$, or equivalently iff $M(\omega) = \top$.

Lemma 8. *Let Δ be the set of all instances of $\mathbf{O}(\omega|\varphi)\wedge\varphi \rightarrow O\omega$ and $\mathbf{F}(\omega|\varphi)\wedge\varphi \rightarrow F\omega$. Then for all $\Gamma \subseteq \mathcal{L}_{\mathrm{LDL}}$ and $\varphi \in \mathcal{L}_{\mathrm{LDL}}$, we have*

$$\Gamma \models_{\mathrm{LDL}} \varphi \iff \Gamma \cup \Delta \models_{\mathrm{CL}} \varphi.$$

Proof. We first note that a propositional valuation $M \subseteq \Omega$ is an LDL model iff $M \Vdash_{\mathrm{CL}} \Delta$. Indeed, considering only obligations (the prohibition case is analogous), M violates LDL-consistency exactly when there exists $\mathbf{O}(\omega|\varphi) \in M$ such that $M \Vdash_{\mathrm{CL}} \varphi$, $O\omega \notin M$, and thus $M \nVdash_{\mathrm{CL}} (\mathbf{O}(\omega|\varphi)\wedge\varphi \rightarrow O\omega)$. Hence, LDL-inconsistency corresponds precisely to falsifying some formula in Δ, and vice versa. It follows that the models of $\Gamma \cup \Delta$ are exactly the LDL models of Γ. Therefore, by Definition 6, $\Gamma \models_{\mathrm{LDL}} \varphi \iff \Gamma\cup\Delta \models_{\mathrm{CL}} \varphi$.

Theorem 9 (Soundness and Completeness). *Let $\Gamma, \varphi \subseteq \mathcal{L}_{\mathrm{LDL}}$. Then*

$$\Gamma \vdash_{\mathrm{LDL}} \varphi \iff \Gamma \models_{\mathrm{LDL}} \varphi.$$

Proof. Let Δ be defined as in Lemma 8. $\Gamma \models_{\mathrm{LDL}} \varphi$ iff $\Gamma \cup \Delta \models_{\mathrm{CL}} \varphi$. By soundness and completeness of classical propositional logic, $\Gamma \cup \Delta \models_{\mathrm{CL}} \varphi$ iff $\Gamma \cup \Delta \vdash_{\mathrm{CL}} \varphi$ (where $\vdash_{\mathrm{CL}}$ is derivability in classical propositional logic over Ω). Finally, note that the Hilbert system is classical propositional logic together with all formulas in Δ as additional axiom instances. Therefore, $\Gamma \cup \Delta \vdash_{\mathrm{CL}} \varphi \iff \Gamma \vdash_{\mathrm{LDL}} \varphi$. Thus, $\Gamma \models_{\mathrm{LDL}} \varphi \iff \Gamma \vdash_{\mathrm{LDL}} \varphi$.

Lemma 8 yields a polynomial reduction from LDL-satisfiability and LDL-validity to classical propositional satisfiability and validity over the normative atoms occurring in the input formula and their detachment consequents. For $\varphi \in \mathcal{L}_{\mathrm{LDL}}$, let Δ^{φ} contain the detachment implications corresponding to all explicit normative atoms occurring recursively in φ. Then φ is LDL-satisfiable iff $\varphi \wedge \bigwedge \Delta^{\varphi}$ is classically satisfiable, and LDL-valid iff $(\bigwedge \Delta^{\varphi}) \rightarrow \varphi$ is classically valid. Since hardness already holds for the propositional fragment, we obtain:

Theorem 10. *Satisfiability in LDL is NP-complete, and validity is co-NP-complete.*

5 A Doxastic Extension of LDL

We now define a direct combination of LDL with multi-agent doxastic logic KD45 [16, 15]. This allows us to express norms conditioned on what agents believe and what they believe to be possible, as well as to express beliefs about norms.

Let Ag be a finite set of agents. The language $\mathcal{L}_{\mathrm{LDLB}}$ is defined by the following grammar, where $p \in Prop$ and $i \in Ag$:

$$\varphi ::= \omega \mid B_i\varphi \mid \neg\varphi \mid \varphi \wedge \varphi,$$
$$\omega ::= p \mid O\omega \mid F\omega \mid \mathbf{O}(\omega|\varphi) \mid \mathbf{F}(\omega|\varphi).$$

We read formulas $B_i\varphi$ as *"agent i believes φ"* and define $\hat{B}_i\varphi$ as an abbreviation for $\neg B_i\neg\varphi$ (which can be read "agent i believes it possible that φ"). We further call formulas generated by ω *objective atoms* and write Ω^B for the set of all such atoms. An important difference from LDL is that conditions of explicit obligations and prohibitions may now contain belief operators. The normative contents of obligations and prohibitions, however, are still restricted to formulas generated by ω (in particular, we cannot model the obligation or prohibition to have a certain belief). Moreover, in $\mathcal{L}_{\mathrm{LDLB}}$, we can also specify beliefs about normative facts, such as $B_iO\omega$, $B_iF\omega$, or $B_i\mathbf{O}(\omega|\varphi)$.

5.1 Kripke Semantics over LDL-Valuations

The LDLB semantics combines doxastic Kripke models with LDL-consistency.

Definition 11. *An* LDLB premodel *is a tuple*

$$\mathcal{M} = (W, (R_i)_{i \in Ag}, V)$$

where W is a non-empty set of worlds, $R_i \subseteq W \times W$ is agent i's doxastic accessibility relation*, and $V : W \to 2^{\Omega^B}$ assigns an LDL-style valuation to each world. For $w \in W$, write $R_i(w) = \{v \in W \mid wR_iv\}$. Truth in a premodel is defined as follows:*

$$\begin{aligned}
&\mathcal{M}, w \Vdash_{\text{LDLB}} \omega && \text{iff } \omega \in V(w),\\
&\mathcal{M}, w \Vdash_{\text{LDLB}} B_i\varphi && \text{iff } \mathcal{M}, v \Vdash_{\text{LDLB}} \varphi \text{ for all } v \text{ with } wR_iv,\\
&\mathcal{M}, w \Vdash_{\text{LDLB}} \neg\varphi && \text{iff } \mathcal{M}, w \nVdash_{\text{LDLB}} \varphi,\\
&\mathcal{M}, w \Vdash_{\text{LDLB}} \varphi \wedge \psi && \text{iff } \mathcal{M}, w \Vdash_{\text{LDLB}} \varphi \text{ and } \mathcal{M}, w \Vdash_{\text{LDLB}} \psi.
\end{aligned}$$

