Rights and Practical Reasoning in Deontic Logic

Huimin Dong [1]

Institute of Logic and Cognition, Sun Yat-sen University, China

Dragan Doder

Department of Computer Science, Utrecht University, the Netherlands

Xu Li, Réka Markovich, Leendert van der Torre

Department of Computer Science, University of Luxembourg, Luxembourg

Marc van Zee

Google Research, Denmark

Abstract

This paper brings together two traditions in deontic logic: the theory of normative positions, that is, reasoning about different types of rights, and practical reasoning, which has special relevance from the viewpoint of artificial intelligence (AI). We do this by exploring the role epistemic rights play in practical reasoning. Rights such as the right to know are intended to enable us to make informed decisions. They often play a role in determining what kind of plans we can make. A patient has the right to know his hospital test results so he can choose his treatment after his doctor has fulfilled her duty of informing him about possible risks and outcomes. This paper investigates, from the "database perspective", the role of (epistemic) rights in planning different scenarios from the database perspective and the dynamics of temporal beliefs and intentions. We take this perspective, extend the logic with deontic notions, and illustrate this with a running example.

Keywords: normative reasoning in AI, practical reasoning, normative positions

1 Introduction

Research in deontic logic includes a decades-long investigation into normative positions, benchmarked by Sergot's chapter in the first volume of the Handbook of Deontic Logic [24]. From the perspective of artificial intelligence

[1] Corresponding author: `donghm.logic@gmail.com`

(AI), practical reasoning is one of the most important topics in deontic logic and normative reasoning, benchmarked by Thomason's chapter in the second volume of the Handbook of Deontic Logic [27]. However, the topics of normative positions and practical reasoning are hardly ever brought together. The aim of this paper is to bring practical reasoning as used in AI to the field of deontic logic, with a special focus on the use of reasoning with rights from the database perspective [26,28].

In the tradition of reasoning about rights, the logics developed for normative positions (by Kanger [13] and Lindahl [15]) were initially aimed at mapping the space of logically possible legal relations between two given agents, differentiating between more and more variants [24]. These logics used a very weak action logic (Chellas called this system ET [7]) preventing the derivation of extensive consequences. Several more recent papers focusing on the conceptual elaboration of different notions of right like [17] adopt this approach. These logics thus have limited use for representing how an agent can reason practically about its actions in detail based on its own normative positions.

In contrast, most research on practical reasoning disregards rights and normative relations. BDI (Belief-Desire-Intention) logics (e.g., [8,22,28]) focus on specifying the relations between various mental states such as belief, desire, intention, and goal, but they traditionally ignore normative concepts. BOID (Belief-Obligation-Desire-Intention) [5] later incorporated obligations, but did not do so with normative positions.

This paper contributes to closing this gap by pointing out that in everyday life, we plan our actions by deliberating different scenarios. Our rights can play an important role in this planning, for instance when we come up with an optimal scenario where we have the right to do or get what we want. We start from the "database perspective" [26], a recent proposal that differentiates between a planner and belief-intention databases. The planner is engaged in some form of (temporal) practical reasoning, and in this process updates the databases. The task for the databases is to remain coherent. Van Zee *et al.* [28] formalized the databases using (Par)ameterized-time Action Logic (PAL) logic and providing AGM-like (Alchourrón-Gärdenfors-Makinson) postulates [1] for the revision of beliefs and intentions. Our main research question is: "how to characterize (epistemic) rights in terms of the role they play in practical reasoning". This is broken down into the following three sub-questions:

RQ1: the role and components of rights in practical reasoning;
RQ2: how to extend PAL [28] with the concepts needed for (epistemic) rights in practical reasoning;
RQ3: how to use this formal framework to model (epistemic) rights in practical reasoning.

We will characterize some variants of the right to know—with an emphasis on power—in terms of how they influence the dynamics of planning from the database perspective". The approach we use contributes to several aspects compared to previous research on the dynamic nature of normative positions (i.e., frameworks on the power type of right). For instance, from a conceptual point of view, it emphasizes the practical reasoning aspect, while from a technical point of view, it expresses the dynamics by using two revision operators as two kinds of coherence on a database. We will discuss this in detail at the end of the paper.

The layout of this paper follows the three research questions. In Section 2, we discuss the role of epistemic rights in practical reasoning and introduce the logic of intentions [28]. We extend the logic of intentions with obligation and permission in Section 3, and in Section 4 we apply the new logic to develop a revision operator to characterize Hohfeldian power. Section 5 ends the paper with our conclusions and future work.

2 Background

This section provides the background for this paper. First, we provide a short introduction to the theory of rights within the theory of normative positions, then we introduce a running example that we formalize throughout the paper. Finally, we give a summary of the database perspective.

2.1 (Epistemic) Rights in Deontic Logic: Theory of Normative Positions

The theory of normative positions in deontic logic refers to the tradition of formalizing normative relations between pairs of agents and their resultant relative positions. The theory relies on different meanings of the word "right" and their correlative duties put forward by Hohfeld. The tradition began with the work of Kanger and Kanger [13] and Lindahl [15], and has been developed by many others (e.g., [16,12,17,9]) more recently. The basic idea is that "right" can have different meanings, and the four atomic ones—in Hohfeldian terminology—are claim-right, privilege, power, and immunity. Each comes with its own correlative duty. That is, whenever an agent has one of these right positions, the counterparty has a duty position: duty (in the narrow sense), no-claim, liability, or disability, respectively. Claim-right is a claim that the duty bearer should take a particular action. Duty is the directed version of the classical notion of obligation in deontic logic. Privilege refers to the freedom of the right-holder to take a particular action when the counterparty has no claim to refrain him from doing so. This is the relationalized version of a weak permission. Power is when the right-holder has the possibility of changing the counterparty's normative positions with a special action. If a professor has the right to hand out homework, that means that she can create a duty for her students to do their homework. Immunity means that the counterparty does not have the power to change the right-holder's normative position. The interpretation of epistemic rights with Hohfeldian categories was put forward by the epistemologist Lani Watson [29], and the logical formalization of this interpretation has been articulated in some recent papers [18,19,14]. The formalization of the right to know in [18] and [19] uses the weak action logic referred to above.

