

Thou Shalt is not You Will

Guido Governatori
NICTA, Australia
guido.governatori@nicta.com.au

ABSTRACT

In this paper we discuss some reasons why temporal logic might not be suitable to model real life norms. To show this, we present a novel deontic logic contrary-to-duty/derived permission paradox based on the interaction of obligations, permissions and contrary-to-duty obligations. The paradox is inspired by real life norms.

CCS Concepts

•Theory of computation → Modal and temporal logics; •Applied computing → Law;

Keywords

Linear Temporal Logic, Compliance, Deontic Logic, Deontic Paradox

1. INTRODUCTION

The aim of this note is to discuss the reasons why temporal logic, specifically Linear Temporal Logic [12] might not be suitable to check whether the specifications of a system comply with a set of normative requirements.

The debate whether it is possible to use temporal logic for the representation of norms is not a novel one (see for example [14]), and while the argument had settled for a while, the past decade saw a resurgence of the topic with many works in the fields of normative multi-agents and business process compliance advocating temporal logic as the formalism to express normative constraints on agent behaviours and process executions. One of the reasons behind this could be the success of model checking for temporal logic in verifying large scale industry applications¹.

The problem in normative multi-agents systems and business process compliance is to determine whether the actions an agent is going to perform (encoded as a plan, corresponding to a sequence of

¹The fathers of model checking for temporal logic, i.e., Edmund Clarke, E. Allen Emerson and Joseph Sifakis, were the recipient of the Turing award in 2007 for their role in developing Model-Checking into a highly effective verification technology that is widely adopted in the hardware and software industries.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

ICAIL'2015, June 8–12, 2015, San Diego, CA, USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3522-5/15/06 ...\$15.00.

DOI: <http://dx.doi.org/10.1145/2746090.2746105>.

actions) or the tasks to be executed by a business process conform with a set of normative constraints regulating their possible (legal) behaviours. In both cases we have sequences of actions/tasks leading to sequences of states and constrains over what states and sequences of states are deemed legal according to a set of normative constraints. From a formal point of view both the behaviours and the constraints are represented by temporal logic formulas and an agent or process are compliant if the set of formulas is consistent. Temporal logic is definitely capable to model the sequences of states corresponding the behaviours of agents and processes, but the issue whether it is able to represent normative constraints (i.e., obligations and prohibition) in a conceptually sound way has been neglected. We believe that this is a crucial issue to be addressed before these techniques can be proposed for practical real life cases. Without a positive answer the work based on temporal logic for the representation of norms remains a futile formal exercise.

The short discussion above boils down to the following question:

Are normative constraints (i.e., obligations and prohibitions) regulating the behaviours of their subjects different from other types of constraints?

In case of a negative answer we have to identify what are the differences, and how to model them in temporal logic. Furthermore, we have to identify what are the issues with the resulting modelling.

Obligations and prohibitions are constraints that limit the scope of actions of the bearer subject to them. However, there is a very important difference between obligations and prohibitions and other types of constraints: violations do not result in inconsistencies. This means that they can be violated without breaking the systems in which they appear. Accordingly, a better understanding of obligations and prohibitions is that they define what is legal (in a particular system) and what is illegal. Based on this reading a violation simply indicates that we ended up in an illegal situation or state. A further aspect we have to consider, and that has been by large neglected by investigations on how to formalise and reason with deontic concepts, is that violations can be compensated for, and a situation where there is a violation but there is a compensation for the violation is still deemed legal (even if, from a legal point of view, less ideal than the situation where the violation does not occur).

The paper is organised as follows: in the next section we introduce a legal scenario (a fragment of a hypothetical privacy act) illustrating some of the aspects differentiating norms from other types of constraints, and we shortly discuss what the outcomes of cases related to this scenario should be. Then in Section 3 we briefly recall the basics of Linear Temporal Logic (LTL). In Section 4 we discuss how to formalise the scenario in LTL. We point out various shortcomings for the representation of norms in LTL, and we show that LTL captures only some of the aspects of the scenario (suggesting that it is not

able to model real life norms), or it leads to paradoxical results.²

2. LEGAL MOTIVATION

Suppose that a Privacy Act contains the following norms:³

Section 1. The collection of personal information is forbidden, unless acting on a court order authorising it.