A premodel is an LDLB *model iff it satisfies the following two conditions.*

(i) For every $i \in Ag$, R_i is serial*,* transitive*, and* euclidean*.*
(ii) For every $w \in W$, every $\omega \in \Omega^B$, and every $\varphi \in \mathcal{L}_{\text{LDLB}}$,

$$\begin{aligned}
\mathbf{O}(\omega|\varphi) \in V(w) \text{ and } \mathcal{M}, w \Vdash_{\text{LDLB}} \varphi \quad &\Longrightarrow \quad O\omega \in V(w),\\
\mathbf{F}(\omega|\varphi) \in V(w) \text{ and } \mathcal{M}, w \Vdash_{\text{LDLB}} \varphi \quad &\Longrightarrow \quad F\omega \in V(w).
\end{aligned}$$

The first condition gives the doxastic logic KD45 for each agent: beliefs are consistent; that is, $B_i\varphi \to \hat{B}_i\varphi$ is valid. Furthermore, agents have positive and negative introspection: $B_i\varphi \to B_iB_i\varphi$ and $\neg B_i\varphi \to B_i\neg B_i\varphi$.[4]

The second condition enforces *LDL-consistency* at each individual world. This effectively lifts the factual detachment of LDL to the mental states of the agents: If an agent believes (or believes it possible) that at the same time an explicit norm is given and its condition is satisfied, then the agent believes (or believes it possible) that the norm is in effect.

For $\Gamma \subseteq \mathcal{L}_{\text{LDLB}}$ and $\varphi \in \mathcal{L}_{\text{LDLB}}$, we write $\Gamma \models_{\text{LDLB}} \varphi$ iff every world of every LDLB model satisfying all formulas in Γ also satisfies φ.

Example 12. Let r mean that an autonomous car moves to the right to yield (e.g., if an emergency vehicle is approaching), and c that a cyclist is to the right of the car. The language distinguishes $\mathbf{F}(r|B_ac)$ from $\mathbf{F}(r|\hat{B}_ac)$; only the latter captures "if a cyclist cannot be ruled out, moving right is forbidden", that is,

$$\{\mathbf{F}(r|\hat{B}_ac), \hat{B}_ac\} \models_{\text{LDLB}} Fr \quad \text{while} \quad \{\mathbf{F}(r|B_ac), \hat{B}_ac\} \not\models_{\text{LDLB}} Fr.$$

We can also reason about an agent's own perspective (i.e., a believes both the norm $\mathbf{F}(r|\hat{B}_ac)$ and that c is possible): $\{B_a\mathbf{F}(r|\hat{B}_ac), \hat{B}_ac\} \models_{\text{LDLB}} B_aFr$ (via $B_a\hat{B}_ac$ by introspection), but $\{\mathbf{F}(r|\hat{B}_ac), \hat{B}_ac\} \not\models_{\text{LDLB}} B_aFr$.

[4] Positive and negative introspection are controversial assumptions. In the case of artificial agents, however, they can be justified under the assumption that each agent has a consistent belief base, full access, and the ability to reason perfectly about it.

Proposition 13. LDLB *is axiomatized by adding to multi-agent KD45 the following LDL detachment axioms:*

(O-Det) $\mathbf{O}(\omega|\varphi) \wedge \varphi \rightarrow O\omega$*, for all* $\omega \in \Omega^B$ *and* $\varphi \in \mathcal{L}_{\text{LDLB}}$.
(F-Det) $\mathbf{F}(\omega|\varphi) \wedge \varphi \rightarrow F\omega$*, for all* $\omega \in \Omega^B$ *and* $\varphi \in \mathcal{L}_{\text{LDLB}}$.

Proof (sketch). Soundness follows from the KD45 frame conditions in (i) and the detachment condition (ii). For completeness, we define a canonical model of maximal LDLB-consistent sets, with wR_iv iff $\{\psi \mid B_i\psi \in w\} \subseteq v$, and $V(w) = w \cap \Omega^B$. Seriality, transitivity, and Euclideanness follow from axioms (D), (4), and (5), respectively. The truth lemma follows by structural induction. It remains to check condition (ii). Suppose $\mathbf{O}(\omega|\varphi) \in V(w)$ and φ holds at w. By the truth lemma, $\varphi \in w$, so (O-Det) gives $O\omega \in w$ and hence $O\omega \in V(w)$. The prohibition case is analogous. Thus the canonical model is an LDLB model, and completeness follows by the truth lemma.

6 Optimizing Norm Violations

We now turn to decision making itself. An acting agent $a \in Ag$ must choose among several decisions, each of which may violate some of the applicable norms, each having a possibly different weight. We formalize this as the Doxastic Norm Compliance Optimization Problem (DNCOP): given a finite description of the facts, the agents' beliefs, and a set of weighted norms, find a decision that minimizes total violation weight.

Two perspectives are possible. The objective perspective evaluates violations against the true state of the world; the subjective perspective evaluates them against the acting agent's beliefs. We provide conditions under which these two agree and under which optimal decisions can be found by a polynomial reduction to weighted partial MaxSAT [2].

6.1 State Vocabulary

A DNCOP is described using a fixed, finite set of atoms – the *state vocabulary*. Besides plain facts, we need a compact way to talk about nested beliefs. Following restricted modal literals and proper epistemic knowledge bases [18, 23], beliefs are represented only by nested belief operators applied to a single objective atom.

As agents' beliefs are introspective, a nested belief such as $B_1B_1B_2\neg B_2B_2B_1\omega$ is LDLB-equivalent to $B_1\neg B_2B_1\omega$: repeated operators for the same agent collapse. We therefore only need belief chains that alternate agents.

Definition 14. *A* restricted doxastic-deontic atom (RDDA) *is either an objective atom* $\omega \in \Omega^B$*, or a formula* $X_1 \dots X_n\omega$ *(*$n \geq 1$*), where each* X_j *is either* B_{a_j} *or* $B_{a_j}\neg$*, consecutive operators refer to different agents (*$a_j \neq a_{j+1}$*), and* ω *is an objective atom. We denote the set of all RDDAs by* RDDA.

