UvA-DARE (Digital Academic Repository)

A Dynamic Approach to Temporal Normative Logic

Ju, F.; Grilletti, G.

DOI

[10.1007/978-3-662-55665-8_35](https://doi.org/10.1007/978-3-662-55665-8_35)

Publication date

2017

Document Version

Submitted manuscript

Published in

Logic, Rationality, and Interaction

[Link to publication](#)

Citation for published version (APA):

Ju, F., & Grilletti, G. (2017). A Dynamic Approach to Temporal Normative Logic. In A. Baltag, J. Seligman, & T. Yamada (Eds.), *Logic, Rationality, and Interaction: 6th International Workshop, LORI 2017, Sapporo, Japan, September 11-14, 2017 : proceedings* (pp. 512-525). (Lecture Notes in Computer Science; Vol. 10455), (FoLLI Publications on Logic, Language and Information). Springer. https://doi.org/10.1007/978-3-662-55665-8_35

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

A Dynamic Approach to Temporal Normative Logic

Fengkui Ju^{1,2} and Gianluca Grilletti^{3,4}

¹ School of Philosophy, Beijing Normal University

² `fengkui.ju@bnu.edu.cn`

³ Institute for Logic, Language & Computation, University of Amsterdam

⁴ `grilletti.gianluca@gmail.com`

Abstract. State commands refer to states, not actions. They have a temporal dimension explicitly or implicitly. They indirectly change what we are permitted, forbidden or obligated to do. This paper presents DTNL, a deontic logic meant to handle state commands based on the branching-time temporal logic PCTL*. The models of DTNL are trees with *bad states*, which are identified by a propositional constant *b* introduced in the language. To model state commands, a dynamic operator that adds states to the extension of *b* is introduced.

Keywords: deontic logic, temporal logic, commands, deadlines

1 Background

There are two types of commands that refer to *actions* and *states* respectively. The former can be called action commands and the latter state ones. For example *never touch the button* is an action command since it imposes a restriction on which actions can be legally performed, while *ensure that the table is clean before the meeting*; *everything be in order until I get back* [12]; *nobody sit in the first row* [12] are state commands as they impose conditions on the future state of affairs.

Commands of both types can change what the agent is *permitted*, *forbidden* or *obligated* to do. After the action command *never touch the button*, the agent is not allowed to touch the button any more. After the state command *ensure that the table is clean*, the agent has the obligation to make it true that the table is clean.

Unlike action commands, state commands change what the agent can do indirectly. The state command *ensure that the table is clean* is not ordering the agent to clean the table himself. All it requires is that the agent makes it the case that the table is clean. The agent can execute the command by letting someone else clean the table.

Commands have a temporal dimension. *Ensure that the table is clean before the meeting* imposes a deadline conditional on future events: the command is fulfilled if the table is clean just before the meeting; *everything be in order until I get back* imposes a certain obligation *until* some other condition is fulfilled;

nobody sit in the first row imposes a permanent condition: the agent should prevent *from now on* that someone sits in the first row.

There are two perspectives concerning how the agent should behave after a command: the perspective of the agent himself and the external perspective. They make a difference for the commands without a *deadline*. We take *ensure that the table is clean some day* as an example. From the external perspective, after the command, the agent should make it true that the table is clean, although he can do it in whichever day he wants. But from the perspective of the agent, the command creates no obligation, as it will not be violated at any point in time. This difference was already pointed out by [6] in the context of *deadlines of norms*.

Computation Tree Logic (CTL), proposed in [5], is a branching-time temporal logic broadly used in computer science. Examples of properties expressible by CTL are *it is sure that ϕ will happen in the next moment*; *it is sure that ϕ will eventually happen*; *it is possible that ϕ will hold until ψ* . By generalizing CTL, [1] presents a deontic logic, called Normative Temporal Logic (NTL). The models of this logic are transition systems plus illegal transitions. An example of properties expressible in NTL is *the agent is allowed to act to make ϕ true in the next moment*. Compared with conventional deontic logics, NTL has two interesting features. Firstly, it makes the idea explicit that the agent acts to make things true. Secondly, normative notions expressible by it have a temporal dimension.

NTL is conceptually suitable to handle state commands, but technically has two problems. Firstly, NTL, as CTL, has a syntactic restriction on applications of temporal operators: they have to be immediately preceded by path quantifiers. This implies that the temporal formulas such as *ϕ will eventually happen* are not well-formed formulas and iterations of temporal operators are not allowed. So it is hard to express and interpret state commands within NTL. Secondly, NTL uses transition systems with illegal transitions to handle normative notions, but these models are *historyless* in the sense that whether a transition is illegal or not has nothing to do with what the agent has done in the past. However, normative notions with a temporal dimension essentially involve past actions. For example, assume a scenario in which a child has to collect 100 coins in a piggy bank, and only then retrieve them by crashing the container. As the difference made by putting a coin into the bank can not be seen, the child has to rely on his memory of the past actions to know when to break the bank.

Full Computation Tree Logic (CTL*), introduced in [7], is an extension of CTL that does not have the syntactic restriction mentioned previously. PCTL* is a further extension of CTL* with two past operators whose completeness is shown in [9]. In what follows we present DTNL (“Dynamic Temporal Normative Logic”), a deontic logic based on PCTL*. This logic takes *trees with bad states* as its models, instead of general transition systems with *bad transitions*. A special propositional constant **b** is introduced to indicate bad states. Using this constant, the normative notions of permission, prohibition and obligation can be defined as in [8] and [3]. A dynamic operator representing state commands is also intro-

duced. Its function is to update the model, adding those states that violate the command to the extension of $\mathbf{b}$. The logic follows the agent's perspective, and commands without a deadline might not create the corresponding obligations. Our way of viewing commands follows [10] and [2] which think that the meaning of commands lies in how they change the agent's internal state.