Section 2. The destruction of illegally collected personal information before accessing it is a defence against the illegal collection of the personal information.

Section 3. The collection of medical information is forbidden, unless the entity collecting the medical information is permitted to collect personal information.

In addition the Act specifies what personal information and medical information are, and they turn out to be disjoint.

Suppose an entity, subject to the Act, collects some personal information without being permitted to do so; at the same time they collect medical information. The entity recognises that they illegally collected personal information (i.e., they collected the information without being authorised to do so by a Court Order) and decides to remediate the illegal collection by destroying the information before accessing it. Is the entity compliant with the Privacy Act above? Given that the personal information was destroyed the entity was excused from the violation of the first section (illegal collection of personal information). However, even if the entity was excused from the illegal collection, they were never entitled (i.e., permitted) to collect personal information⁴, consequently they were not permitted to collect medical information; thus the prohibition of collecting medical information was in force. Accordingly, the collection of medical information violates the norm forbidding such an activity.

Let us examine the structure of the act:

Section 1 establishes two conditions:

- i. Typically the collection of personal information is forbidden; and
- ii. The collection of personal information is permitted, if there is a court order authorizing the collection of personal information.

Section 2 can be paraphrased as follows:

- iii. The destruction of personal information collected illegally before accessing it excuses the illegal collection.

Similarly to Section 2, Section 3 states two conditions:

- iv. Typically the collection of medical information is forbidden; and
- v. The collection of medical information is permitted provided that the collection of personal information is permitted.

²Following Åqvist's [1] presentation, a paradox arises in a deontic logic Δ either when there is a formula ϕ derivable in Δ but for which the translation does not seem derivable within the natural normative language, or there is a formula ϕ which is not derivable in Δ but for which the translation seems derivable within our natural normative language.

³The Privacy Act presented here, though realistic, is a fictional one. However, (i) it is based on the novel Australian Privacy Principles (APP), Privacy Amendment (Enhancing Privacy Protection) Act 2012, and (ii) sections with the same logical structure as the clauses of this fictional act are present in the APP Act.

⁴If they were permitted to collect personal information, then the collection would have not been illegal, and they did not have to destroy it.

Based on the above discussion, if we abstract from the actual content of the norms, the structure of the act can be represented by the following set of norms (extended form):

- E1. A is forbidden.
- E2. A is permitted given C (alternatively: if C , then A is permitted).
- E3. The violation of A is compensated by B
- E4. D is forbidden.
- E5. If A is permitted, so is D .

To compensate a violation we have to have a violation the compensation compensates. Moreover, to have a violation we have to have an obligation or prohibition, the violation violates. Accordingly, it makes sense to combine E1 and E3 in a single norm, obtaining thus the following set of norms (condensed form):

- C1. A is forbidden; its violation is compensated by B .
- C2. A is permitted given C (alternatively: if C , then A is permitted).
- C3. D is forbidden.
- C4. If A is permitted, so is D .

Based on the discussion so far the logical structure of the act is (logical form):

- L1. Forbidden A ; if Forbidden A and A , then Obligatory B .
- L2. if C , then Permitted A .
- L3. Forbidden D .
- L4. If Permitted A , then Permitted D .

Notice the way we modelled the violation of the prohibition of A in L1, namely as the conjunction of A and the prohibition of A .⁵ Then we model that B is the compensation of the violation of A as an implication from the violation of A to the obligation of B .