A finite set Σ *of such atoms is a* state vocabulary *if it is closed under sub-atoms, i.e., (let* $\mathsf{Bool}(\Sigma)$ *denote all Boolean combinations over members of* Σ*)*

(i) if $B_i\alpha \in \Sigma$ *or* $B_i\neg\alpha \in \Sigma$ *(for any RDDA* α*), then* $\alpha \in \Sigma$*,*
(ii) if $O\omega \in \Sigma$*,* $F\omega \in \Sigma$*,* $\mathbf{O}(\omega|\varphi) \in \Sigma$*, or* $\mathbf{F}(\omega|\varphi) \in \Sigma$*, then* $\omega \in \Sigma$*; and*

(iii) if $\mathbf{O}(\omega|\varphi) \in \Sigma$ *or* $\mathbf{F}(\omega|\varphi) \in \Sigma$*, then* $\varphi \in \mathsf{Bool}(\Sigma)$.

Finally, we denote by Σ^a *the subset of* Σ *consisting of RDDAs that have* B_a *as their outermost operator (i.e., the beliefs of the acting agent).*

A state vocabulary therefore contains the atoms needed to finitely describe a norm compliance optimization problem. Later, these atoms will be classified as fixed by the environment or controlled by the acting agent.

6.2 Doxastic Norm Compliance Optimization Problems (DNCOPs)

We formulate the optimization problem faced by a fixed acting agent $a \in Ag$. The goal is to choose a decision that minimizes (weighted) norm violations.

The state vocabulary Σ is partitioned into atoms fixed true by the environment $\mathcal{T}$ (e.g., whether an emergency vehicle is approaching), fixed false $\mathcal{F}$, and $\overline{\mathcal{D}}$ atoms under the agent's control (its decisions and their knowledge[5]).

Definition 15. *A* DNCOP *is a tuple* $\Delta = (\Sigma, \mathcal{T}, \mathcal{F}, \mathcal{D}, \chi, N, v)$*, where*

- Σ *is a finite state vocabulary,*
- $\mathcal{T} \subseteq \Sigma$ *(resp.* $\mathcal{F} \subseteq \Sigma$*) contains the atoms fixed true (resp. false) independently of the decision the agent makes,*
- $\mathcal{D} \subseteq \Omega^B$ *is the set of* decision atoms *the acting agent sets directly, and we write* $\overline{\mathcal{D}} = \mathcal{D} \cup \{B_a\delta, B_a\neg\delta \mid \delta \in \mathcal{D}\}$ *for these atoms together with the agent's (automatically correct) beliefs about them;*
- $\chi \in \mathsf{Bool}(\mathcal{D})$ *is the* decision constraint*, defining the possible decisions* $\mathsf{Dec}(\Delta) = \{d \subseteq \mathcal{D} \mid d \Vdash_{\mathrm{CL}} \chi\}$
- $N \subseteq \Sigma$ *is the set of* relevant norms *(of the form* $\mathbf{O}(\omega|\varphi)$*,* $\mathbf{F}(\omega|\varphi)$*,* $O\omega$*,* $F\omega$*).*
- v *is a function that assigns a positive integer weight to each norm* $n \in N$*.*

We require that $\Sigma = \mathcal{T} \dot\cup \mathcal{F} \dot\cup \overline{\mathcal{D}}$ *is a disjoint union.*

Unless specified otherwise, we assume $\chi = \top$, so every subset of $\mathcal{D}$ is a possible decision. We also omit Σ when clear from the context.

Example 16. Let r and c be as in Example 12, and let e denote that an emergency vehicle is approaching from behind. Consider the norms $\mathbf{O}(r|e)$ ("move right if an emergency vehicle approaches") and $\mathbf{F}(r|c)$ ("don't move right", once a cyclist is present). Fix e and the agent's correct belief in it; leave open whether c is fixed true or false. Assume $v(\mathbf{O}(r|e)) = 1$ and $v(\mathbf{F}(r|c)) = 2$. We define two variants of the scenario, one where no cyclist is present (c is false and fixed) and one where a cyclist is present (c is true and fixed). We define the DNCOP where there is a cyclist as Δ_c (with $\mathcal{T}_c$, $\mathcal{F}_c$), and the DNCOP without a cyclist as $\Delta_{\neg c}$ (with $\mathcal{T}_{\neg c}$, $\mathcal{F}_{\neg c}$). For both versions $\mathcal{D} = \{r\}$, and $N = \{\mathbf{O}(r|e), \mathbf{F}(r|c)\}$.

$$\mathcal{T}_{\neg c} = \{e, \mathbf{O}(r|e), \mathbf{F}(r|c), Or, B_a e, B_a\mathbf{O}(r|e), B_a\mathbf{F}(r|c)\},\ \mathcal{F}_{\neg c} = \{c, Fr, B_a c, B_a\neg c\},$$
$$\mathcal{T}_c = \{e, c, \mathbf{O}(r|e), \mathbf{F}(r|c), Or, Fr, B_a e, B_a\mathbf{O}(r|e), B_a\mathbf{F}(r|c)\},\ \mathcal{F}_c = \{B_a c, B_a\neg c\}.$$

[5] Since LDLB does not have knowledge, we use true beliefs.

In either case, we assume that a believes an emergency vehicle is approaching from behind and that both $\mathbf{O}(r|e)$ and $\mathbf{F}(r|c)$ hold, and therefore that Or holds. In Δ_c also Fr holds. All these beliefs are *true beliefs*.

By the default $\chi = \top$, the possible decisions are $\mathsf{Dec}(\Delta_{\neg c}) = \mathsf{Dec}(\Delta_c) = \{\emptyset, \{r\}\}$, where $\emptyset$ represents not moving right and $\{r\}$ represents moving right.

6.3 Compatible Models and Violation Counts

Given a decision d in a DNCOP Δ, we can ask which models of the LDLB logic are consistent with it – either assuming full, correct information (*objective compatibility*) or only what the acting agent believes (*subjective compatibility*).

Intuitively, compatible models represent the possible situations against which a decision is evaluated. The objective perspective fixes all information specified by the DNCOP, whereas the subjective perspective fixes only the information available through the acting agent's beliefs. Consequently, several objectively different situations may remain possible from the subjective perspective.