2 Language

Let Φ_0 be a countable set of atomic propositions and p range over Φ_0 . Define the language Φ_{DTNL} as the following:

$$\phi ::= p \mid \top \mid \mathbf{b} \mid \neg\phi \mid (\phi \wedge \phi) \mid Y\phi \mid (\phi S\phi) \mid X\phi \mid (\phi U\phi) \mid \mathbf{A}\phi \mid [!\phi]\phi$$

The featured formulas of this language are read as follows:

1. $\mathbf{b}$: this is a *bad* state.
2. $Y\phi$: ϕ was the case in the *last* moment.
3. $\phi S\psi$: ϕ has been the case *since* ψ .
4. $X\phi$: ϕ will be the case in the *next* moment.
5. $\phi U\psi$: ϕ will be the case *until* ψ .
6. $\mathbf{A}\phi$: no matter how the agent will act in the future, ϕ is the case *now*.
7. $[!\phi]\psi$: ψ is the case after the command *make ϕ true* is given.

Note that the language Φ_{DTNL} is the language of PCTL* plus the propositional constant $\mathbf{b}$ and the dynamic operator $[!\phi]$.

It seems strange to say that no matter how the agent will act in the future, ϕ is the case now, but in fact this is fine. Whether a sentence that involves time relations is true or not now might be dependent on how the agent will act in the future. For example, whether a student will pass an exam is dependent on how he will study. In order to make a sentence true now, the agent has to act in a certain way in the future.

The other usual propositional connectives and the falsum $\perp$ are defined in the usual way:

1. $\mathbf{f} := \neg\mathbf{b}$: this is a *fine* state.
2. $P\phi := (\top S\phi)$: ϕ was the case.
3. $H\phi := \neg P\neg\phi$: ϕ has been the case.
4. $F\phi := (\top U\phi)$: ϕ will be the case.
5. $G\phi := \neg F\neg\phi$: ϕ will always be the case.
6. $\mathbf{E}\phi := \neg\mathbf{A}\neg\phi$: the agent has a way to act in the future s.t. ϕ is the case now;
 $\mathbf{E}\phi$ intuitively means that making ϕ true is *achievable*.

Let $X^n\phi$ denote $X \dots X\phi$ where n is the number of occurrences of X . The state commands mentioned in the beginning can be expressed as follows:

1. *Ensure that the table is clean before the meeting*: $!X^{k-1}c$

2. *Everything be in order until I get back*: $!(oUb)$
3. *Nobody sit in the first row*: $!G \neg s$

For the first example, we suppose that the starting time of the meeting is already fixed and there are k units of time from now to it.

3 Models

Let W be a nonempty set of states and R a binary relation on it. A sequence $w_0 \dots w_n$ of states (possibly of length one) is called an *R-sequence* if $w_0 R \dots R w_n$. (W, R) is a *tree* if there is a $r \in W$, called the *root of the tree* s.t. for any w , there is a unique *R*-sequence from r to w . Immediate consequences of the definition are that the root is unique and R is irreflexive. R is *serial* if for any w , there is a u s.t. Rwu (there are no end points).

A serial tree (W, R) is understood as a time structure encoding an agent's actions (the transitions) and states in time (the nodes). At any state w , the *history* of the agent up to that point is represented by the path connecting the root to w (the actions performed). The seriality condition corresponds to the fact that the agent can always perform an action at any given time, while a branching in the tree is interpreted as a situation in which the agent *can choose between different possible actions*.

Fix a serial tree (W, R) . Here are some auxiliary notations. An *R*-sequence $w_0 \dots w_n$ starting from the root is an *history* of w_n . For any w and u , u is a *historical* state of w if there is an *R*-sequence $u_0 \dots u_n$ s.t. $0 < n$, $u_0 = u$ and $u_n = w$. w is a *future* state of u if u is a historical state of w . Note that a state can not be a historical or future state of itself.

An infinite *R*-sequence is a *path*. A path starting at the root is a *timeline*. A path $w_0 \dots$ *passes by* a state x if $x = w_i$ for some $i > 0$. Let π be a path. We use $\pi(i)$ to denote the $i+1$ -th element of π , ${}^i\pi$ the prefix of π to the $i+1$ -th element and π^i the suffix of π from the $i+1$ -th element. For example, if $\pi = w_0 \dots$, then $\pi(2) = w_2$, ${}^2\pi = w_0 w_1 w_2$ and $\pi^2 = w_2 \dots$. For any history $w_0 \dots w_n$ and path $u_0 \dots$ s.t. $w_n = u_0$, let $w_0 \dots w_n \otimes u_0 \dots$ denote the timeline $w_0 \dots w_n u_1 \dots$.

A tuple $\mathfrak{M} = (W, R, r, B, V)$ is a *model* if

1. (W, R) is a serial tree with r as the root
2. B is a subset of W meeting the following conditions:
 - (a) if $w \in B$, then $u \in B$ for any u s.t. Rwu
 - (b) if $w \notin B$, then there is a u s.t. Rwu and $u \notin B$
3. V is a function from Φ_0 to 2^W

B is called the set of *bad* states and $W - B$ the set of *fine* ones. Intuitively, a transition (w, u) is *illegal* if u is a bad state. The first constraint on B is called *persistence of liability*; it indicates that if we reached a state in which we failed to fulfill a command, then this holds for all its successors. This constraint implies that if a state is bad, then all of its future states are bad, and if a state is fine, then all of its historical states are fine. The second constraint on B is called

seriality of legality; it means that if a state is fine, then at least a successor of it is fine. The conjunction of the two constraints is called *normative coherence*. **Figure 1** illustrates a model.