Let us consider what are the situations compliant with the above set of norms. Clearly, if C does not hold, then we have that the prohibition of A and prohibition of D are in force. Therefore, a situation where $\neg A$, $\neg C$, and $\neg D$ hold is fully compliant (irrespective whether B holds or not). If C holds, then the permission of A derogates the prohibition of A , thus situations with either A holds or $\neg A$ holds are compliant with the first two norms; in addition, the permission of A allows us to derogate the prohibition of D . Accordingly, situations with either D or $\neg D$ comply with the third norm. Let us go back to scenarios where C does not hold, and let us suppose that we have A . This means that the prohibition of A has been violated; nevertheless the set of norms allows us to recover from such a violation by B . However, as we just remarked above to have a violation we have to have either an obligation or a prohibition that has been violated: in this case the prohibition of A . Given that the prohibition of A and the permission of A are mutually incompatible, we must have, to maintain a consistent situation, that A is not permitted. But if A was not permitted D is not permitted either; actually, according to the third norm, D is forbidden. To sum up, a scenario where $\neg C$, A , B and $\neg D$ hold is still compliant (even if to a lesser degree given the compensated violation of the prohibition of A). In any case, no situation where both $\neg C$ and D hold is compliant.

⁵Similarly, the violation of the obligation of A is the conjunction of obligation A and the negation of the content of the obligation, that is, $\neg A$.

Table 1 summarises the compliant and not compliant situations. We only report the minimal sets required to identify whether a situation is compliant or not. For non-minimal sets the outcome is determined by the union of the status for the minimal subsets.

Minimal Set	Compliance Status
C	compliant
$\neg C, A, B$	weakly compliant: compensated violation of the prohibition of A
$\neg C, A, \neg B$	not compliant: uncompensated violation of the prohibition of A
$\neg C, D$	not compliant: violation of prohibition of D
$\neg C, \neg A, \neg D$	compliant

Table 1: Compliance Status for the Privacy Act

3. LOGIC BACKGROUND

Linear Temporal Logic [12] is equipped with three unary temporal operators:

- $X\phi$: next ϕ (ϕ holds at the next time);
- $F\phi$: eventually ϕ (ϕ holds sometimes in the future); and
- $G\phi$: globally ϕ (ϕ always holds in the future).

In addition we have the following binary operators:

- $\phi U \psi$: ϕ until ψ (ϕ holds until ψ holds);
- $\phi W \psi$: ϕ weak until ψ (ϕ holds until ψ holds and ψ might not hold).

The operators above are related by the following equivalences establishing some interdefinability among them:

- $F\phi \equiv \top U \phi$,
- $G\phi \equiv \neg F\neg\phi$,
- $\phi W \psi \equiv (\phi U \psi) \vee G\phi$.

The semantics of LTL can be given in terms of transition systems. A *transition system* TS is a structure

$$TS = \langle S, R, v \rangle \quad (1)$$

where

- S is a (non empty) set of states
- $R \subseteq S \times S$ such that $\forall s \in S \exists t \in S: (s, t) \in R$
- v is a valuation function $v: S \mapsto 2^{Prop}$

where $Prop$ is the set of atomic propositions.

Formulas in LTL are evaluated against fullpaths (also called traces or runs). A *fullpath* is a sequence of states in S connected by the transition relation R . Accordingly, $\sigma = s_0, s_1, s_2 \dots$ is a fullpath if and only if $(s_i, s_{i+1}) \in R$. Given a fullpath σ , σ_i denotes the subsequence of σ starting from the i -th element, and $\sigma[i]$ denotes the i -th element of σ .

Equipped with the definitions above, the valuation conditions for the various temporal operators are:

- $TS, \sigma \models p$ ($p \in Prop$) iff $p \in v(\sigma[0])$;
- $TS, \sigma \models \neg\phi$ iff $TS, \sigma \not\models \phi$;

- $TS, \sigma \models \phi \wedge \psi$ iff $TS, \sigma \models \phi$ and $TS, \sigma \models \psi$;
- $TS, \sigma \models X\phi$ iff $TS, \sigma_1 \models \phi$;
- $TS, \sigma \models \phi U \psi$ iff $\exists k: k \geq 0, TS, \sigma_k \models \psi$ and $\forall j: 0 \leq j < k, TS, \sigma_j \models \phi$;
- $TS, \sigma \models G\phi$ iff $\forall k \geq 0, TS, \sigma_k \models \phi$;
- $TS, \sigma \models F\phi$ iff $\exists k \geq 0, TS, \sigma_k \models \phi$.