Definition 17. *Let $\Delta = (\Sigma, \mathcal{T}, \mathcal{F}, \mathcal{D}, \chi, N, v)$ be a DNCOP and $d \subseteq \mathcal{D}$. Let*

- $\mathsf{Choice}(\Delta, d) = \bigcup_{\delta \in d}\{\delta, B_a\delta, \neg B_a\neg\delta\} \cup \bigcup_{\delta \in \mathcal{D}\setminus d}\{\neg\delta, \neg B_a\delta, B_a\neg\delta\}$,
  (adding B_a because the agent believes their decision and thus also $\hat{B}_a$).
- $\mathsf{Obj}(\Delta, d) = \mathcal{T} \cup \{\neg\alpha \mid \alpha \in \mathcal{F}\} \cup \mathsf{Choice}(\Delta, d)$
- $\mathsf{Subj}(\Delta, d) = (\mathcal{T} \cap \Sigma^a) \cup \{\neg\alpha \mid \alpha \in \mathcal{F} \cap \Sigma^a\} \cup \mathsf{Choice}(\Delta, d)$
  i.e., $\{$only the acting agent's own beliefs about $\mathcal{T}$ and $\mathcal{F}\} \cup \mathsf{Choice}(\Delta, d)$.

We characterize the LDLB-models $(\mathcal{M}, w)$ that are subjectively and objectively compatible with a DNCOP Δ and a decision d. For the subjective perspective, we assume that the acting agent trusts their own beliefs. To achieve this, we restrict ourselves to models with wR_aw for the actual world w.[6]

Definition 18. *A pointed* LDLB *model $(\mathcal{M}, w)$ is* objectively compatible *with (Δ, d) if $\mathcal{M}, w \Vdash_{\text{LDLB}} \mathsf{Obj}(\Delta, d)$. It is* subjectively compatible *with (Δ, d) if $\mathcal{M}, w \Vdash_{\text{LDLB}} \mathsf{Subj}(\Delta, d)$ and the agent trusts itself at w, i.e., wR_aw.*

We denote the corresponding sets of models by $\mathsf{Mod}^{\text{obj}}(\Delta, d)$ and $\mathsf{Mod}^{\text{sub}}(\Delta, d)$. Nonemptiness of the respective set defines objective or subjective admissibility.

In some DNCOPs none of the decisions are objectively admissible (due to inconsistencies in beliefs, norms, or both). In other DNCOPs, only specific decisions are. Consider an example from the perspective of the designer, where $p \in \mathcal{T}$ is fixed true and $Ox, \mathbf{O}(x|p) \in \mathcal{D}$ are decision atoms. By detachment, making $\mathbf{O}(x|p)$ true requires making Ox true as well. In this example $\{\mathbf{O}(x|p), Ox\}$ is objectively admissible but $\{\mathbf{O}(x|p)\}$ is not. Since facts in $\mathcal{T}$ are inaccessible to the acting agent, the objective admissibility of $\{\mathbf{O}(x|p)\}$ cannot be determined subjectively (and in particular is not implied by subjective admissibility).

Hence, we restrict our attention to *decision-sound* DNCOPs. This strong but necessary promise allows optimization from the agent's perspective.

[6] This ensures that $(\mathcal{M}, w) \Vdash_{\text{LDLB}} \varphi$ whenever $(\mathcal{M}, w) \Vdash_{\text{LDLB}} B_a\varphi$.

Definition 19. *We call a DNCOP* Δ decision-sound *if* $\mathsf{Dec}(\Delta) \neq \emptyset$ *and every* $d \in \mathsf{Dec}(\Delta)$ *is objectively admissible.*

Importantly, as shown below, every objectively admissible decision is also subjectively admissible. This justifies the use of subjectively compatible models when optimizing from the subjective perspective.

Proposition 20. *Let* $\Delta = (\Sigma, \mathcal{T}, \mathcal{F}, \mathcal{D}, \chi, N, v)$ *and* $d \subseteq \mathcal{D}$*. If* d *is objectively admissible, then it is subjectively admissible.*

Proof. Choose $(\mathcal{M}, w) \in \mathsf{Mod}^{\mathrm{obj}}(\Delta, d)$. By seriality, $wR_a v$ for some v. Transitivity and euclideanness give $R_a(w) = R_a(v)$ and $vR_a v$. Hence every RDDA whose outermost agent is a has the same truth value at w and v, so the non-decision part of $\mathsf{Subj}(\Delta, d)$ holds at v. As w and v have the same a-successors, all belief literals in $\mathsf{Choice}(\Delta, d)$ that hold at w also hold at v. Moreover, $B_a\delta$ or $B_a\neg\delta$ ensures that the corresponding decision literal δ or $\neg\delta$ holds at v. Hence $\mathcal{M}, v \Vdash_{\mathrm{LDLB}} \mathsf{Subj}(\Delta, d)$. Together with $vR_a v$, this gives $(\mathcal{M}, v) \in \mathsf{Mod}^{\mathrm{sub}}(\Delta, d)$. Therefore $\mathsf{Mod}^{\mathrm{sub}}(\Delta, d) \neq \emptyset$, so d is subjectively admissible.

We show how to count and optimize norm violations for DNCOPs. We first formalize a violation condition χ_n for norms: An explicit or implicit obligation is violated when it exists, its content is false, and (only in the case of explicit obligations) its condition holds. Prohibitions are treated analogously.

We use this to define the weighted norm violation count $\mathsf{viol}(\mathcal{M}, w)$ for a given model $(\mathcal{M}, w)$, and subsequently the number of violations $\mathsf{viol}^{\mathrm{obj}}(\Delta, d)$ and $\mathsf{viol}^{\mathrm{sub}}(\Delta, d)$ for admissible decisions d in a DNCOP Δ, as predicted by the objective or subjective perspective. Since there may be multiple compatible models, our definitions use a pessimistic worst-case[7] assumption.

From the subjective perspective, the available information may leave several situations possible, with different norm violations. We therefore evaluate a decision conservatively, by the largest violation cost among all compatible situations. We then choose a decision that minimizes this worst-case cost.