2.2 Running Example

Avery (also called “Patient”, or simply P) suspects he has an illness that makes him eligible for early retirement, but he doesn’t want to apply until he is sure. He intends to get tested, knowing that he has the right to know the results as this is one of the patient’s listed rights under the law. After the tests, Avery exercises his right to know by asking for the results. The ‘right to know’ is understood as a Hohfeldian power by default. When Avery asks for the results, this puts an obligation on the doctor to inform him. That obligation means that Avery’s right to know becomes a claim-right. The doctor may intend to ignore the request, violating her obligation. This could make Avery give up his plan to apply for early retirement. Or he could decide to complain to the hospital director with the expectation that he would then get the results. After all, he still believes he has the right to be informed.

2.3 The Database Perspective

The logic of intentions has been studied in the fields of theory of mind and artificial intelligence. Van Zee et al. [28] provided a logic for reasoning about the dynamics of intentions and beliefs in time, formalizing Shoham’s database perspective [26]. This approach uses a temporal branching time logic called Parameterized-time Action Logic (PAL).

Definition 2.1 [The PAL Language] Let $\mathsf{Act} = \{a, b, c, \dots\}$ be a finite set of deterministic primitive actions and let $\mathsf{Prop} = \{p, q, r, \dots\} \cup \{pre(\bar{a}), post(a)\}$ be a finite set of propositions where $\bar{a} = (a_1, a_2, \dots)$ is a non-empty sequence of actions and $\{a, a_1, a_2, \dots\} \subseteq \mathsf{Act}$ are actions. The language $\mathscr{L}$ of the logic is as follows:

$$\varphi ::= \chi_t \mid do(a)_t \mid \Box_t \varphi \mid \varphi \wedge \varphi \mid \neg\varphi,$$

where $\chi \in \mathsf{Prop}$, $a \in \mathsf{Act}$, and $t \in \mathbb{N}$.

Intuitively, p_t means that p is true at time t, and $do(a)_t$ means that action a is executed at time t. Then, $pre(\bar{a})_t$ means that the precondition for a sequence of actions $\bar{a} = (a_1, ..., a_n)$ at time t is satisfied. Preconditions are defined on action sequences to ensure that it is possible to do the all the intended actions together (see the original paper [28] for more details). For instance, $pre(a_1, a_2, a_3)_0$ indicates that the precondition for doing a_1 at time 0, a_2 at time 1 and a_3 at time 2 is true. Then, $post(a)_t$ represents the postcondition for a at time t. The modal operator $\Box_t$ is interpreted as a *temporal* necessity for the planner, so a formula of the form $\Box_t p_{t'}$ means “it is necessary at time t that p is true at time t'”. This necessity means that no matter which actions are executed between time t and time t', p will hold in t'.

This provides a sound and strongly complete axiomatization. Due to space constraints, we only provide axioms relevant to this paper.

Definition 2.2 [Axiomatization (Partial)] Here are some PAL axioms. The full axiomatization can be found in the work of Van Zee [28], Section 2.3.

(A5) $\Box_t\varphi \rightarrow \Box_{t+1}\varphi$

(A6) $\bigvee_{a\in\mathsf{Act}} do(a)_t$

(A7) $do(a)_t \rightarrow \neg do(b)_t$, where $b \neq a$

(A8) $do(a)_t \rightarrow post(a)_{t+1}$

(A9) $pre(a)_t \rightarrow \Diamond_t do(a)_t$

(NEC) From φ, infer $\Box_0\varphi$

The intuitive meaning of some of the above axioms is explained below. Axiom A5 indicates continuity along the progression of time. If something is necessary at time t, then it remains necessary at the next time point $t+1$. Axiom A6 and Axiom A7 together state that at any given time point, one and only one action can be executed.

Due to space constraints, we omit the technical details of the semantics, and provide a short description only (see the work of Van Zee *et al.* [28], Section 2.2, for full details). PAL semantics is similar to that of computation tree logic (CTL)* [23] except that each transition between two consecutive states, the transition is also labeled by an action. A model (T,π) consists of a tree T and a path π. Trees have their root at time 0. Then, $T,\pi \models p_t$ means that proposition p is in the valuation function of the state corresponding to path π at time t (denoted as π_t). It follows that $T,\pi \models do(a)_t$ means that the transition from state π_t to π_{t+1} is labeled with action a. And $T,\pi \models \Box_t\varphi$ means that φ is true for all paths that are equivalent to π up to time t (i.e., they have the same states as π up to time t) in tree T. The other truth definitions are defined as per usual.

Note that the semantics distinguishes regular or strong beliefs from weak beliefs, which are beliefs contingent on the intended actions. The set $\mathbb{SB}$ of all *strong beliefs* is generated by Boolean combinations of $\Box_0\varphi$ where $\varphi \in \mathscr{L}$. A strong belief is an element of $\mathbb{SB}$. A set SB of strong beliefs is the deductive closure of a subset of $\mathbb{SB}$ such that $SB = Cn(\Sigma)$ where $\Sigma \subseteq \mathbb{SB}$. Semantically, a strong belief is a formula that is true for all the paths of the tree, meaning that they are independent of a specific future or *plan* (i.e., a specific sequence of intentions).

Definition 2.3 [Belief-Intention Database] A *belief-intention database* (SB, I) consists of a belief database SB and an intention database I:

- $SB \in \mathbb{SB}$ is a set of strong beliefs closed under consequence: $SB = Cn(SB)$;
- $I = \{(a_1,t_1),(a_2,t_2),\dots\}$ is a set of intentions (a_i,t_i) where $a_i \in \mathsf{Act}$ such that no two intentions exist at the same time point, i.e., if $i \neq j$ then $t_i \neq t_j$.

Weak beliefs are obtained by adding intentions to the strong beliefs and closing the result under consequence. Thus, a weak belief is closely related to a contingent or specific *plan.*

Definition 2.4 [Weak Beliefs] Given a belief-intention database (SB, I), weak beliefs are defined as follows:

$$WB(SB,I) = Cn(SB \cup \{do(a)_t | (a,t) \in I\}).$$

Commitment to intentions is characterized using a coherence condition stating that it is possible to perform all the intended actions.

Definition 2.5 [Coherence] Given an intention database $I = \{(b_{t_1},t_1),\dots,(b_{t_n},t_n)\}$ with $t_1 < \ldots < t_n$, let

$$Cohere(I) = \Diamond_0 \bigvee_{\substack{a_t \in Act : t \notin \{t_1,\ldots,t_n\} \\ a_t = b_t : t \in \{t_1,\ldots,t_n\}}} pre(a_{t_1}, a_{t_1+1}, \ldots, a_{t_n})_{t_1}.$$

More precisely, when we have a set of intended actions at non-consecutive time points $t_1, \ldots, t_n$, it is always possible at the initial time point 0 to carry out these intended actions by incorporating additional actions in the remaining time points. We say that a given belief-intention database (SB, I) is *coherent* iff SB is consistent with $Cohere(I)$, i.e., $SB \nvdash \neg Cohere(I)$.