A formula ϕ is true in a fullpath σ iff it is true at the first element of the fullpath. Next we define what it means for a formula ϕ to be true in a state $s \in S$ ($TS, s \models \phi$).

$$TS, s \models \phi \text{ iff } \forall \sigma: \sigma[0] = s, TS, \sigma \models \phi. \quad (2)$$

4. SCENARIO FORMALISED

The first problem we have to address is how to model obligations and permissions in Linear Temporal Logic. When one considers the temporal lifecycle obligations, obligations can be classified as *achievement* and *maintenance* obligations [6]. After an obligation enters into force, the obligation remains in force for an interval of time. A maintenance obligation is an obligation whose content must hold for every instant in the interval in which the obligation is in force. On the other hand, for an achievement obligation, the content of the obligation has to hold at least once in the interval of validity of the obligation. Accordingly, a possible solution is to use G to model maintenance obligations⁶ and F for achievement obligations. A drawback of this proposal is that G and F are the dual of each other, i.e., $G\alpha \equiv \neg F\neg\alpha$. In Deontic Logic permission is typically defined as the lack of the obligation to the contrary and the deontic operators O and P to model obligations and permissions are defined to be the dual of each other, namely $O\alpha \equiv \neg P\neg\alpha$. In addition, most deontic logics assume the following axiom (Axiom D)⁷

$$O\alpha \rightarrow P\alpha \quad (3)$$

to ensure consistency of sets of norms. The axiom is equivalent to $O\alpha \rightarrow \neg O\neg\alpha$ meaning that if α is obligatory, then its opposite ($\neg\alpha$) is not. Prohibitions can be modelled as negative obligations, thus α is forbidden if its opposite is obligatory, that is $O\neg\alpha$. Furthermore, it has been argued that maintenance obligations are suitable to model prohibitions.

Based on the discussion above, considering that the normative constraints in the scenario of Section 2 are actually prohibitions, we formalise the scenario using G for maintenance obligations (actually prohibitions) and F for permissions. We temporarily suspend judgement whether using an operator suitable to model achievement obligations to model the dual permission for maintenance obligation is appropriate or not. All we remark here is that any formalism meant to model real life norms should account for both obligations and permissions as first class citizens.

A first possible *prima facie* formalisation of the conditions set out in the Privacy Act is:

1. $G\neg A, (G\neg A \wedge A) \rightarrow GB$;
2. $C \rightarrow FA$;
3. $G\neg D$;

⁶We can use U instead of G to capture that an obligation is in force in an interval.

⁷In terms of Kripke possible world semantics Axiom D is characterised by *seriality*, i.e., $\forall x \exists y (xRy)$, and this is the property imposed on the transition relation R over the set of states S in a transition system for LTL.

4. $FA \rightarrow FD$.

The set of formulas above exhibits some problems. First of all, in a situation where we have C we get a contradiction from 1. and 2., i.e., $G\neg A$ and FA , and then a second from 3., and 2. and 4., namely $G\neg D$ and FD . This is due to the fact that normative reasoning is defeasible. Shortly and roughly a conclusion can be asserted unless there are reasons against it. In addition, to get the expected results, we have to consider that the scenario uses strong permissions, where the permissions derogates the obligations to the contrary, or, in other terms, that the permissions are exceptions to the obligations. To accomplish this we have to specify that 2. *overrides* 1., and 4. *overrides* 3. Technically, the overrides relationship can be achieved using the following procedure:⁸