Fig. 1. This figure indicates a model. w_0 is the root. Dotted circles denote bad states and solid circles fine states. Arrows are transitions.

The following is an intuitive interpretation of the models: an agent's possible actions are encoded by a serial tree (as mentioned above). At any moment, the agent has a set of rules he should respect, and these rules are encoded by the bad states: the agent is allowed to travel to the fine states but not to the bad ones. If the agent has not done anything illegal, there is always something legal for him to do, and if he has done something illegal, there is nothing legal for him to do. We will get back to the last point again in **Section 6**.

4 Semantics

Following the dynamic approach, we define by mutual recursion the truth of formulas with respect to a model and an update operation to interpret the dynamic operator $!$.

$\mathfrak{M}, \pi, i \Vdash \phi$, the formula ϕ being true at the state $\pi(i)$ relative to the timeline π in the model $\mathfrak{M}$, is defined as follows:

$\mathfrak{M}, \pi, i \Vdash p$	$\Leftrightarrow$	$\pi(i) \in V(p)$
$\mathfrak{M}, \pi, i \Vdash \top$		
$\mathfrak{M}, \pi, i \Vdash \mathbf{b}$	$\Leftrightarrow$	$\pi(i) \in B$
$\mathfrak{M}, \pi, i \Vdash \neg\phi$	$\Leftrightarrow$	not $\mathfrak{M}, \pi, i \Vdash \phi$
$\mathfrak{M}, \pi, i \Vdash \phi \wedge \psi$	$\Leftrightarrow$	$\mathfrak{M}, \pi, i \Vdash \phi$ and $\mathfrak{M}, \pi, i \Vdash \psi$
$\mathfrak{M}, \pi, i \Vdash Y\phi$	$\Leftrightarrow$	$i > 0$ and $\mathfrak{M}, \pi, i-1 \Vdash \phi$
$\mathfrak{M}, \pi, i \Vdash \phi S \psi$	$\Leftrightarrow$	there is a $j \leq i$ s.t. $\mathfrak{M}, \pi, i-j \Vdash \psi$ and $\mathfrak{M}, \pi, i-k \Vdash \phi$ for any $k < j$
$\mathfrak{M}, \pi, i \Vdash X\phi$	$\Leftrightarrow$	$\mathfrak{M}, \pi, i+1 \Vdash \phi$
$\mathfrak{M}, \pi, i \Vdash \phi U \psi$	$\Leftrightarrow$	there is a j s.t. $\mathfrak{M}, \pi, i+j \Vdash \psi$

$$\begin{array}{ll}
& \text{and } \mathfrak{M}, \pi, i + k \Vdash \phi \text{ for any } k < j \\
\mathfrak{M}, \pi, i \Vdash \mathbf{A}\phi & \Leftrightarrow \text{for any path } \rho \text{ starting at } \pi(i), \mathfrak{M}, {}^i\pi \otimes \rho, i \Vdash \phi \\
\mathfrak{M}, \pi, i \Vdash [!\phi]\psi & \Leftrightarrow \mathfrak{M}_{\pi(i)}^{!\phi}, \pi, i \Vdash \psi
\end{array}$$

$\mathbf{A}$ is a universal quantifier over possible timelines. The evaluation of some formulas does not depend on the whole path but only on the point of the path with the selected index (e.g., p , $\top$, $\mathbf{b}$ and $\mathbf{A}\phi$). We will call a formula for which this property holds in every model a *state formula*, while a formula whose semantical interpretation depends on other points of the path (e.g., $\mathbf{Y}\phi$, $\phi\mathbf{S}\psi$, $\mathbf{X}\phi$ and $\phi\mathbf{U}\psi$) will be called a *temporal formula*. In particular, if ψ is a state formula then $[!\phi]\psi$ is a state formula too.

A path $w_0 \dots$ is *legal* if w_i is a fine state for every $i > 0$. The update model $\mathfrak{M}_{\pi(i)}^{!\phi}$ is defined as follows.

Definition 1 (Update with commands). Let $\mathfrak{M} = (W, R, r, B, V)$ be a model, ϕ a formula and w a state. Let $w_0 \dots w_i$ be the history of w . Define a set $X_w^{!\phi}$ of states as follows: for any $x \in W$, $x \in X_w^{!\phi} \Leftrightarrow$ (i) x is a future state of w , (ii) x is a fine state, and (iii) there is no legal path ρ starting at w and passing by x s.t. $\mathfrak{M}, w_0 \dots w_i \otimes \rho, i \Vdash \phi$. Let $\mathfrak{M}_w^{!\phi} = (W, R, r, B \cup X_w^{!\phi}, V)$ if $B \cup X_w^{!\phi}$ is normatively coherent, or else $\mathfrak{M}_w^{!\phi} = \mathfrak{M}$. $\mathfrak{M}_w^{!\phi}$ is called the result of updating $\mathfrak{M}$ at w with the command $!\phi$.

Proposition 1. Fix $\mathfrak{M} = (W, R, r, B, V)$, π, i and ϕ . $B \cup X_{\pi(i)}^{!\phi}$ is not normatively coherent $\Leftrightarrow \pi(i)$ is a fine state in $\mathfrak{M}$ and there is no legal path ρ starting at $\pi(i)$ s.t. $\mathfrak{M}, {}^i\pi \otimes \rho, i \Vdash \phi$.