A proposition relating weak beliefs to coherence [28] is shown below.

Proposition 2.6 *Given a belief-intention database* (SB, I)*, if* (SB, I) *is coherent, then* $WB(SB, I)$ *is consistent.*

Revision operators are then defined for both beliefs and intentions. The ones presented here are almost the same as those of Van Zee *et al.* but are slightly simpler. [2]

Definition 2.7 [Intention revision function] An *intention revision function* $\otimes$ maps a belief-intention database and an intention to a belief-intention database such that

$$(SB, I) \otimes i = (SB, I'),$$

where the following postulates hold:
$(P1)$ (SB, I') is coherent;
$(P2)$ If $(SB, \{i\})$ is coherent, then $i \in I'$;
$(P3)$ If $(SB, I \cup \{i\})$ is coherent, then $I \cup \{i\} \subseteq I'$;
$(P4)$ $I' \subseteq I \cup \{i\}$;
$(P5)$ For all I'' with $I' \subset I'' \subseteq I \cup \{i\}$:$(SB, I'')$ is not coherent.

Postulate (P2) states that new intention i takes precedence over all other current intentions. If possible, it should be added even if all current intentions must be discarded. Postulate (P3) and (P4) together state that if it is possible to simply add the intention, then this is the only change that is made. These two postulates are comparable to the inclusion and vacuity of AGM. Finally, (P5) states that we do not discard intentions unnecessarily.

Definition 2.8 [Belief revision function] A *belief revision function* $\circ$ maps a belief-intention database and a strong belief formula φ to a belief-intention database such that

$$(SB, I) \circ \varphi = (SB', I'),$$

where:

- SB' is the result of revising SB with a φ that satisfies the AGM postulates [1],

[2] Our revision operators differ from those of Van Zee *et al.* in three ways. 1) They bind their revision operators up to a time point t, which is a mere technical detail to prove a representation theorem, so we leave this out. 2) For technical reasons, they represent a belief set SB as a propositional formula ψ such that $SB = \{\varphi | \psi \vdash \varphi\}$, but we simply use SB directly. 3) They define a revision operator for revising with the pair (φ, i), which is slightly more general than our variant but is used only for edge cases.

- I' is the result of revising the new beliefs with the empty intention ϵ so that coherence is restored, i.e., $(SB', I) \otimes \epsilon = (SB', I')$.

Note that, by this definition, the revision of strong beliefs cannot be triggered by intention revision, but it can trigger intention revision. Intuitively, this makes sense: one would not wish to change one's strong beliefs after adopting an intention, but might want to update one's intentions after learning new information.

3 Formalizing Obligation (and Claim-Right)

In this section, we formalize obligation—and thus also claim-right, its corresponding notion in the normative position theory—from the database perspective, while not extending them in any way. It turns out that we are able to model these concepts quite naturally using only beliefs and intentions. In the next section, we extend the coherence condition so that we are able to use deontic notions when revising with new information.

We model the doctor-patient example with only one belief base and one intention base. In this case, the beliefs may be seen as shared or common beliefs, and the intentions may be seen as shared intentions. [3]

In our minimal formalization, we introduce only some special actions such as test, ask, and inform. And we introduce only some special propositions such as pre/postconditions for actions and a violation constant for obligations.

3.1 Actions

To model the action that agent i informs agent j about proposition p, we use the action $inform(i, j, p)$. And we use $is\text{-}informed\text{-}whether(i, p)$ propositions (abbreviated as $iiw(i, p)$) to model whether agent i is informed about the truth or falsehood of p. We have that $iiw(i, p)$ is a precondition of $inform(i, j, p)$ and $iiw(j, p)$ is a postcondition of that action.

We assume that the doctor can not only learn whether p is true or false by being informed but can also carry out tests to find out. So $test(i, p)$ has postcondition $iiw(i, p)$.

Since p (whether the patient is ill) is the focus and is always repeated in our running example, we simplify things by omitting it from the actions and propositions below.

Example 3.1 [Running example] Let $\mathsf{Prop} = \{iiw(\mathtt{D}), iiw(\mathtt{P}), \mathtt{v}\}$ and $\mathsf{Act} = \{test(\mathtt{D}), ask(\mathtt{P}, \mathtt{D}), inform(\mathtt{D}, \mathtt{P}), ignore(\mathtt{D}, \mathtt{P}), complain(\mathtt{P}, \mathtt{HD})\}$. These are interpreted as follows:

- $test(\mathtt{D})$: the doctor tests whether the patient is ill;
- $ask(\mathtt{P}, \mathtt{D})$: the patient asks the doctor whether he is ill or not;
- $inform(\mathtt{D}, \mathtt{P})$: the doctor informs the patient whether he is ill or not;

[3] Note that this means that the revision operators aren't revision operators for a particular agent but for the entire system. Thus, if we revise intentions related to a particular agent, this may affect the intentions of other agents.

- $ignore(\mathtt{D},\mathtt{P})$: the doctor ignores the patient's request;
- $complain(\mathtt{P},\mathtt{HD})$: the patient complains to the director;
- $iiw(\mathtt{D})$ / $iiw(\mathtt{P})$: the doctor/patient is informed whether the patient is ill;
- v: a violation occurs.

While PAL defines pre- and postconditions as primitive propositions, we introduce the following abbreviations in our running example:

(i) $post(test(\mathtt{D})) = iiw(\mathtt{D})$: after the doctor has carried out the tests, she knows whether the patient is ill or not;
(ii) $pre(inform(\mathtt{D},\mathtt{P})) = iiw(\mathtt{D})$: the doctor can only inform the patient if she knows whether the patient is ill or not;
(iii) $post(inform(\mathtt{D},\mathtt{P})) = iiw(\mathtt{P})$: after the doctor has informed the patient, he knows whether he is ill or not;
(iv) $post(ask(\mathtt{P},\mathtt{D})) = pre(ignore(\mathtt{D},\mathtt{P}))$: the doctor can only ignore the request if the patient has made the request.
(v) $post(ask(\mathtt{P},\mathtt{D})) = pre(inform(\mathtt{D},\mathtt{P}))$: the doctor can inform the patient whether he is ill or not upon request;
(vi) $post(ignore(\mathtt{D},\mathtt{P})) = pre(complain(\mathtt{P},\mathtt{HD}))$: the patient can only complain to the director if the doctor ignores his request.