Definition 21. *Let* $\Delta = (\Sigma, \mathcal{T}, \mathcal{F}, \mathcal{D}, \chi, N, v)$ *be a DNCOP. For each norm* $n \in N$ *we define the violation condition as*

$$\chi_n = \begin{cases} \mathbf{O}(\omega|\varphi) \wedge \varphi \wedge \neg\omega, & n = \mathbf{O}(\omega|\varphi), \\ \mathbf{F}(\omega|\varphi) \wedge \varphi \wedge \omega, & n = \mathbf{F}(\omega|\varphi), \\ O\omega \wedge \neg\omega, & n = O\omega, \\ F\omega \wedge \omega, & n = F\omega. \end{cases}$$

For a pointed model $(\mathcal{M}, w)$*, we define its weighted norm violation count as*

$$\mathsf{viol}(\mathcal{M}, w) = \sum_{\substack{n \in N \\ \mathcal{M}, w \Vdash_{\mathrm{LDLB}} \chi_n}} v(n).$$

[7] Assuming the best case is unrealistic and assuming the average is unjustified since compatible models need not be equally likely.

If $d \subseteq \mathcal{D}$ is objectively admissible, we define its objective violation count as

$$\mathsf{viol}^{\mathrm{obj}}(\Delta, d) = \max_{(\mathcal{M},w)\in\mathsf{Mod}^{\mathrm{obj}}(\Delta,d)} \mathsf{viol}(\mathcal{M}, w).$$

If $d \subseteq \mathcal{D}$ is subjectively admissible, we define its subjective violation count as

$$\mathsf{viol}^{\mathrm{sub}}(\Delta, d) = \max_{(\mathcal{M},w)\in\mathsf{Mod}^{\mathrm{sub}}(\Delta,d)} \mathsf{viol}(\mathcal{M}, w).$$

Both maxima exist because, by admissibility, the respective sets of compatible models are non-empty and every violation count belongs to $\{0, \ldots, \sum_{n\in N} v(n)\}$.

Remark 5. The use of weights allows our framework to handle the priority of more specific norms over more general ones (*lex specialis*), and contrary-to-duty (CTD) obligations. Priority between norms can be encoded by assigning higher weights to more specific norms, thereby incentivizing the agent to comply with them whenever they apply. For CTD obligations, independently of the concrete positive weights assigned to the primary obligation (whose violation triggers the secondary obligation) and to the secondary obligation, violating both always incurs a higher cost than violating only the primary obligation. When multiple interacting norms are present, however, appropriate weights must be chosen to reflect the desired priorities, in line with the utilitarian approach in [24] to design norm-compliant reinforcement learning agents.

By minimizing over all possible decisions, we now define the objectively and subjectively optimal number of violations $\mathsf{viol}^{\mathrm{obj}}_{\mathrm{min}}(\Delta)$ and $\mathsf{viol}^{\mathrm{sub}}_{\mathrm{min}}(\Delta)$, as well as the objectively and subjectively optimal decisions $\mathsf{optdec}^{\mathrm{obj}}(\Delta)$ and $\mathsf{optdec}^{\mathrm{sub}}(\Delta)$. This is well-defined for decision-sound DNCOPs since they have at least one possible decision and all decisions are admissible.

Definition 22. *Let Δ be a decision-sound DNCOP. We define its objectively and subjectively optimal norm violation counts as*

$$\mathsf{viol}^{\mathrm{obj}}_{\mathrm{min}}(\Delta) = \min_{d\in\mathsf{Dec}(\Delta)} \mathsf{viol}^{\mathrm{obj}}(\Delta, d), \quad \textit{and} \quad \mathsf{viol}^{\mathrm{sub}}_{\mathrm{min}}(\Delta) = \min_{d\in\mathsf{Dec}(\Delta)} \mathsf{viol}^{\mathrm{sub}}(\Delta, d).$$

The corresponding sets of optimal decisions are defined as

$$\mathsf{optdec}^{\mathrm{obj}}(\Delta) = \{d \in \mathsf{Dec}(\Delta) \mid \mathsf{viol}^{\mathrm{obj}}(\Delta, d) = \mathsf{viol}^{\mathrm{obj}}_{\mathrm{min}}(\Delta)\},$$
$$\mathsf{optdec}^{\mathrm{sub}}(\Delta) = \{d \in \mathsf{Dec}(\Delta) \mid \mathsf{viol}^{\mathrm{sub}}(\Delta, d) = \mathsf{viol}^{\mathrm{sub}}_{\mathrm{min}}(\Delta)\}.$$

For a DNCOP $\Delta = (\Sigma, \mathcal{T}, \mathcal{F}, \mathcal{D}, \chi, N, v)$ and every $d \subseteq \mathcal{D}$, the objective description determines a propositional valuation $\nu_{\Delta,d} = \mathsf{Obj}(\Delta, d) \cap \Sigma$, since it contains exactly one of α and $\neg\alpha$ for every $\alpha \in \Sigma$. All objectively compatible models assign the same truth values to the atoms in Σ. Since each χ_n is a Boolean combination of these atoms, it can be evaluated under $\nu_{\Delta,d}$. Thus, whenever d is objectively admissible, the maximization in the definition of $\mathsf{viol}^{\mathrm{obj}}(\Delta, d)$ is vacuous and the violation count can be computed directly from $\nu_{\Delta,d}$.

Proposition 23. *Let* $\Delta = (\Sigma, \mathcal{T}, \mathcal{F}, \mathcal{D}, \chi, N, v)$ *be a DNCOP and let* $d \subseteq \mathcal{D}$ *be objectively admissible. For every pointed model* $(\mathcal{M}, w) \in \mathsf{Mod}^{\mathrm{obj}}(\Delta, d)$*, we have* $\mathsf{viol}(\mathcal{M}, w) = \sum_{\substack{n \in N \\ \nu_{\Delta,d} \Vdash_{\mathrm{CL}} \chi_n}} v(n)$*. Consequently,* $\mathsf{viol}^{\mathrm{obj}}(\Delta, d) = \sum_{\substack{n \in N \\ \nu_{\Delta,d} \Vdash_{\mathrm{CL}} \chi_n}} v(n)$*.*

The maximization over compatible models is still relevant from the subjective perspective, since atoms relevant to the computation of violation counts may be open under the beliefs of the acting agent.

Example 24. Consider the DNCOPs $\Delta_{\neg c}$ and Δ_c from Ex. 16. Both DNCOPs are decision-sound and in both variants, the agent believes e and $\mathbf{O}(r|e)$ – so they always believe the obligation is in force. Since c is left open in the agent's beliefs, moving right risks the cost-2 violation in the worst case, while not moving right always costs 1 (violating $\mathbf{O}(r|e)$); not moving right is therefore subjectively optimal regardless of the actual truth of c. Objectively, the optimal decision matches the true value of c: if no cyclist is present, moving right is objectively optimal; if one is present, not moving right is objectively optimal, as it is subjectively.