Proof. ($\Leftarrow$) Assume that $\pi(i)$ is as in the hypothesis. Since $\pi(i)$ is not reachable from $\pi(i)$, by definition of $X_{\pi(i)}^{!\phi}$ it follows that $\pi(i) \notin B \cup X_{\pi(i)}^{!\phi}$.

Consider now a successor u of $\pi(i)$. If $u \notin B$, by hypothesis we know that there is no legal path ρ starting at $\pi(i)$ s.t. $\mathfrak{M}, {}^i\pi \otimes \rho, i \Vdash \phi$ and passing by u . By definition of $X_{\pi(i)}^{!\phi}$ this implies $u \in X_{\pi(i)}^{!\phi}$, and because u was arbitrary, it follows that all the successors of $\pi(i)$ lie in $B \cup X_{\pi(i)}^{!\phi}$, and so the set is not normatively coherent.

($\Rightarrow$) The result is proved by contraposition. We want to show that $B \cup X_{\pi(i)}^{!\phi}$ is normatively coherent assuming that $\pi(i)$ is a bad state or that there is a legal path ρ starting at $\pi(i)$ s.t. $\mathfrak{M}, {}^i\pi \otimes \rho, i \Vdash \phi$.

The former case: assume that $\pi(i)$ is a bad state in $\mathfrak{M}$. Then $X_{\pi(i)}^{!\phi} = \emptyset$ and $B \cup X_{\pi(i)}^{!\phi} = B$. Then $B \cup X_{\pi(i)}^{!\phi}$ is trivially normatively coherent.

The latter case: assume that there is a legal path ρ as described above. Spelling out the definition of normative coherence, we have to show that the successors of a state in $B \cup X_{\pi(i)}^{!\phi}$ are again in $B \cup X_{\pi(i)}^{!\phi}$ (persistence of liability) and that given a state $w \notin B \cup X_{\pi(i)}^{!\phi}$ then it has a successor u s.t. $u \notin B \cup X_{\pi(i)}^{!\phi}$ (seriality of legality).

The first condition it's easily checked. If $w \in B$ then all its successors are in B because we assumed that $\mathcal{M}$ is normative coherent. Otherwise if $w \in X_{\pi(i)}^{! \phi}$, consider u s.t. Rwu ; clearly we have

$$\begin{aligned} w \in X_{\pi(i)}^{! \phi} &\Rightarrow R\pi(i)w \text{ and for all } \rho \text{ starting at } \pi(i) \\ &\quad \text{and passing by } w: \mathcal{M}, {}^i\pi \otimes \rho, i \not\models \phi \\ &\Rightarrow R\pi(i)u \text{ and for all } \rho \text{ starting at } \pi(i) \\ &\quad \text{and passing by } u: \mathcal{M}, {}^i\pi \otimes \rho, i \not\models \phi \\ &\Rightarrow u \in B \text{ or } u \in X_{\pi(i)}^{! \phi} \\ &\Rightarrow u \in B \cup X_{\pi(i)}^{! \phi} \end{aligned}$$

For the second condition, consider $w \notin B \cup X_{\pi(i)}^{! \phi}$ and suppose that $R\pi(i)w$ (otherwise the result trivially follows). We have, by definition of $X_{\pi(i)}^{! \phi}$, that there exists a *legal* path ρ starting at $\pi(i)$ passing by w such that $\mathcal{M}, {}^i\pi \otimes \rho, i \models \phi$. And now it easily follows that for any successor u of w in the path ρ we have $u \notin B$ (as ρ is legal) and $u \notin X_{\pi(i)}^{! \phi}$ (as the path $\pi(i) \otimes \rho$ witness this). $\square$

The proof of this proposition is omitted due to limit of space. That there is no legal path ρ starting at $\pi(i)$ s.t. $\mathfrak{M}, \pi \otimes \rho, i \models \phi$ means that making ϕ true at $\pi(i)$ is forbidden. $X_w^{! \phi}$ is understood as follows. Assume that the agent is at the state w and the command *make ϕ true* is given to him. If w is a fine state but it is not allowed to make ϕ true, then the agent considers the command strange and ignores it. Assume otherwise. Then the agent scans the fine states that he can reach from w one by one. He marks a state bad if he finds this: *if he travels to it, there would be no legal way to make ϕ true at w , no matter where he goes afterwards*. $X_w^{! \phi}$ is the collection of the states that he marks bad. After marking, the agent behaves by taking the new bad states into consideration. **Figure 2** illustrates how a command updates a model.

Note that the set $X_w^{! \phi}$ is defined w.r.t. $\mathfrak{M}$. Updating $\mathfrak{M}$ at w with $! \phi$ only changes the future states of w .

A formula ϕ is *valid* if for any $\mathfrak{M}, \pi$ and i , $\mathfrak{M}, \pi, i \models \phi$. Let Γ be a set of formulas and ϕ a formula. $\Gamma \models \phi$, Γ *entails* ϕ , if for any $\mathfrak{M}, \pi$ and i , if $\mathfrak{M}, \pi, i \models \Gamma$, then $\mathfrak{M}, \pi, i \models \phi$. We in the sequel use DTNL to denote the set of valid formulas.

5 A static deontic logic by reduction

Without the dynamic operator, the static part of DTNL is just PCTL* plus the propositional constant **b**. As mentioned in the introduction, there is already a complete axiomatization of PCTL* in the literature. We can get a complete axiomatization for the static part of DTNL by adding the axioms **b** $\rightarrow$ **AXb** and **f** $\rightarrow$ **EXf** to PCTL*. The two formulas respectively express the two constraints on the models of DTNL: persistency of liability and seriality of legality.