While PAL defines preconditions for action sequences as primitive propositions, we use the following inductive definition so that we can also include the precondition formulas above in preconditions for action sequences: $pre(a,\bar{b})_t = pre(a)_t \wedge \Diamond_t(do(a)_t \rightarrow pre(\bar{b})_{t+1})$.

We can use PAL axiomatization (Def. 2.2) and the above formulas to derive new formulas:

- $do(test(\mathtt{D}))_t \rightarrow iiw(\mathtt{D})_{t+1}$ (A8, (i));
- $\Box_0(do(inform(\mathtt{D},\mathtt{P}))_t \rightarrow iiw(\mathtt{P})_{t+1})$ (A8, (iii), NEC);
- $iiw(\mathtt{D})_t \rightarrow \Diamond_t do(inform(\mathtt{D},\mathtt{P}))_t$ ((ii), A9);
- $do(ask(\mathtt{P},\mathtt{D}))_t \rightarrow$
  $(\Diamond_{t+1} do(ignore(\mathtt{D},\mathtt{P}))_{t+1} \wedge \Diamond_{t+1} do(inform(\mathtt{D},\mathtt{P}))_{t+1}$ (A8, (iv), (v), A9).

Example 3.2 [Running Example (cont'd.)] Avery suspects he has an illness, so he intends to get tested, knowing he has a right to know the results. We formalize this as the following strong belief formula:

$$RK = \Box_0 \left[do(test(\mathtt{D}))_0 \wedge do(ask(\mathtt{P},\mathtt{D})_1) \rightarrow \Box_2(\neg do(inform(\mathtt{D},\mathtt{P})_2 \rightarrow \mathsf{v}_3)\right].$$

RK should be understood as: the doctor ought to inform the patient of the test results if the patient has had the tests and has asked for his test results; otherwise, a violation occurs. That is, this is a power type of right: the duty occurs once the patient asks for the results. Note that this is not supposed to be a general definition of the power to know; it describes actions that are preconditions for the duty to hold in this setting. One is the duty-creating action of the patient, the other is a practical precondition: the patient has to get tested before he can be informed of any kind of result.

Using this formalization, we now provide the initial belief-intention database for our running example.

Example 3.3 [Running Example (cont'd.)] Initially, there are no intentions, and the only belief under consideration is: Avery has the right to know whether he is ill. Since we would like to be able to reason about obligations and what happens when a violation occurs, we use RK from Example 3.2 to formalize the right to know. Formally, the initial belief-intention database is (SB, I), where

$$SB_0 = Cn(RK) \text{ and } I_0 = \emptyset.$$

Because the patient has no action he intends to carry out, his set of weak beliefs $WB(SB_0, I_0) = SB_0$ is the set of strong beliefs.

Next, we add two intentions using the intention revision operator. Notice that we actually have two agents. Avery is the agent we consider from the planning point of view, and he reasons about the doctor's obligation when planning. He derives the doctor's obligation from his weak beliefs since he still needs to ask to be informed, and he reasons from his strong beliefs after he has made his request.

Example 3.4 [Running Example, revision with intentions (cont'd.)] After a process of planning, the following two intentions are added: $i_1 = (test(\mathtt{D}), 0)$ and $i_2 = (ask(\mathtt{P}, \mathtt{D}), 1)$. Since both these intentions cohere with the current beliefs, they can simply be added to the intention database.

More formally, using postulates [P3] and [P4] (Def. 2.7), we obtain

$$((SB_0, I_0) \otimes i_1) \otimes i_2 = (SB_1, I_1),$$

where $SB_1 = SB_0 = Cn(RK)$ (revision of intentions cannot change strong beliefs), and $I_1 = \{(test(\mathtt{D}), 0), (ask(\mathtt{P}, \mathtt{D}), 1)\}$.

Note that $WB(SB_1, I_1) \vdash \Box_2(\neg do(inform(\mathtt{D}, \mathtt{P})_2 \rightarrow \mathtt{v}_3)$ (Def. 2.4), which means that that the doctor should inform the patient of his test results at time 2; otherwise a violation occurs at time 3. We simply consider the obligation derived from weak beliefs as the result of exercising a legal power.

This is the point of power type of rights: one can plan with them with the knowledge that by carrying out these actions, the other party will have a duty. Hence, if carrying out the action (of asking) is among my intentions, then the obligation of the other person will be among those postconditions that depend on the actions I intend to carry out. That is, the obligation will be derivable from weak beliefs.

Next, we model the belief database with an action a executed at time 0. This is something that was not investigated by Van Zee *et al.* [28]. We model this simply by adding the strong belief $\Box_0 do(a)_0$, which states that some action is necessarily carried out. Intuitively, this ensures that everything that follows from executing a at time 0 is now a strong belief. So, for instance, $\Box_0 post(a)_1$ now also holds, as well as everything that follows from that.

Example 3.5 [Running Example, revision with strong beliefs (cont'd.)] Next, the doctor carries out the tests, which we model by adding the strong belief

$$(SB_1, I_1) \circ \Box_0 do(test(\mathtt{D}))_0 = (SB_2, I_2),$$

where

- $SB_2 = Cn(\{RK, \Box_0 do(test(\mathtt{D}))_0\})$;
- $I_2 = I_1 = \{(test(\mathtt{D}), 0), ask(\mathtt{P}, \mathtt{D}), 1)\}$ (adding that the strong belief did not invalidate any intentions).

We can now infer the following power relationship between the patient and the doctor: if the doctor does not provide the test result upon request, there is a violation. In other words, the doctor is obliged to provide the test results:

$$SB_2 \vdash \Box_0(do(ask(\mathtt{P}, \mathtt{D})_1 \rightarrow \Box_2(\neg do(inform(\mathtt{D}, \mathtt{P})_2) \rightarrow \mathsf{v}_3).$$

Next, the patient requests his test results:

$$(SB_2, I_2) \circ \Box_0 do(ask(\mathtt{P}, \mathtt{D}))_1 = (SB_3, I_3),$$

where

- $SB_3 = Cn(\{RK, \Box_0 do(test(D))_0, \Box_0 do(ask(\mathtt{P}, \mathtt{D}))_1\})$;
- $I_3 = I_2 = I_1 = \{(test(\mathtt{D}), 0), ask(\mathtt{P}, \mathtt{D}), 1)\}$.