6.4 When the Two Perspectives Coincide

The mismatch described in the above example (subjective and objective optimum disagreeing) arises because the norms' conditions (e and c) are objective facts the agent may be unsure about. If instead every norm and every norm's condition is stated directly in terms of the agent's own beliefs, its truth value is automatically settled by the agent's beliefs at every subjectively compatible model, making the subjective violation cost determinate. If these beliefs are also truthful, the two perspectives coincide.

Definition 25. *A DNCOP* $\Delta = (\Sigma, \mathcal{T}, \mathcal{F}, \mathcal{D}, \chi, N, v)$ *is* local *if every norm* $n \in N$*, writing* ω *for its content and* φ *for its condition (in the case of explicit norms), satisfies the following conditions:*

(i) either $n \in \mathcal{D}$*, or exactly one of* $B_a n$ *and* $B_a \neg n$ *belongs to* $\mathcal{T}$*;*
(ii) either $\omega \in \mathcal{D}$*, or exactly one of* $B_a \omega$ *and* $B_a \neg \omega$ *belongs to* $\mathcal{T}$*; and*
(iii) $\varphi \in \mathsf{Bool}(\Sigma^a)$ *(if* n *is an explicit norm).*

A DNCOP is truthful *if the acting agent's beliefs about fixed atoms are correct. Formally, for every* $\alpha \in \mathcal{T} \cup \mathcal{F}$*, we require* $B_a \alpha \in \mathcal{T} \Longrightarrow \alpha \in \mathcal{T}$ *and* $B_a \neg \alpha \in \mathcal{T} \Longrightarrow \alpha \in \mathcal{F}$*.*

Locality reflects a common design pattern for actionable norms: rather than conditioning obligations on facts the agent may only partially observe, the norm is phrased directly in terms of the agent's own (introspectively accessible) beliefs. The next example restates the norms of Example 24 in this style, and we will see that this removes the discrepancy between the two examples. Truthfulness is a complementary idealization, assuming the agent's beliefs are accurate.

Example 26. Replace in Ex. 16 the norms with $n_O = \mathbf{O}(r|B_a e \wedge \neg \hat{B}_a c)$ ("move right if you believe an emergency vehicle is coming and do not consider a cyclist possible")

and $n_F = \mathbf{F}(r|\hat{B}_a c)$ ("do not move right if you cannot rule out a cyclist") with $v(n_O) = v(n_F) = 1$. Hence $N = \{n_O, n_F\}$, $\mathcal{D} = \{r\}$, and

$$\mathcal{T} = \{e, n_O, n_F, B_a e, B_a n_O, B_a n_F\}, \quad \mathcal{F} = \{c, B_a c, B_a \neg c\}.$$

The resulting DNCOP is decision-sound. Objectively, an emergency vehicle is approaching from behind and no cyclist is present to the right of the car. The agent believes that the emergency vehicle is approaching but has no belief about a cyclist. Since $B_a \neg c \in \mathcal{F}$, the formula $\hat{B}_a c$ is true: moving right has objective and subjective cost one, whereas not moving right has cost zero. Thus, not moving right is uniquely optimal from both perspectives. Both norms condition only on the agent's own beliefs, so the DNCOP is local; since the agent's beliefs are in fact accurate, it is truthful too. The two conditions are mutually exclusive, so no priority between the norms is ever needed.

We show below that in a local, decision-sound DNCOP Δ, all subjectively compatible models for a given decision agree on the violation count: the worst case in $\mathsf{viol}^{\mathrm{sub}}(\Delta, d)$ is not actually in doubt.

Proposition 27. *Let Δ be a decision-sound local DNCOP. For every $d \in \mathsf{Dec}(\Delta)$ and all $(\mathcal{M}, w), (\mathcal{M}', w') \in \mathsf{Mod}^{\mathrm{sub}}(\Delta, d)$, $\mathsf{viol}(\mathcal{M}, w) = \mathsf{viol}(\mathcal{M}', w')$. Hence the maximization defining $\mathsf{viol}^{\mathrm{sub}}(\Delta, d)$ is vacuous.*

Proof. For each $n \in N$, the choice or, by reflexivity, the corresponding belief fixes both n and its content in every subjectively compatible model. The subjective fixed and choice descriptions fix all RDDAs in its local condition. Hence all such models agree on whether each n is violated, and therefore on the total count.

If the DNCOP is additionally truthful, the subjective and objective violation counts coincide for every decision, and hence so do the optimal decisions and their violation counts under both perspectives.

Proposition 28. *Let Δ be a decision-sound, local, and truthful DNCOP. For every $d \in \mathsf{Dec}(\Delta)$, we have $\mathsf{viol}^{\mathrm{sub}}(\Delta, d) = \mathsf{viol}^{\mathrm{obj}}(\Delta, d)$. Consequently, $\mathsf{viol}^{\mathrm{sub}}_{\min}(\Delta) = \mathsf{viol}^{\mathrm{obj}}_{\min}(\Delta)$ and $\mathsf{optdec}^{\mathrm{sub}}(\Delta) = \mathsf{optdec}^{\mathrm{obj}}(\Delta)$.*

Proof. For decision atoms, the objective and subjective descriptions agree by construction. For the other atoms, locality provides a positive or negative belief, and truthfulness makes a's beliefs agree with reality. Both descriptions also agree on every RDDA in norm conditions. Thus every subjectively compatible model evaluates each χ_n according to $\nu_{\Delta,d}$. Prop. 23 gives equality of the costs, and both optimizations range over $\mathsf{Dec}(\Delta)$.

6.5 Reduction to MaxSAT

For local DNCOPs, every subjective violation is a Boolean function of the decision. We can therefore optimize subjective violations directly with MaxSAT. Under truthfulness, the same construction also optimizes objective violations.

Definition 29. *Let $\Delta = (\Sigma, \mathcal{T}, \mathcal{F}, \mathcal{D}, \chi, N, v)$ be a local DNCOP. We directly use the decision atoms $\delta \in \mathcal{D}$ as Boolean variables for the MaxSAT problem. Apart from these decision atoms, the translation uses only the acting agent's belief part of the DNCOP, given by $\mathcal{T} \cap \Sigma^a$ and $\mathcal{F} \cap \Sigma^a$.*

We first translate the decision constraint χ polynomially to a CNF using a Tseitin transformation [28]. The resulting clauses are used as hard constraints.