1. rewrite the formulas involved as conditionals. Thus $G\neg A$ can be rewritten as $\top \rightarrow G\neg A$.
2. add the negation of the antecedent of the overriding formulas to the antecedent of the formulas overridden formula. Accordingly $\top \rightarrow G\neg A$ is transformed into $\neg C \rightarrow G\neg A$.⁹

The second aspect we concentrate on is the form of the formulas in 1., in particular on the expression

$$(G\neg A \wedge A) \rightarrow GB. \quad (4)$$

To start with they bear resemblance to the so called *contrary-to-duty obligations*. A contrary-to-duty obligation states that an obligation/prohibition is in force when the opposite of an obligation/prohibition holds. The template for contrary-to-duty obligations is given by the pair (a) $O\alpha$ and (b) $\neg\alpha \rightarrow O\beta$. Contrary-to-duty obligations are typically problematic for deontic logic and the source of inspiration for a wealth of research in the field (see [13, 5]). The formula under scrutiny is indeed related, but there is a difference: it explicitly requires a violation, while the structure in (b) does not. In the context of the Privacy Act scenario (b) would mean that an entity has the obligation to destroy collected personal information without accessing simply because they collected it (even in the case the collection was legal, or even when they had the mandate to collect it and eventually preserve it).

Accordingly, we introduce the class of *compensatory* (contrary-to-duty) *obligations*. A compensatory obligation states that an obligation/prohibition is in force as the result of the violation of another obligation/prohibition. Thus the obligation triggered in response to the violation (secondary obligation) compensates the violation of the violated obligation (primary obligation). In other words a situation where the primary obligation is violated, but the secondary obligation is fulfilled is still deemed legal, even if it is less ideal than the case where the primary obligation is fulfilled.¹⁰ The language employed in the Privacy Act suggests that the conditions stated in Section 1 and Section 2 of the Act correspond to a case of compensatory obligation.

We turn now our attention to the issue of how to formalise compensatory obligations in LTL. The first concern we have when we look at 4 we notice that its antecedent is always false, i.e.,

⁸The focus of this paper is not how to implement defeasibility or non-monotonicity in LTL or in another monotonic logic, thus we just exemplify a possible procedure.

⁹A side-effect of this procedure, which is harmless for the purpose of this paper, is that now the combination of 3. and 4. makes FA and FD equivalent, namely $FA \equiv FD$.

¹⁰We do not exclude the case that there are situations where norms have the form of what we call compensatory obligations, but where the obligation in response to the violation does not (legally) compensate the violation.

$G\neg A \wedge A \equiv \perp$, since $G\neg A$ implies that A is false in all worlds following the world where the formula is evaluated including that world, but at the same time A is required to be true at that world. The second issue is that the compensation is assumed to be a maintenance obligation while the textual provision suggests it is an achievement obligation. We shortly discuss that achievement obligation should be represented by F , but F is used to model permissions.

To avoid the issues just discussed we introduce a new binary (temporal) operator \otimes for compensatory obligations¹¹. What we have to do for this end is to identify the conditions under which a maintenance obligation is violated. The maintenance obligation $O\alpha$ is violated if there is a instant in the interval of validity of the obligation where α does not hold, namely $\neg\alpha$ holds. The second thing is to define what it means to compensate a violation. Suppose that we are told that the violation of α is compensated by β . A natural intuition for this is that there is an instant in the interval of validity of $O\alpha$ where $\neg\alpha$ holds, and there is an instant successive to the violation where the course of action described by β holds. Based on the intuition just described LTL seems well suited to this task. Here is the evaluation condition for \otimes :^{12, 13}

$$TS, \sigma \vDash \phi \otimes \psi \text{ iff } \forall i \geq 0, TS, \sigma_i \vDash \phi; \text{ or} \\ \exists j, k : 0 \leq j \leq k, TS, \sigma_j \vDash \neg\phi \text{ and } TS, \sigma_k \vDash \psi. \quad (5)$$

We are now ready to provide the formalisation of the Privacy Act.