We can infer the next claim-right relationship between the patient and the doctor: if the doctor does not inform the patient, there is a violation:

$$SB_3 \vdash \Box_0(\neg do(inform(\mathtt{D}, \mathtt{P})_2) \rightarrow \mathsf{v}_3).$$

We will formalize the obligation, claim-right, and legal power involved in the above examples more precisely in the next section.

3.2 Obligations and Claim-Rights in the Logic of Intentions

In deontic logic, deontic concepts such as obligation and permission are considered to be deontic variants of necessity and possibility [11]. Following this tradition, our database framework represents modalities for deontic concepts utilizing temporal modalities, taking *deontic* necessity and possibility as temporal modalities of necessity or possibility *to plan what is normative*. We will define obligation, permission and prohibition in the sense of "ought to do" [11], representing them as deontic modalities on individual actions. They are defined in the style of Anderson reduction [2].

Definition 3.6 [Obligation, Permission, and Prohibition] Given $t \in \mathbb{N}$ and $a \in \mathsf{Act}$:

- an action a that is allowed to be carried out at time t, denoted as $P(a)_t$, is defined as $\Diamond_t(do(a)_t \wedge \neg \mathsf{v}_{t+1})$;
- an action a that ought to be carried out at time t, denoted as $O(a)_t$, is defined as $\Box_t(\neg do(a)_t \rightarrow \mathsf{v}_{t+1})$;
- an action a that is prohibited from being carried out at time t, denoted as $F(a)_t$, is defined as $\Box_t(do(a)_t \rightarrow \mathsf{v}_{t+1})$.

These three deontic modalities are defined on single actions but not consecutive actions (denoted by $\overline{a}$, see Def. 2.1). For instance, if $\overline{a} = (a_1, a_2)$ with $a_1, a_2 \in$ Act, then $O(\overline{a})_3$ is not a correct expression.

In the PAL language, $P(a)_t$ means that it is possible at time t to do action a and not have a violation in the next time point. Then, $O(a)_t$ means that it must be the case that if at time t action a is not executed, there is a violation in the next time points, and $F(a)_t$ means that it must be the case that if action a is executed, there is a violation in the next time point.

Next, we extend the logic of Van Zee *et al.* with a new axiom stating that it is always possible to avoid a violation.

Definition 3.7 [Avoiding Violation Axiom] We add the following axiom to the axiomatization of Van Zee *et al.* (see [28], Section 2.3) [4]: $\Diamond_t \neg v_{t+1}$.

We now obtain the following proposition. We omit the proof since it follows straightforwardly from the definition of $O(a)_t$, $P(a)_t$ and the new axiom.

Proposition 3.8 (Obligation Implies Permission) *If we add the Avoiding Violation Axiom, $O(a)_t \rightarrow P(a)_t$ is a theorem of the logic.*

To capture claim-rights, we show how our deontic concepts can be included in the strong beliefs given a belief-intention database.

Example 3.9 [Running example, claim and privilege (cont'd.)] Recall from the previous example that $SB_3 = Cn(\{RK, do(test(\mathtt{D}))_0, do(ask(\mathtt{P},\mathtt{D}))_1\}$ and that we could then infer the following:

$$SB_3 \vdash \Box_0(\neg do(inform(\mathtt{D},\mathtt{P})_2) \rightarrow \mathsf{v}_3).$$

Using Def. 3.6 and Axiom A5 (Def. 2.2), it follows that an obligation is inferred:

$$SB_3 \vdash O(inform(\mathtt{D},\mathtt{P}))_2.$$

Thus, after the patient has asked for his result at time point 1, the doctor has an obligation to inform the patient of the result at time 2. Therefore, Avery now has a claim-right that the doctor should inform him of the result.

We obtain other types of deontic concepts if we update the databases differently. For instance, assume the following strong belief formula b_1:

$$\Box_0(do(test(\mathtt{D}))_0 \rightarrow \Diamond_1((inform(\mathtt{D},\mathtt{P}))_1 \wedge \neg \mathsf{v}_2)),$$

and suppose we update the belief-intention database, after carrying out the tests specified in the planner (Example 3.5), as follows:

$$(SB_2, I_2) \circ b_1 = (SB_2', I_2').$$

Now the following permission can be inferred:

$$SB_2' \vdash P(inform(\mathtt{D},\mathtt{P}))_1,$$

[4] Due to space constraints, we omit the semantics here, but if we add a property to the definition of the model (see [28], Def. 6) stating that in each state there exists an action transition such that in the next time moment $\neg v$ holds, then we can straightforwardly prove that the logic remains sound and strongly complete.

which states that the doctor has a permission to inform the patient of the result, which then indicates in this belief base that the patient has the privilege of requesting that the doctor informs him of the result.

Similarly, if we add the following strong belief b_2:

$$\Diamond_0(do(test(\mathtt{D}))_0 \wedge \neg \mathsf{v}_1),$$

we have this revision of the belief-intention:

$$(SB_2, I_2) \circ b_2 = (SB_2'', I_2'').$$

Now we conclude with another permission as a strong belief in this database:

$$SB_2'' \vdash P(test(\mathtt{D}))_0.$$

So the patient has the privilege, given his strong beliefs SB_2'', of expecting the doctor to carry out the tests.

The permission and prohibition of an action cannot simply be reduced to an action obligation, as shown in the following proposition. Proposition 3.10 shows how obligation, permission, and prohibition can be connected. In particular, Proposition 3.10 (iii) and (iv) shows that a variant of the dual relation between obligation and permission exists.

Proposition 3.10 *Given $t \in \mathbb{N}$ and $a \in \mathsf{Act}$, the following propositions are theorems in our logic.*

(i) $F(a)_t \leftrightarrow \neg P(a)_t$;

(ii) $F(a)_t \rightarrow \bigvee_{b \neq a} P(b)_t$;

(iii) $P(a)_t \rightarrow \bigwedge_{b \neq a} \neg O(b)_t$;

(iv) $O(a)_t \leftrightarrow \bigwedge_{b \neq a} \neg P(b)_t$.

Note that the last part of Proposition 3.10(iv) implies that if an action is obligatory, then no other action can be permitted. In our logic, the property is a consequence of a practical interpretation of A7. It clarifies our key understanding about actions from the database perspective: if an action is executed at some time point, no other action can be performed at the same time. This leads to the conclusion that if we are obligated to do action a, we are not allowed to engage in other actions as that would prevent us from executing a. This property does not necessarily fit the understanding on norms or the law from a deontic point of view, but it fits well from a database perspective.