For each $n \in N$, we define a Boolean formula $\widehat{\chi}_n$ over $\mathcal{D}$ characterizing the condition χ_n under which n is violated. We obtain this formula by replacing each RDDA occurring directly in χ_n, i.e., each RDDA occurring as a propositional atom of χ_n, as follows:

(i) any δ or $B_a\delta$ with $\delta \in \mathcal{D}$ is replaced by δ,
(ii) any $B_a\neg\delta$ with $\delta \in \mathcal{D}$ is replaced by $\neg\delta$,
(iii) any remaining objective atom $\alpha \in \Omega^B$ (such as n itself and its content ω) is replaced by $\top$ if $B_a\alpha \in \mathcal{T} \cap \Sigma^a$, and by $\bot$ if $B_a\neg\alpha \in \mathcal{T} \cap \Sigma^a$, and
(iv) any remaining $\alpha \in \Sigma^a$ is replaced by $\top$ if $\alpha \in \mathcal{T} \cap \Sigma^a$, and by $\bot$ if $\alpha \in \mathcal{F} \cap \Sigma^a$.

Locality guarantees that every RDDA occurring in χ_n is covered by exactly one of these cases. For each norm $n \in N$, we then introduce a fresh variable z_n and impose $z_n \leftrightarrow \widehat{\chi}_n$ as a hard constraint, adding $\neg z_n$ as a soft clause of weight $v(n)$. Thus, weight $v(n)$ is incurred exactly when n is violated. All hard constraints are converted to CNF using a Tseitin transformation.

Example 30. Consider the local DNCOP from Example 26. Its violation conditions are $\chi_{n_O} = n_O \wedge B_a e \wedge \neg\hat{B}_a c \wedge \neg r$ and $\chi_{n_F} = n_F \wedge \hat{B}_a c \wedge r$.

The fixed beliefs make n_O, n_F, and $B_a e$ true, while $B_a\neg c$ is false and hence $\hat{B}_a c = \neg B_a\neg c$ is true. Therefore, the translation gives $\widehat{\chi}_{n_O} = \bot$ and $\widehat{\chi}_{n_F} = r$. Since the decision constraint is $\top$, the hard clauses simplify to $\neg z_{n_O}$, $\neg z_{n_F} \vee r$, and $z_{n_F} \vee \neg r$, and the soft clauses are $(\neg z_{n_O}, 1)$ and $(\neg z_{n_F}, 1)$. Setting r false satisfies both soft clauses, while setting it true falsifies the second at cost one. Thus the unique optimal assignment encodes $d = \emptyset$, as found in Ex. 26.

Theorem 31. *For a decision-sound local DNCOP Δ, the construction above is a polynomial-time reduction to weighted partial MaxSAT. The restrictions of its optimal assignments to $\mathcal{D}$ encode exactly the subjectively optimal decisions. If Δ is also truthful, they also encode exactly the objectively optimal decisions.*

Proof. For $d \subseteq \mathcal{D}$, let I_d be the valuation defined by $I_d(\delta) = 1$ iff $\delta \in d$. By correctness of the Tseitin transformation, I_d can be extended to a valuation satisfying the hard clauses obtained from χ iff $d \Vdash_{\mathrm{CL}} \chi$, that is, iff $d \in \mathsf{Dec}(\Delta)$.

Now consider any $n \in N$. Cases (i) and (ii) reproduce the values fixed by the decision, case (iii) reproduces the acting agent's fixed beliefs about the remaining objective atoms, and case (iv) reproduces the fixed values of the remaining belief atoms. By locality, these cases cover every RDDA occurring directly in χ_n. A structural induction on χ_n therefore shows that $I_d \models \widehat{\chi}_n$ iff n is subjectively violated under d. The hard constraint $z_n \leftrightarrow \widehat{\chi}_n$ thus makes z_n true exactly when n is subjectively violated. Consequently, the total weight of the falsified soft clauses in any hard-satisfying extension of I_d equals $\mathsf{viol}^{\mathrm{sub}}(\Delta, d)$ by Proposition 27. Conversely, every assignment satisfying the hard constraints restricts to

I_d for some $d \in \mathsf{Dec}(\Delta)$. Minimizing the weight of the falsified soft clauses therefore yields exactly the subjectively optimal decisions.

If Δ is truthful, Prop. 28 gives $\mathsf{viol}^{\mathrm{sub}}(\Delta, d) = \mathsf{viol}^{\mathrm{obj}}(\Delta, d)$ for every $d \in \mathsf{Dec}(\Delta)$. Hence the subjective and objective optimal decisions coincide. Finally, the replacements increase the formula size at most linearly, and the Tseitin transformations are polynomial.

7 Conclusions and Future Work

We introduced a simple logical framework for norm-guided decision making under uncertainty. Starting from a restricted fragment of Chellas' Minimal Deontic Logic, we identified a language where deontic reasoning reduces to classical propositional reasoning while retaining obligations and prohibitions over atomic actions. We extended this fragment with explicitly represented conditional norms and factual detachment, obtaining a lightweight deontic logic.

To reason under incomplete information, we combined the framework with multi-agent doxastic logic, allowing norms to depend on agents' beliefs and beliefs about norms. We introduced the Doxastic Norm Compliance Optimization Problem (DNCOP), where an agent selects decisions minimizing weighted norm violations. By distinguishing optimization relative to the agent's beliefs from optimization relative to the objective state of the world, the framework captures a key aspect of autonomous decision making under uncertainty. We identified locality and truthfulness conditions under which these perspectives coincide, guaranteeing that subjectively optimal decisions are also objectively optimal.

The framework also admits efficient computation. For local and decision-sound DNCOPs, we showed that optimal decisions can be found by a polynomial-time reduction to weighted partial MaxSAT, leveraging mature optimization techniques while preserving logical rigour and transparency.

The framework deliberately favors simplicity over expressiveness. Restricting normative contents to atomic propositions enables reductions to propositional reasoning and straightforward integration with epistemic formalisms, while still covering a broad range of practical decision-making scenarios. Its modular design leaves room for richer normative and epistemic extensions.