The formula XGf indicates that from the next moment on, the state will always be fine. It expresses legal paths in the following sense: for any model $\mathfrak{M}$

Fig. 2. This figure illustrates how a model is updated by a command. The valuation of two propositions p and q is depicted. $\mathfrak{M}_{w_1}^{!Xp}$ is the result of updating $\mathfrak{M}$ at w_1 with the command *make p true in the next moment* ($!Xp$).

and history $w_0 \dots w_i$, a path ρ is legal iff $\mathfrak{M}, w_0 \dots w_i \otimes \rho, i \Vdash \text{XGf}$. Using this formula and the path quantifiers **A** and **E**, we can define some deontic notions.

1. $\mathcal{P}\phi := \mathbf{E}(\text{XGf} \wedge \phi)$: the agent has a legal way to act in the future s.t. ϕ is the case now. That is, the agent is *permitted* to make ϕ true now.
2. $\mathcal{F}\phi := \mathbf{A}(\text{XGf} \rightarrow \neg\phi)$: no matter how the agent will legally act in the future, ϕ is not the case now. That is, the agent is *forbidden* to make ϕ true now.
3. $\mathcal{O}\phi := \mathbf{A}(\text{XGf} \rightarrow \phi)$: no matter how the agent will legally act in the future, ϕ is the case now. That is, the agent is *obligated* to make ϕ true now.

The truth conditions of the normative formulas can be specified by use of legal paths.

$$\begin{aligned} \mathfrak{M}, \pi, i \Vdash \mathcal{P}\phi &\Leftrightarrow \text{there is a legal path } \rho \text{ from } \pi(i) \text{ s.t. } \mathfrak{M}, {}^i\pi \otimes \rho, i \Vdash \phi \\ \mathfrak{M}, \pi, i \Vdash \mathcal{F}\phi &\Leftrightarrow \text{there is no legal path } \rho \text{ from } \pi(i) \text{ s.t. } \mathfrak{M}, {}^i\pi \otimes \rho, i \Vdash \phi \\ \mathfrak{M}, \pi, i \Vdash \mathcal{O}\phi &\Leftrightarrow \text{for any legal path } \rho \text{ from } \pi(i), \mathfrak{M}, {}^i\pi \otimes \rho, i \Vdash \phi \end{aligned}$$

The deontic operators can be treated as quantifiers over legal paths.

We obtained a deontic logic, namely the static part of **DTNL**, for which normative formulas have a temporal dimension. For example, $\mathcal{O}Fp$ says that the agent ought to make p true at some point in the future.

A quite different case is $\mathcal{O}p$. This formula does not mean that the agent ought to make p true, but that the agent ought to *act in the future to make p true at the present moment*. As the condition is independent from the future events, this should count as a trivial obligation, and in fact it can be verified that $\text{f} \rightarrow (p \leftrightarrow \mathcal{O}p)$ is a valid formula.

As defined above, a path is legal if it consists only of legal transitions, but there are other properties of paths expressible in **DTNL** that are interesting in normative contexts. One of them is *containing finitely many illegal transitions*,

which is expressed by $\text{FG}\mathfrak{f}$. This property can be intuitively understood as *mainly* consisting of legal transitions and is *slightly* weaker than legal paths. Using this property, we can define different normative notions in a similar fashion as the one presented above. This issue deserves a closer look.

6 Explanations

The models of DTNL are trees with bad states. Note that in trees, defining bad states or bad transitions results in equivalent semantics and there is no real difference between the two approaches. We mentioned in **Section 1** that bad transitions in general transition systems are *historyless* and not suitable to handle normative notions with a temporal dimension. A concrete example is given in **Figure 3**.

Fig. 3. The structure in this figure represents a conceivable scenario. Let $\phi = \mathcal{O}(X\neg p \wedge \text{XXG}p)$. This formula says what follows is obligatory: firstly make p false; then make p true; then keep p true forever. In fact, it can be checked that there is no way to arrange illegal transitions in this structure s.t. ϕ is true at w . This shows the importance of keeping track of the past actions.

Note that technically, trees with bad states are equivalent to general transition systems with *history-dependent* bad transitions. By history-dependent, we mean whether a transition is bad or not is dependent on specific histories, that is, finite transition sequences. Roughly, the arguments for the equivalence can go as follows: *point generated submodels* preserve truth; pointed transition systems with history-dependent bad transitions can be safely *unwound* to trees with bad states. We refer to [4] for the details of generated submodels and unwinding. However, if we work with transition systems with history-dependent bad transitions, the definitions, especially the definition for the update with commands, would be very complicated. Working with trees plus bad states makes things much easier.

The special constraint persistency of liability on the models of DTNL denotes that if a state is bad, then all of its successors are bad as well. We use this constraint for two reasons. One reason concerns offering state permissions such as *you may let the prisoners go today or tomorrow*. Offering permissions tends to make bad states fine. This paper focuses on giving commands and does not deal with offering permissions, but once we want to handle it in this framework,

it becomes clear why the constraint is needed. Another reason is conceptual. If a command has been violated, then this fact will remain true also in the future. The constraint *persistence of liability* is coherent with this.

7 Commands without a deadline

The formula $[\!|\phi|\!]O\phi$ is not generally valid. For example, $[\!|Fp|\!]OFp$ is not valid, as shown in **Figure 4**. Therefore, commands do not always cause the expected effects, sometimes they *fail*. Note that the command $!Fp$ does not have a *deadline* and so it can not be checked if it has been violated after a finite amount of time. Actually, this failure only happens for this kind of commands.