- N1. $\neg C \rightarrow (\neg A \otimes B)$;
- N2. $C \rightarrow FA$;
- N3. $G\neg A \rightarrow G\neg D$;
- N4. $FA \rightarrow FD$.

Transition systems can be used to model runs of systems, possible ways in which business processes can be executed, the actions of an agent or more in general the dynamic evolution of a system or the world. Norms are meant to regulate the behaviour of systems, how organisations run their business, the actions of agents and so on. So, how do we check if a particular course of actions (modelled by a transition system) complies with a set of norms (where the norms are formalised in LTL)? Simply, if the transition system is a model for the set of formulas representing the norms.

Consider a transition system $TS = \langle S, R, v \rangle$ where

1. $S = \{t_i : i \in \mathbb{N}\}$,
2. $R = \{(t_i, t_{i+1}) : i \in \mathbb{N}\}$,
3. $\neg C \in v(t_i)$ for all $i \in \mathbb{N}$, $A \in v(t_1)$, $D \in v(t_1)$ and $B \in v(t_2)$.

The transition system is such that

$$TS, t_i \vDash \neg C, \quad TS, t_1 \vDash A, \quad TS, t_1 \vDash D, \quad TS, t_2 \vDash B. \quad (6)$$

This transition system implements the scenario where at no time there is a Court Order authorising the collection of personal information

¹¹The idea of using a specific operator for compensatory (contrary-to-duty) obligations is presented in [9].

¹²Again the focus of the paper is not on how to properly model compensatory (contrary-to-duty) obligations. The operator presented here does its job in the context of the paper. For alternative definitions in the context of temporal logic or inspired by temporal logic see [11, 2]; for a semantic approach not based on temporal logic see [3].

¹³This condition implements compensatory obligations when the primary obligation is a maintenance obligation and the secondary obligation is an achievement obligation. Similar definitions can be given for other combinations of primary and secondary obligations.

($\neg C$ for all t_i), an entity collects personal information (A at time t_1) and successively destroys it (B at time t_2), and at the same time when personal information was collected medical information was collected (D at time t_1).

It is immediate to verify that the transition system TS is a model of N1–N4, namely:

$$\forall t \in S: TS, t \models N1 \wedge N2 \wedge N3 \wedge N4. \quad (7)$$

Accordingly, TS is compliant with N1–N4. However, there is state t_1 where both $\neg C$ and D hold. In Section 2 we argued that a situation where $\neg C$ and D both hold is not compliant. Therefore, we have a paradox, the formalisation indicates that the scenario is compliant, the course of actions described by the transition system does not result in a contradiction, so no illegal action is performed (or better, the collection of personal information is illegal, but its compensation, destruction of the personal information, makes full amends to it), but our legal intuition suggests that the collection of medical information in the circumstances of the scenario is illegal.¹⁴

5. CONCLUSION

The contribution of this note is twofold. First we presented a novel paradox for Deontic Logic inspired by real life norms. In particular the logical structures used in the paradox appear frequently in real life (legal) norms. The second contribution was a short analysis of how to represent norms in Linear Temporal Logic, and that the proposed formalisation results in a paradox, showing that LTL might not be suitable to model norms and legal reasoning.

We would like to point out that the discussion in the previous section just shows that a particular formalisation based on LTL is not suitable to represent the scenario, not that LTL per se is not able to represent the scenario. Indeed one could create all possible full paths in a transition system not breaching the norms, and then use the paths to synthesise the norms that regulate the transition system. However, we believe that such *ex post* analysis is useless. First humans have to perform the reasoning to determine which norms hold and when and then which paths violate the norms. In a compliance perspective, i.e., in situations where one wants to determine if the specifications of a system comply with a set of norms, this approach requires to have an oracle able to discern the compliant executions (traces) from the non-compliant ones [10], and based on the oracle analysis to remove the traces resulting in violations. This means that LTL is not used for reasoning about the norms and the transition system, and verifying whether the transition system complies with the norms. But the strength of LTL is the ability to verify specifications against transition systems. However, in such a case, given that the specifications (i.e., the formalisation of the norms) are derived from the transition systems (which have been determined to be compliant by the oracle) the verification is always positive and totally uninformative. Furthermore, we believe that the formalisation we proposed, while naive, is extremely intuitive. The major objection, as we remarked in Section 4, is that permissions are modelled using F, and we hinted that F might be suitable to model achievement obligation, and using a particular type of obligation to