4 Optimality and Power

In the previous section, we formalized static deontic concepts such as obligation and permission using a violation constant. But because the coherence condition of Van Zee *et al.* does not use this information, we were not able to use it when revising with new beliefs or intentions. In this section, we propose a new condition, stronger than coherence, called "optimality": if a belief-intention database is optimal, then it is coherent, and it avoids violation states. We show that this new coherence condition can be used to revise belief-intention databases satisfying the deontic notions we proposed in the previous section.

Definition 4.1 [Optimality] Given an intention database $I = \{(b_{t_1}, t_1), \dots, (b_{t_n}, t_n)\}$ with $t_1 < \cdots < t_n$, let

$$Opt(I) = \Diamond_0 \bigvee_{\substack{a_t \in Act : t \notin \{t_1,\dots,t_n\} \\ a_t = b_t : t \in \{t_1,\dots,t_n\}}} (pre(a_{t_1}, a_{t_1+1} \dots, a_{t_n})_{t_1} \wedge \bigwedge_{t_1 \leq i \leq t_n} (do(a_i)_i \rightarrow \neg \mathsf{v}_{t_{i+1}})).$$

For a given belief-intention database (SB, I), we say that it is optimal iff SB is consistent with $Opt(I)$, i.e., $SB \nvdash \neg Opt(I)$.

Note that the above definition requires not only that the actions intended don't lead to a new violation state but also that the other possible actions that may be carried out should act as a bridge on the path from t_1 to t_n. It ensures that no new violation can occur from t_2 to t_{n+1}. For example, $Opt(\{(a, 1), (c, 3)\})$ requires the execution of some action b at time 2 bridging a at time 1 and c at time 3 without any new violations from time 2 to time 4.

Definition 4.2 [Postulates of Optimal Revision] An *intention revision function* $\bullet$ maps a belief-intention database and an intention to a belief-intention database such that

$$(SB, I) \bullet i = (SB, I')$$

where the postulates that hold for optimality are similar to the postulates for intention revision (Def. 2.7) $(P1)$–$(P5)$, except that the condition of coherence is replaced by the condition of optimality.

In order to specify the distinction between the coherence and optimality conditions, we continue our discussion of the running example and now consider the action $ignore(\mathtt{D}, \mathtt{P})$, which means that the doctor ignores the patient's request.

Example 4.3 [Coherent Intentions vs. Optimal Intentions (Ctd.)] Recall that $SB_2 = Cn(RK) \circ \Box_0(do(test)_0)$ (Example 3.5). We consider two possible intention databases:

- $I = \{(inform(\mathtt{D}, \mathtt{P}), 2)\}$;
- $I' = \{(ignore(\mathtt{D}, \mathtt{P}), 2)\}$.

Now we have the following implications:

$$Cohere(I) = \Diamond_0 pre(inform(\mathtt{D}, \mathtt{P}))_2;$$
$$Cohere(I') = \Diamond_0 pre(ignore(\mathtt{D}, \mathtt{P}))_2.$$

Thus, both I and I' cohere with strong belief SB_2. However, only I is optimal with SB_2. The intention database I' is not optimal because there is a violated state that necessarily occurs after the intended action is executed:

$$SB_2 \vdash \Box_0(do(ignore(\mathtt{D}, \mathtt{P}))_2 \rightarrow \mathsf{v}_3).$$

This formula follows from SB_2 because, informally, it means that for all the paths in which $ignore(\mathtt{D},\mathtt{P})$ is executed at time 2, violations will occur at time 3. This is true because in each such path, with $ask(\mathtt{D},\mathtt{P})$ occurring at time 1, the doctor can ignore the patient's request for his results (recall that $post(ask(\mathtt{P},\mathtt{D})) = pre(ignore(\mathtt{D},\mathtt{P}))$ and A9).

So intention base I is optimal but intention base I' is not:

$$SB_2 \vdash Opt(I) \text{ and } SB_2 \vdash \neg Opt(I').$$

Consequently, the optimal revision of the belief-intention database (SB_2, I_0) (recall that $I_0 = \emptyset$) will not incorporate the action $ignore(\mathtt{D},\mathtt{P})$ at time 2, unlike the coherent revision:

- $(SB_2, I_0) \bullet (ignore(\mathtt{D},\mathtt{P}), 2) = (SB_2, I_0)$;
- $(SB_2, I_0) \otimes (ignore(\mathtt{D},\mathtt{P}), 2) = (SB_2, \{(ignore(\mathtt{D},\mathtt{P}), 2)\})$.

We introduced optimal revision because it prevents an artificial agent (like a robot) from remaining committed to an intended action that leads to violations and helps it to make and revise legal plans. On the other hand, violations do occur in practice, and therefore we should also allow reasoning about the dynamics of intentions (like contrary-to-duty reasoning [21]) to account for those situations. We will use a coherence condition (see Example 4.5) for this purpose.

Example 4.4 [Running example, power (cont'd.)] We know that the strong beliefs set SB_2 does not contain the following two deontic concepts:

$$SB_2 \nvdash O(inform(\mathtt{D},\mathtt{P}))_2 \text{ and } SB_2 \nvdash F(ignore(\mathtt{D},\mathtt{P}))_2.$$

Now by updating the database with intention $(ask(\mathtt{D},\mathtt{P}), 1)$, we can see that an obligation exists in the weak beliefs of the updated database $(SB_2, \{(ask(\mathtt{D},\mathtt{P}), 1)\})$:

$$WB((SB_2, \{(ask(\mathtt{D},\mathtt{P}), 1)\})) \vdash O(inform(\mathtt{D},\mathtt{P}))_2;$$
$$WB((SB_2, \{(ask(\mathtt{D},\mathtt{P}), 1)\})) \vdash F(ignore(\mathtt{D},\mathtt{P}))_2.$$

After the tests have been carried out, patient Avery has a Hohfeldian power. If Avery exercises that power by asking for the results, then he will have a claim-right that the doctor informs him of the result. If the patient does not intend to ask for the result, the doctor cannot be obliged to inform the patient.

If, instead of forming the intention $(ask(\mathtt{D},\mathtt{P}), 1)$, the planner has the action $ask(\mathtt{D},\mathtt{P})$ that is actually executed at time 1, then we obtain the revised strong beliefs set SB_3 (see Example 3.5). The same obligation and prohibition exist, but since the claim-right of Avery (and thus the corresponding duty of the doctor) was created by his request, now the obligation (and the prohibition) follows from the strong beliefs:

$$SB_3 \vdash O(inform(\mathtt{D},\mathtt{P}))_2 \text{ and } SB_3 \vdash F(ignore(\mathtt{D},\mathtt{P}))_2.$$

The question arises: what happens if the doctor ignores the request, violating her duty? This scenario leads to contrary-to-duty reasoning [21]. It is very intuitive to say that Avery's right to know his results must include a "solution" for when the newly created claim-right's corresponding duty is violated. Indeed, Avery has a new intention for which this violation is a precondition: a complaint to the hospital director.[5] The example below shows what can be done in the current version of the logic.