Fig. 4. This figure explains why $[\!|Fp|\!]OFp$ is not valid. The model $\mathfrak{M}$ does not contain any bad state and all the paths in it are legal. The command $!Fp$ at w_0 in $\mathfrak{M}$ does not change anything, as no matter where the agent goes from w_0 , there is a way to make Fp true at w_0 . Therefore, $w_0w_1\dots$ is a legal path of $\mathfrak{M}_{w_0}^{!Fp}$. However, it can be seen that not $\mathfrak{M}_{w_0}^{!Fp}, w_0w_1\dots, 0 \Vdash Fp$. Then not $\mathfrak{M}_{w_0}^{!Fp}, w_0w_1\dots, 0 \Vdash OFp$, that is, not $\mathfrak{M}, w_0w_1\dots, 0 \Vdash [\!|Fp|\!]OFp$.

We say that a formula ϕ is *colorless* if it contains no occurrence of $\mathfrak{b}$. Colorless formulas are not sensitive to badness of states.

Proposition 2. *Let ϕ be a colorless formula. Then $\mathfrak{M}, \pi, i \Vdash [\!|\phi|\!]O\phi \Leftrightarrow$ for any path ρ starting at $\pi(i)$, if $\mathfrak{M}, {}^i\pi \otimes \rho, i \Vdash \{X^nEY^n(XG\mathfrak{f} \wedge \phi) \mid n \in \mathbb{N}\}$, then $\mathfrak{M}, {}^i\pi \otimes \rho, i \Vdash \phi$.*

Proof. ($\Leftarrow$) By contraposition, assume $\mathfrak{M}, \pi, i \not\Vdash [\!|\phi|\!]O\phi$. By spelling out the semantic clauses, this means that there exists a legal path ρ of $\mathfrak{M}_{\pi(i)}^{! \phi}$ starting at $\pi(i)$ s.t. $\mathfrak{M}_{\pi(i)}^{! \phi}, {}^i\pi \otimes \rho, i \not\Vdash \phi$. It is straightforward that ρ is also a legal path of $\mathfrak{M}$ as well. Moreover, as ϕ is colorless, $\mathfrak{M}, {}^i\pi \otimes \rho, i \not\Vdash \phi$.

It remains to show that $\mathfrak{M}, {}^i\pi \otimes \rho, i \Vdash X^nEY^n(XG\mathfrak{f} \wedge \phi)$ for every value of $n \in \mathbb{N}$. Fix the notation $\rho = u_0u_1\dots$ and $j \geq 0$. As u_j is a fine state in

$\mathfrak{M}_{\pi(i)}^{! \phi}$, there exists a legal path τ in $\mathfrak{M}$ starting at $\pi(i)$ and passing by u_j s.t. $\mathfrak{M}, i\pi \otimes \tau, i \Vdash \phi$ (notice that τ and ρ coincide up to the n -th state because our structure is a *tree*). Then clearly $\mathfrak{M}, i\pi \otimes \tau, i \Vdash \text{XGf} \wedge \phi$ and consequently $\mathfrak{M}, i\pi \otimes \rho, i \Vdash X^j \text{EY}^j (\text{XGf} \wedge \phi)$. Since j was arbitrary, it follows $\mathfrak{M}, i\pi \otimes \rho, i \Vdash \{X^n \text{EY}^n (\text{XGf} \wedge \phi) \mid n > 0\}$, as wanted.

($\Rightarrow$) By contraposition, assume there is a path ρ starting at $\pi(i)$ s.t. $\mathfrak{M}, i\pi \otimes \rho, i \Vdash \{X^n \text{EY}^n (\text{XGf} \wedge \phi) \mid n \in \mathbb{N}\}$ but $\mathfrak{M}, i\pi \otimes \rho, i \not\Vdash \phi$. Note that, as ϕ is colorless, this is equivalent to $\mathfrak{M}_{\pi(i)}^{! \phi}, i\pi \otimes \rho, i \not\Vdash \phi$. So in order to show $\mathfrak{M}, \pi, i \not\Vdash [\phi] \mathcal{O}\phi$, it suffices to prove that ρ is a fine path of $\mathfrak{M}_{\pi(i)}^{! \phi}$.

Let $\rho = u_0 u_1 \dots$ and $j \geq 0$. It is easy to show, using again that ϕ is colorless, that the hypothesis $\mathfrak{M}, i\pi \otimes \rho, i \Vdash X^j \text{EY}^j (\text{XGf} \wedge \phi)$ implies that every u_j is a fine in $\mathfrak{M}$ and is not an element of $X_{\pi(i)}^{! \phi}$, and this means exactly that u_j is a fine state of $\mathfrak{M}_{\pi(i)}^{! \phi}$. Since j was arbitrary, it follows that ρ is a fine path, as wanted. $\square$

A formula ϕ is *co-compact* if $\{X^n \text{EY}^n \phi \mid n \in \mathbb{N}\} \models \phi$. An example of such a formula is $\text{G}p$, while $\text{F}p$ is not. Co-compact formulas are exactly the formulas that can be *falsified in finite steps*: if the formula is not true at a path, this is attested by an initial segment. For example, if $\text{G}p$ is false at π , then in order to know this, we just have to scan π up to a state in which p is false. Note that if $\text{G}p$ is true, we might need to scan the whole path: it is possible that the agent will never execute $! \phi$ but we are unable to attest this. In other words, if ϕ is not co-compact, then $! \phi$ is a *command without a deadline*.