Several directions remain for future work. We plan to analyze LDLB in more detail; we conjecture that LDLB-satisfiability is PSPACE-complete, as is plain multi-agent KD45 [15]. An immediate objective is to extend the optimization framework beyond local DNCOPs, where subjective optimization may require reasoning over multiple compatible belief models. We also plan to enrich the representation of agents' beliefs, allowing belief bases to capture more complex propositional information, including constitutive norms. Appropriate candidates are belief bases in the style of Lorini [19]. Another promising direction is to study decentralized settings in which multiple agents make decisions simultaneously while reasoning about one another's beliefs and intended actions. Finally, we intend to implement the MaxSAT translation and experimentally evaluate its scalability on representative normative decision-making benchmarks.

Acknowledgments. This research was funded in whole or in part by the Vienna Science and Technology Fund (WWTF) project ICT22-023 and the Austrian Science Fund (FWF) 10.55776/COE12 and 10.55776/PAT1815025.

Disclosure of Interests. The authors have no competing interests to declare.

AI Declaration The authors used generative AI tools for language editing and for checking and refining their theoretical results and proofs. They verified all mathematical content and take full responsibility for the paper.

References

1. Alferes, J.J., Gonçalves, R., Leite, J.: Equivalence of defeasible normative systems. J. Appl. Non Class. Logics **23**(1-2), 25–48 (2013)
2. Ansótegui, C., Bonet, M.L., Levy, J.: Solving (weighted) partial maxsat through satisfiability testing. In: SAT. Lecture Notes in Computer Science, vol. 5584, pp. 427–440. Springer (2009)
3. Åqvist, L.: Deontic logic. In: Gabbay, D., Guenthner, F. (eds.) Handbook of Philosophical Logic: Volume II, pp. 605–714. Springer (1984)
4. van Berkel, K., Ciabattoni, A., Freschi, E., Gulisano, F., Olszewski, M.: Deontic paradoxes in mīmāṃsā logics: There and back again. J. Log. Lang. Inf. **32**(1), 19–62 (2023)
5. Broersen, J.M., Dastani, M., van der Torre, L.W.N.: Beliefs, obligations, intentions, and desires as components in an agent architecture. Int. J. Intell. Syst. **20**(9), 893–919 (2005)
6. Cabalar, P., Ciabattoni, A., van der Torre, L.: Deontic equilibrium logic with explicit negation. In: JELIA. Lecture Notes in Computer Science, vol. 14281, pp. 498–514. Springer (2023)
7. Chellas, B.F.: Modal Logic - An Introduction. Cambridge University Press (1980)
8. Ciabattoni, A., Istenič Urh, B., Parent, X.: Deontic explanations in åqvist's systems. In: DEON. pp. 79–98. College Publications (2025)
9. Cooper, M.C., Herzig, A., Maffre, F., Maris, F., Perrotin, E., Régnier, P.: A lightweight epistemic logic and its application to planning. Artif. Intell. **298**, 103437 (2021)
10. Gabbay, D., Horty, J., Parent, X., van der Meyden, R., van der Torre, L. (eds.): Handbook of Deontic Logic and Normative Systems. College Publications (2013)
11. Gabbay, D., Horty, J., Parent, X., van der Meyden, R., van der Torre, L. (eds.): Handbook of Deontic Logic and Normative Systems, Volume 2. College Publications (2021)
12. Giordano, L., Martelli, A., Theseider Dupré, D.: Temporal deontic action logic for the verification of compliance to norms in ASP. In: ICAIL. pp. 53–62. ACM (2013)
13. Governatori, G., Olivieri, F., Rotolo, A., Scannapieco, S.: Computing strong and weak permissions in defeasible logic. J. Philos. Log. **42**(6), 799–829 (2013)
14. Governatori, G., Rotolo, A.: BIO logical agents: Norms, beliefs, intentions in defeasible logic. Auton. Agents Multi Agent Syst. **17**(1), 36–69 (2008)
15. Halpern, J.Y., Moses, Y.: A guide to completeness and complexity for modal logics of knowledge and belief. Artif. Intell. **54**(3), 319–379 (1992)
16. Hintikka, J.: Knowledge and Belief: An Introduction to the Logic of the Two Notions. Cornell University Press (1962)
17. Kowalski, R.A., Satoh, K.: Obligation as optimal goal satisfaction. J. Philos. Log. **47**(4), 579–609 (2018)
18. Lakemeyer, G., Lespérance, Y.: Efficient reasoning in multiagent epistemic logics. In: ECAI. Frontiers in Artificial Intelligence and Applications, vol. 242, pp. 498–503. IOS Press (2012)
19. Lorini, E.: Rethinking epistemic logic with belief bases. Artif. Intell. **282**, 103233 (2020)

20. Makinson, D., van der Torre, L.W.N.: Input/output logics. J. Philos. Log. **29**(4), 383–408 (2000)
21. Makinson, D., van der Torre, L.W.N.: Constraints for input/output logics. J. Philos. Log. **30**(2), 155–185 (2001)
22. Mill, J.S.: Utilitarianism. Parker, Son, and Bourn (1863)
23. Muise, C.J., Belle, V., Felli, P., McIlraith, S.A., Miller, T., Pearce, A.R., Sonenberg, L.: Planning over multi-agent epistemic states: A classical planning approach. In: AAAI. pp. 3327–3334. AAAI Press (2015)
24. Neufeld, E.A., Ciabattoni, A., Tulcan, R.F.: Combining MORL with restraining bolts to learn normative behaviour. In: IJCAI. pp. 4615–4623. ijcai.org (2025)
25. Parent, X.: Preference semantics for Hansson-type dyadic deontic logic: A survey of results. In: Gabbay, D., Horty, J., Parent, X., van der Meyden, R., van der Torre, L. (eds.) Handbook of Deontic Logic and Normative Systems, Volume 2, pp. 7–69. College Publications (2021)
26. Ross, A.: Imperatives and logic. Philos. Sci. **11**(1), 30–46 (1944)
27. Sergot, M.J., Sadri, F., Kowalski, R.A., Kriwaczek, F., Hammond, P., Cory, H.T.: The british nationality act as a logic program. Commun. ACM **29**(5), 370–386 (1986)
28. Tseitin, G.S.: On the complexity of derivation in propositional calculus. Zap. Nauchn. Sem. LOMI **8**, 234–259 (1968)
29. von Wright, G.H.: Deontic logics. Am. Philos. Q. **4**(2), 136–143 (1967)