¹⁴We run a pseudo empirical validation of the scenario by proposing the scenario and the Privacy Act to about a dozen legal professionals ranging from corporate legal counsellors, to high court judges to law professors. They all agree without any hesitation that the collection of medical information under the circumstances described by the scenario is illegal. However, a true validation can be only given either by a law court adjudication of a case where the norms at hand are isomorphic to the Privacy Act, or by any body with the power to give a true interpretation of an act isomorphic to the act we proposed for the scenario.

model permissions is not appropriate and counter-intuitive outcomes are to be expected. We fully agree with this objection, but if we agree that a permission is the lack of an obligation to the contrary, then F is the natural choice for permissions for prohibitions (maintenance obligations). The other issue is that if we do not use F, the issue is how to model permission, and the alternative is that LTL does not support permissions. The act we presented clearly shows that there are acts where permissions must be represented and that permissions play an important role in determining which obligations are in force and when they are in force. Hence, any formalisation excluding permission is doomed to be unable to represent the vast majority of real life legal norms. Thus, a natural question is whether branching time logic with path quantifiers such as CTL and CTL* are more apt for this task.¹⁵ In such logics permissions could be formalised by EF. While modelling permissions using path quantifiers seems a better option and provides more flexibility for modelling norms, it does not solve the problem with the scenario we proposed, given that the problem requires just a single (non) branching trace to arise, and thus path quantifiers are essentially irrelevant.

We want to remark that the paradox is not restricted to LTL or other temporal logics. It can be easily replicated in Standard Deontic Logic (and it is well known that Standard Deontic Logic is plagued with many other contrary-to-duty paradoxes). A root-cause analysis of the paradox is that a violation of a compensable obligation results in a sub-ideal state. Hence, there is a state with a violation that is still deemed legal. This means, that there is a (somehow) legal state, and if permission is evaluated as being in at least one legal state, then the violation has to be evaluated as (somehow) permitted. Part of the problem is that in such somehow legal states there might be other true legitimate permissions which are not the violation of compensable obligations. Accordingly, we conjecture, that logics using truth of a formula in at least one (somehow) legal state to determine whether something is permitted have counterparts of the paradox we presented. However, a careful analysis of existing deontic logics is needed to evaluate if they are actually affected by the paradox. Notable exceptions are the logics incorporating the non-boolean (sub-structural) \otimes operator introduced by Governatori and Rotolo [9] to model compensatory obligations. In [9] norms are represented by instances of a “normative” consequence relation; that is, a norm is represented by $\Gamma \vdash_O \alpha$, where Γ is a set of formula, and α is a formula of the form $\alpha_1 \otimes \alpha_2$. More specifically, $\Gamma \vdash_O \alpha$ means that if Γ holds, then the obligation α is in force; in case α is $\alpha_1 \otimes \alpha_2$, then, the meaning is that if Γ holds the violation of the obligation of α_1 is compensated by α_2 . Permissions are handled by a second normative consequence relation (\vdash_P) that inherits from the consequence relation for obligation, namely it is possible to derive the (permissive) norm $\Gamma \vdash_P \alpha$ from the norm $\Gamma \vdash_O \alpha$. This means that a permission can only be derived if there is a norm explicitly permitting it or from an existing obligation. However, there are no mechanisms to conclude that something is permitted because it is in a state somehow deemed as legal. The logic in [9] was developed as a (sub-structural) Gentzen system without a Tarski style semantics, The treatment of \otimes as proposed in [9] was integrated with defeasibility in [7] for reasoning with contracts, and in general, with norms in a computationally oriented fashion and used extensively and successfully in the field of business process compliance. The treatment of various permissions was thoroughly investigated in [8]. More recently, [3, 4] proposed a novel complete, sequence based, possible world semantics for the \otimes , obligation and permission operators. The semantics shows that the approach based

¹⁵For a discussion of the pros and cons of linear and branching time logics, we refer to [15], though such discussion is not on their capability of modelling norms.

on the sub-structural \otimes operator and consequence relation is not affected by the paradox presented in this papers and it is able to deal with the other deontic logic paradoxes.