Proposition 3. *Let ϕ be a colorless formula. Then $[\phi] \mathcal{O}\phi$ is valid $\Leftrightarrow \phi$ is co-compact.*

Proof. By **Proposition 2**, it suffices to show that $\{X^n \text{EY}^n (\text{XGf} \wedge \phi) \mid n \in \mathbb{N}\} \models \phi \Leftrightarrow \{X^n \text{EY}^n \phi \mid n \in \mathbb{N}\} \models \phi$. As $X^n \text{EY}^n (\text{XGf} \wedge \phi)$ implies $X^n \text{EY}^n \phi$ for any n , the direction from right to left is trivial.

For the other direction, reasoning by contraposition, assume $\{X^n \text{EY}^n \phi \mid n \in \mathbb{N}\} \not\models \phi$. Then there is $\mathfrak{M}, \pi$ and i s.t. $\mathfrak{M}, \pi, i \Vdash \{X^n \text{EY}^n \phi \mid n \in \mathbb{N}\}$ but $\mathfrak{M}, \pi, i \not\Vdash \phi$. Note that, as ϕ is colorless, we can assume that *all the states in $\mathfrak{M}$ are fine*, and so it follows trivially $\mathfrak{M}, \pi, i \Vdash \{X^n \text{EY}^n (\text{XGf} \wedge \phi) \mid n \in \mathbb{N}\}$, thus showing that $\{X^n \text{EY}^n (\text{XGf} \wedge \phi) \mid n \in \mathbb{N}\} \not\models \phi$ as wanted. $\square$

Something worth mentioning is that a command containing a co-compact formula might have *different deadlines in different situations* and so a bound to the length of the initial segment falsifying the formula can not be given. An example is $!\text{G}p$. These commands just have implicit deadlines. We say that a formula ϕ is *explicitly co-compact* if there is a n s.t. $X^n \text{EY}^n \phi \rightarrow \phi$ is valid. The commands containing this sort of formulas have explicit deadlines. For example, $!(Xp \vee \text{XX}p)$ has an explicit deadline, that is, two steps; if p is not true in two steps, then the command is violated.

It can be seen that past operators let us easily state the condition of $[!\phi]\mathcal{O}\phi$ being true and the definition of co-compact formulas. This is the reason that we introduce them.

8 Connections with other work

We mentioned in **Section 1** the deontic logic NTL which is based on CTL and presented in [1]. The language of NTL, Φ_{NTL} , is defined as follows⁵:

$$\phi ::= p \mid \top \mid \neg\phi \mid (\phi \wedge \phi) \mid \mathcal{P}X\phi \mid \mathcal{P}(\phi U \phi) \mid \mathcal{O}X\phi \mid \mathcal{O}(\phi U \phi)$$

Note that Φ_{NTL} is a fragment of Φ_{DTNL} that contains only state formulas.

A model of NTL is a structure (W, R, η, V) where W and V are as usual, R is a serial relation on W , and η is a subset of R whose complement is serial, called the *set of illegal transitions*⁶. A path $w_0 w_1 \dots$ is called *legal* if $(w_i, w_{i+1}) \notin \eta$ for any i . The semantics of Φ_{NTL} is a special case of the semantics of Φ_{DTNL} , but now $\mathcal{P}$ and $\mathcal{O}$ can be regarded as an existential and universal quantifier over legal paths respectively.

Say that a formula ϕ in Φ_{DTNL} is *f-valid* if for any $\mathfrak{M}$, π and i , if $\pi(i)$ is a fine state, then $\mathfrak{M}, \pi, i \Vdash \phi$. NTL can be embedded into DTNL under the notion of f-validity.

Let $\mathfrak{M} = (W, R, \eta, V)$ be a NTL model s.t. (W, R) is a tree. Let r be the root of (W, R) . Let $B = \{y \in W \mid (x, y) \in \eta \text{ for some } x\}$ and B' be the smallest set containing B and closed under R . It can be verified that for any w , if $w \notin B'$, then there is a u s.t. Rwu and $u \notin B'$. Define $\mathfrak{M}^\times$ as the structure (W, R, r, B', V) , which is a DTNL model.

Lemma 1. *Let $\mathfrak{M} = (W, R, \eta, V)$ be a NTL model s.t. (W, R) is a tree. Let w be a fine state of $\mathfrak{M}^\times$. Then for any $\phi \in \Phi_{\text{NTL}}$, $\mathfrak{M}, w \Vdash \phi \Leftrightarrow \mathfrak{M}^\times, w \Vdash \phi$.*

The lemma follows immediately by noticing that for any fine state u of $\mathfrak{M}^\times$, a path π starting at u is legal in $\mathfrak{M}$ iff it is legal in $\mathfrak{M}^\times$.

Proposition 4. *For any ϕ of Φ_{NTL} , ϕ is valid in NTL $\Leftrightarrow \phi$ is f-valid in DTNL.*

Proof. Assume that ϕ is not valid in NTL. Then there is an NTL model $\mathfrak{M} = (W, R, \eta, V)$ and a state w s.t. $\mathfrak{M}, w \not\Vdash \phi$. Let $\mathfrak{M}' = (W', R', \eta', V')$ be the *unwinding* of $\mathfrak{M}$ from w . It follows immediately that $\mathfrak{M}'$ is also an NTL model, that $\mathfrak{M}', w \not\Vdash \phi$ and that w is a fine state of $(\mathfrak{M}')^\times$. By **Lemma 1**, $(\mathfrak{M}')^\times, w \not\Vdash \phi$ and so ϕ is not f-valid in DTNL.