Example 4.5 [Running example, contrary-to-duty reasoning (cont'd.)] Assume that following Avery's request to the doctor at time 1, the doctor intends to ignore him:

$$(SB_3, I_0) \otimes (ignore(\mathtt{D}, \mathtt{P}), 2) = (SB_3, I_4),$$

where $I_4 = \{(ignore(\mathtt{D}, \mathtt{P}), 2)\}$. To recover from this bad situation, Avery will have a new intention: complain to the hospital director. Intuitively, this corresponds to contrary-to-duty scenarios in deontic logic literature [21], which is about how to recover when the primary obligation is violated. Therefore we have:

$$(SB_3, I_4\}) \otimes (complain(\mathtt{P}, \mathtt{HD}), 3) = (SB_3, I_5),$$

where $I_5 = \{(ignore(\mathtt{D}, \mathtt{P}), 2), (complain(\mathtt{P}, \mathtt{HD}), 3)\}$. Recall that we assume that the database is shared. After the doctor understands that Avery intends to complain to the hospital director, she revises her intention and decides to let Avery know his test results:

$$(SB_3, I_5\}) \otimes (inform(\mathtt{D}, \mathtt{P}), 2) = (SB_3, I_6).$$

Here we have that $I_6 = \{(inform(\mathtt{D}, \mathtt{P}), 2)\}$. The intention $(ignore(\mathtt{D}, \mathtt{P}), 2)$ is dropped because only one intention is possible at time 2, and the new intention takes priority (according to P2 from Def. 2.7). Then $(ignore(\mathtt{D}, \mathtt{P}), 2)$ must be dropped as well because its precondition will not hold at time 3.

It would be rather intuitive to allow Avery to model conditional planning by adding to the database both his intention to complain and his intention to submit a request at time 3, depending on how the situation develops (i.e., whether the doctor informs him or ignores his request). But they have incompatible preconditions. The precondition for the complaint action is the postcondition of the ignore action, while the precondition for applying for early retirement requires that Avery is informed[6]). The agent will drop the action whose precondition is not met. In any case, the current system does not allow two intentions at the same time point, so we leave this as future work.

5 Conclusions and Future Work

Rights, including epistemic rights, influence our plans and thus the intentions we assume or discard. Avery wouldn't have gotten tested if he hadn't believed that he would get the information he needed to apply for early re-

[5] This complaint action is very similar to asking for test results; it imposes a duty on the hospital director to inform Avery (or make the doctor inform Avery). This duty can also be violated, but we do not go that far into the reasoning in this paper.

[6] For the sake of simplicity, we haven't formally added the action *submit*(P) and its pre- and postconditions to the language since we haven't used them in the example.

tirement. In order to accommodate reasoning about normative positions in a framework, we need basic deontic concepts such as obligation and permission: these could be introduced through a violation constant. We also need some formalism to express the nature of power: that some specific power action can result in changes to normative positions. We could express this by updating obligations so that weak beliefs become strong beliefs once the duty-bound action has been carried out. Additionally, one of the most characteristic features of the theory of normative positions is that we consider pairs of agents and their relations. In this paper, we considered only two agents, thus the relation between their normative positions could be handled tacitly.

We employed the PAL temporal logic of intentions [28] to reason about obligations, permissions, and rights by modeling the dynamics of intentions and beliefs. We were able to model obligations and claim-rights directly in the PAL framework and without extending it in any way. However, we did extend the revision of belief-intention databases in two ways. First, we introduced Optimal Revision, which revises the databases so that no violation can occur and prevents artificial agents from having illegal intentions. Secondly, we introduced revision of databases after actions have been carried out in order to model the nature of power (transforming weak beliefs into strong beliefs). This framework thus introduces a new way of characterizing Hohfeldian rights in practical reasoning.

In conclusion, this paper contributes to closing the gap between reasoning about rights and practical reasoning. On the one hand, the deontic concepts introduced to the framework make it possible to align the plans of artificial agents with norms. These agents are (or will be) subject to normative expectations and will have normative positions based on the deontic concepts and optimality condition involved in planning to make these possible. On the other hand, deontic logic, including the theory of normative positions, is ultimately about defining and reasoning about the normative aspect of actions. A richer action logic contributes to fulfilling its full potential.

Our future research needs to address the "ought to be" question. When it comes to rights, it seems very natural at first to talk about actions, and so "ought to do" appears to be an adequate concept to work with. In deontic logic, it is also very natural to consider "ought to be" and compare this to "ought to do" [11]. It is particularly relevant if we consider the planning aspect: the normative goal is taken as an "ought to be", and it is the role of the planner to assign the obligation to an agent to fulfill the normative plan.[7] However, from a technical point of view, defining "ought to be" in the database is more complicated than defining "ought to do". We cannot simply represent "It ought to be the case that χ at time t" as $\Box_t(\neg\chi_t \rightarrow \mathsf{v}_t)$, because axiom $\mathsf{A8}$ makes the temporal modality redundant. To maintain the temporal necessity of

[7] In fact, this also fits what happens with rights. For instance, a legislative agent that signs the Convention of Human Rights is obliged to assign corresponding duties in its own legal system.

planning normatively while avoiding temporal redundancy, one could consider the following definition of "ought to be": $O(\chi_t) := \Box_{t-1}(\neg\chi_t \to \mathsf{v}_t)$. This states: "It is necessary to plan at time $t-1$ that χ will be the case at time t if no violation occurs", which makes sense as a way to describe χ as a normative goal for the planner. However, the proper formalization of "ought to be" remains to be studied.

As compared to existing theories of agents and norms, our proposal highlights the crucial role of belief and intention in normative reasoning. Traditional logic-based methods, including dynamic deontic logic [20], see-to-it-that (STIT) logics [3], and labeled transition systems [25], encompass a wide range of deontic and temporal operators, which are interpreted using semantic models like CTL*. Our logic also uses CTL*-like models and fairly simple syntax based on PAL [28]. The framework is expressive enough to model rights and define deontic operators but is simple enough to perform AGM-style revision of belief and intention, and is therefore suitable for practical reasoning. It can be extended to address issues related to physical or normative constraints, such as environmental persistence [25], multi-agent interaction within the context of personal intentions [4,10], and the trade-off between violation and compliance [6]. We leave these topics for future work.