ACKNOWLEDGEMENTS

I thank Antonino Rotolo and Giovanni Sartor for fruitful comments on previous drafts of this paper. I also thank the participants to NorMAS 2014 for the discussions and valuable suggestions.

NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

REFERENCES

- [1] L. Åqvist. Deontic logic. In D. M. Gabbay and F. Guentner, editors, *Handbook of Philosophical Logic. 2nd Edition*. Vol. 8, pp. 147–264. Springer, 2002.
- [2] J. van Benthem, D. Grossi, and F. Liu. Priority Structures in Deontic Logic. *Theoria*, 802:116–152, 2014.
- [3] E. Calardo, G. Governatori, and A. Rotolo. A Preference-based Semantics for CTD Reasoning. In F. Cariani, D. Grossi, J. Meheus, and X. Parent, eds. *Deontic Logic in Computer Science (DEON 2014)*. Lecture Notes in Computer Science 8554, pp. 49–64. Springer, 2014.
- [4] E. Calardo, G. Governatori, and A. Rotolo. A Sequence Semantics for Deontic Logic. Tech. rep. 8580. NICTA, 2015.
- [5] J. Carmo and A. J. Jones. Deontic Logic and Contrary-to-Duties. In D. M. Gabbay and F. Guentner, editors, *Handbook of Philosophical Logic. 2nd Edition*. Vol. 8, pp. 265–343. Springer, 2002.
- [6] G. Governatori. Business Process Compliance: An Abstract Normative Framework. *IT – Information Technology*, 556:231–238, 2013.
- [7] G. Governatori. Representing Business Contracts in RuleML. *International Journal of Cooperative Information Systems*, 142-3:181–216, 2005.
- [8] G. Governatori, F. Olivieri, A. Rotolo, and S. Scannapieco. Computing Strong and Weak Permissions in Defeasible Logic. *Journal of Philosophical Logic*, 426:799–829, 2013.
- [9] G. Governatori and A. Rotolo. Logic of Violations: A Gentzen System for Reasoning with Contrary-To-Duty Obligations. *Australasian Journal of Logic*, 4:193–215, 2006.
- [10] G. Governatori and S. Sadiq. The Journey to Business Process Compliance. In J. Cardoso and W. van der Aalst, editors, *Handbook of Research on BPM*, pp. 426–454. IGI Global, 2009.
- [11] G. Piolle. A Dyadic Operator for the Gradation of Desirability. In G. Governatori and G. Sartor, eds. *10th International Conference on Deontic Logic in Computer Science (DEON 2010)*. Lecture Notes in Computer Science 6181, pp. 33–49. Springer, 2010.
- [12] A. Pnueli. The Temporal Logic of Programs. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science (SFCS '77:)* pp. 46–57. IEEE Computer Society, 1977.
- [13] H. Prakken and M. J. Sergot. Contrary-to-Duty Obligations. *Studia Logica*, 571:91–115, 1996.
- [14] R. H. Thomason. Deontic Logic Founded on Tense Logic. In R. Hilpinen, editor, *New Studies on Deontic Logic*, pp. 165–176. Kluwer, 1981.
- [15] M. Y. Vardi. Branching vs. Linear Time: Final Showdown. In T. Margaria and W. Yi, eds. *7th International Conference Tools and Algorithms for the Construction and Analysis of Systems, (TACAS 2001)*. Lecture Notes in Computer Science 2031, pp. 1–22. Springer, 2001.