Assume that ϕ is not f-valid in DTNL. Then there is a DTNL model $\mathfrak{M} = (W, r, R, B, V)$ and a fine state w s.t. $\mathfrak{M}, w \not\Vdash \phi$. Let $\mathfrak{M}' = (W, R, \eta, V)$ where $\eta = \{(x, y) \mid y \in B\}$. Then $\mathfrak{M}'$ is an NTL model and it can be seen that $(\mathfrak{M}')^\times = \mathfrak{M}$. By **Lemma 1**, $\mathfrak{M}', w \not\Vdash \phi$ and so ϕ is not valid in NTL. $\square$

⁵ In [1], every deontic operator has a parameter referring to a specific set of illegal transitions. Here we ignore the parameter, but this is not crucial for the comparison between NTL and DTNL.

⁶ The models of NTL in [1] have initial states that are omitted here.

Standard Deontic Logic (SDL), proposed in [11], is one of the most well-known deontic logics. SDL is a typical normal modal logic, where $\Box\phi$ is interpreted as *it ought to be that ϕ* , and $\Diamond\phi$ as *it may be that ϕ* . The frames of SDL are *serial* relational structures.

As mentioned in [1], SDL is a sublogic of NTL under the translation σ defined as follows:

$$\begin{aligned} p^\sigma &= p \\ (\neg\phi)^\sigma &= \neg\phi^\sigma \\ (\phi \wedge \psi)^\sigma &= \phi^\sigma \wedge \psi^\sigma \\ (\Box\phi)^\sigma &= \mathcal{O}X\phi^\sigma \\ (\Diamond\phi)^\sigma &= \mathcal{P}X\phi^\sigma \quad (\text{derived}) \end{aligned}$$

Under this translation, the formulas of Φ_{SDL} have temporal reading: $\Box\phi$ means that *it ought to be that ϕ is true in the next moment*. Note that for any $\phi \in \Phi_{\text{SDL}}$, ϕ^σ is a state formula.

Actually, this previous fact can be stronger: SDL can be embedded to NTL. As NTL can be embedded to DTNL, SDL can be embedded to DTNL too.

9 Future work

The expressive power of DTNL needs further study. It is still not known whether the dynamic operator can be reduced: if this is the case, then the completeness of DTNL follows by the completeness of PCTL*. Another point worth mentioning is that in defining the commands with a deadline, we use the inference $\{X^n \mathbf{E}Y^n \phi \mid n \in \mathbb{N}\} \models \phi$. It's still not known if this entailment can be expressed by a formula of DTNL.

Offering state permissions is another important way to change the normative state of an agent, but it is harder to capture as it raises some interesting issues such as *free-choice* permission. Formalizing giving permissions is another future work.

The constraint *persistence of liability* on the models of DTNL says that if the agent has done something illegal, then there will be nothing legal for him to do. This implies that DTNL only works for ideal agents who always comply, so this logic is far from being realistic. One way to solve this issue is to introduce more shades, instead of a simple distinction between bad and fine states.

Acknowledgment

We would like to thank the audience of the Workshop of Language and Logic in China 2016, the Tsinghua-Bayreuth Logic Workshop, the LIRa seminar at University of Amsterdam, the seminar at Osnabrück University and the 3rd PIOTR project meeting. Thanks also go to Maria Aloni, Johan van Benthem, Ilaria Canavotto, Jan van Eijck, Davide Grossi, Frank Veltman and Yanjing Wang. Fengkui Ju was supported by the National Social Science Foundation of China (No. 12CZX053), the Fundamental Research Funds for the Central Universities

(No. SKZZY201304) and China Scholarship Council. Gianluca Grilletti was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (No. 680220).

References

1. Ågotnes, T., van der Hoek, W., Rodríguez-Aguilar, J., Sierra, C., Wooldridge, M.: A temporal logic of normative systems. In: Makinson, D., Malinowski, J., Wansing, H. (eds.) *Towards Mathematical Philosophy*, pp. 69–106. Springer Netherlands (2009)
2. Aloni, M.: Free choice, modals, and imperatives. *Natural Language Semantics* 15(1), 65–94 (2007)
3. Anderson, A.: Some nasty problems in the formal logic of ethics. *Noûs* 1(4), 345–360 (1967)
4. Blackburn, P., de Rijke, M., Venema, Y.: *Modal Logic*. Cambridge University Press (2002)
5. Clarke, E., Emerson, E.: Design and synthesis of synchronization skeletons using branching time temporal logic. In: Kozen, D. (ed.) *Proceedings of the Workshop of Logics of Programs*. pp. 52–71. Springer Berlin Heidelberg (1981)
6. Dignum, F., Broersen, J., Dignum, V., Meyer, J.J.: Meeting the Deadline: Why, When and How, pp. 30–40. Springer Berlin Heidelberg (2005)
7. Emerson, E., Halpern, J.: “Sometimes” and “not never” revisited: On branching versus linear time temporal logic. *J. ACM* 33(1), 151–178 (1986)
8. Kanger, S.: New foundations for ethical theory. In: Hilpinen, R. (ed.) *Deontic Logic: Introductory and Systematic Readings*, pp. 36–58. Springer Netherlands (1971)
9. Reynolds, M.: An axiomatization of PCTL*. *Information and Computation* 201(1), 72–119 (2005)
10. Veltman, F.: Imperatives at the borderline of semantics and pragmatics. Manuscript (2010)
11. von Wright, G.: Deontic logic. *Mind* 60(237), 1–15 (1951)
12. Zanuttini, R., Pak, M., Portner, P.: A syntactic analysis of interpretive restrictions on imperative, promissive, and exhortative subjects. *Natural Language & Linguistic Theory* 30(4), 1231–1274 (2012)