Acknowledgments

We would like to express our sincere gratitude for the valuable comments provided by the three referees of DEON 2023. This work was supported by the *Fonds National de la Recherche* of Luxembourg through the project Deontic Logic for Epistemic Rights (OPEN O20/14776480) and the project INTEGRAUTO (INTER/AUDACE/21/16695098). L. van der Torre is financially supported by the CHIST-ERA grant CHIST-ERA19-XAI (G.A. INTER/CHIST/19/14589586) and by the European Union's Justice program under grant agreement 101007420 (ADELE). Huimin Dong, adjunct research fellow at the Institute of Logic and Cognition at Sun Yat-sen University, is supported by the Chinese Ministry of Education (MOE) project of the Key Research Institute of Humanities and Social Sciences at Universities (22JJD720021).

References

[1] Alchourrón, C. E., P. Gärdenfors and D. Makinson, *On the logic of theory change: Partial meet contraction and revision functions*, Journal of Symbolic Logic **50** (1985), pp. 510–530.

[2] Anderson, A. R., *A reduction of deontic logic to alethic modal logic*, Mind **67** (1958), pp. 100–103.
URL `http://www.jstor.org/stable/2251344`

[3] Belnap, N., M. Perloff and M. Xu, "Facing the future: agents and choices in our indeterminist world," Oxford University Press, 2001.

[4] Bratman, M., "Intention, plans, and practical reason," Harvard University Press, 1987.

[5] Broersen, J., M. Dastani, J. Hulstijn, Z. Huang and L. van der Torre, *The boid architecture: Conflicts between beliefs, obligations, intentions and desires*, in: *Proceedings*

of the Fifth International Conference on Autonomous Agents, AGENTS '01 (2001), p. 9–16.
URL `https://doi.org/10.1145/375735.375766`

[6] Broersen, J. and L. van der Torre, *Ten problems of deontic logic and normative reasoning in computer science*, Lectures on Logic and Computation: ESSLLI 2010 Copenhagen, Denmark, August 2010, ESSLLI 2011, Ljubljana, Slovenia, August 2011, Selected Lecture Notes (2012), pp. 55–88.

[7] Chellas, B. F., "Modal Logic. An Introduction." Cambridge University Press, 1980.

[8] Cohen, P. R. and H. J. Levesque, *Intention is choice with commitment*, Artificial intelligence **42** (1990), pp. 213–261.

[9] Dong, H. and O. Roy, *Dynamic logic of legal competences*, Journal of Logic, Language and Information **30** (2021), pp. 701–724.

[10] Gilbert, M., *Shared intention and personal intentions*, Philosophical studies **144** (2009), pp. 167–187.

[11] Hilpinen, R. and P. McNamara, *Deontic logic: A historical survey and introduction*, Handbook of Deontic Logic and Normative Systems **1** (2013), pp. 3–136.

[12] Jones, A. J. I. and M. Sergot, *A Formal Characterisation of Institutionalised Power*, Logic Journal of the IGPL **4** (1996), pp. 427–443.

[13] Kanger, S. and H. Kanger, *Rights and parliamentarism*, Theoria **32** (1966), pp. 85–115.

[14] Li, X., D. Gabbay and R. Markovich, *Dynamic Deontic Logic for Permitted Announcements*, in: *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning*, 2022, pp. 226–235.

[15] Lindahl, L., "Position and change: A study in law and logic," Synthese Library, Springer Dordrecht, 1977, 1 edition.

[16] Makinson, D., *On the formal representation of rights relations*, Journal of Philosophical Logic **15** (1986), pp. 403–425.

[17] Markovich, R., *Understanding Hohfeld and Formalizing Legal Rights: the Hohfeldian Conceptions and Their Conditional Consequences*, Studia Logica **108** (2020).
URL `onlinefirst:https://doi.org/10.1007/s11225-019-09870-5`

[18] Markovich, R. and O. Roy, *Cause of action and the right to know. a formal conceptual analysis of the texas senate bill 25 case*, in: *Legal Knowledge and Information Systems - JURIX 2021*, Frontiers in Artificial Intelligence and Applications **346**, IOS Press, 2021 pp. 217–224.

[19] Markovich, R. and O. Roy, *Formalizing the right to know: Epistemic rights as normative positions*, in: *The First International Workshop on Logics for New-Generation Artificial Intelligence (LNGAI 2021)*, 2021, pp. 154–159.

[20] Meyer, J.-J. C. et al., *A different approach to deontic logic: deontic logic viewed as a variant of dynamic logic.*, Notre Dame J. Formal Log. **29** (1988), pp. 109–136.

[21] Prakken, H. and M. Sergot, *Contrary-to-duty obligations*, Studia Logica **57** (1996), pp. 91–115.

[22] Rao, A. S. and M. P. Georgeff, *Modeling rational agents within a bdi-architecture*, Readings in agents , pp. 317–328.

[23] Reynolds, M., *An axiomatization of full computation tree logic*, The Journal of Symbolic Logic **66** (2001), pp. 1011–1057.

[24] Sergot, M., *Normative Positions*, in: D. Gabbay, J. Horty, X. Parent, R. van der Meyden and L. van der Torre, editors, *Handbook of Deontic Logic and Normative Systems*, College Publications, 2013 pp. 353–406.

[25] Sergot, M. and R. Craven, *The deontic component of action language nc+*, in: *Deontic Logic and Artificial Normative Systems: 8th International Workshop on Deontic Logic in Computer Science, DEON 2006, Utrecht, The Netherlands, July 12-14, 2006. Proceedings 8*, Springer, 2006, pp. 222–237.

[26] Shoham, Y., *Logical theories of intention and the database perspective*, Journal of Philosophical Logic **38** (2009).

[27] Thomason, R. H., *Practical reasoning: Problems and prospects*, in: X. P. R. v. d. M. Dov Gabbay, John Horty and L. van der Torre, editors, *Handbook of Deontic Logic Vol. 2*, College Publications, 2021 pp. 463–498.

[28] van Zee, M., D. Doder, L. van der Torre, M. Dastani, T. Icard and E. Pacuit, *Intention as commitment toward time*, Artificial Intelligence **283** (2020), p. 103270.
URL `https://www.sciencedirect.com/science/article/pii/S0004370220300308`

[29] Watson, L., *The right to know: Epistemic rights, and why we need them*, Australasian Journal of Philosophy **100** (2022), pp. 426–427.