Contrary-To-Duties in Games

Paolo Turrini, Xavier Parent,
Leendert van der Torre, and Silvano Colombo Tosatto

Individual and Collective Reasoning Group (ICR)
Computer Science and Communications (CSC)
Faculty of Sciences, Technology and Communication (FSTC)
University of Luxembourg
6, rue Richard Coudenhove - Kalergi
L-1359 Luxembourg
paolo.turrini,xavier.parent,
leon.vandertorre,silvano.colombotosatto@uni.lu
http://icr.uni.lu

Abstract. The aim of the paper is to bring to the realm of game theory the well-known deontic notion of contrary-to-duty (CTD) obligation, so far not investigated in relation to optimality of strategic decisions. We maintain that, under a game-theoretical semantics, CTDs are well-suited to treat sub-ideal decisions. We also argue that, in a wide class of interactions, CTDs can used as a compact representation of coalitional choices leading to the achievement of optimal outcomes. Finally we investigate the properties of the proposed operators.

Keywords: Deontic logic, games, CTDs, optimality

1 Introduction

Horty's interaction-theoretical account of deontic logic [9] has shown that when classical deontic concepts, such as obligations, permissions and prohibition, are interpreted in game-like structures they acquire new meanings in terms of optimality of choices and shed new light on strategic interaction.[1]

[1] As a matter of fact the models Horty uses to interpret his deontic operators are not *strictu sensu* strategic games in the sense of [12], but can be thought as a strategic game forms — technically, games without preference relations — endowed with a unique utility function (and not one per player, as happens in strategic games), representing an abstract notion of betterness applying to all players. As the following quotation shows, Horty's proposal consists of viewing choices that should be performed as carrying a meaning in terms of an underlying notion of optimality, i.e. as optimal choices at players disposal.

> "In the past, the task of mapping the relations between deontic logic and act utilitarianism has resulted in surprising difficulties, leading some writers to suggest the possibility of a conflict in the fundamental principles underlying the two theories. One source of these difficulties, I believe, is the gap

Up till now several follow-up contributions, starting with Kooi and Tamminga [10], and continuing with Broersen *et al.* [4], and Turrini *et al.* [19], have pushed Horty's view further, working out notions such as moral obligations in the interest of a set of players, socially optimal norms, agreements and contracts, in a full-fledged game-theoretical framework, where Horty's utilitarian approach is made relative to players' (and coalitions') specific preferences.

Along the same line of these last contributions we aim at bringing to the realm of game theory the well-known deontic notion of *contrary-to-duty (CTD) obligation*, which states what should be done in case some primary obligation is already violated, and which has so far not been investigated for the special case of strategic interaction.

Marek Sergot (mainly together with Henry Prakken in [15, 14]) made fundamental advances in the conceptual and logical study of CTDs. These include:

- the definition of a CTD obligation as presupposing a context in which a primary obligation is already violated;
- the compilation of a set of benchmark examples against which CTDs must be assessed;
- the use of consistency provisos to block undesirable consequences stemming from the logical representation of CTDs scenarios;

The examples and the constraints studied by Marek Sergot were not intended to address CTDs in situations where strategic interaction plays an explicit role, which is what the present article is concerned with.

Our main observation is that, when faced with interactive decision-makers, issuing commands of the type "it is obligatory that φ, but if $\neg\varphi$ then it is obligatory that ψ" bears consequences in terms of strategic decisions. In particular we will see that if the equation suggested by the game-theoretical approach to deontic concepts *obligatory actions are rational decisions* holds for classical deontic operators, an alternative equation, i.e. *contrary-to-duty actions are second-best decisions*, holds for contrary-to-duty obligations.

Concretely, we argue that contrary-to-duty obligations can be meaningfully used *to reason about* the achievement of those optimal outcomes that can only be obtained by making second-best choices. One instance of this class is provided by games in which players are in possession of complementary goods and have the possibility of exchanging some of them, as illustrated by the following example.

Example 1 (Left & right shoes). Consider a scenario in which two players, i and j, possess two different, but complementary, types of resources: *left shoes* and

between the subjects of normative evaluation involved in the two areas: while deontic logic has been most successfully developed as a theory of what ought or ought not to be, utilitarianism is concerned with classifying actions, rather than states of affairs, as right or wrong. The present account closes this gap, developing a deontic logic designed to represent what agents ought to do within a framework that allows, also, for the formulation of a particular variant of act utilitarianism, the dominance theory" [9, p.70].

right shoes. The starting situation consists of player i possessing two left shoes and player j two right shoes. The underlying assumption is that players gain more utility by possessing more resources. In addition it is better for players to possess one resource of each type than possessing two resources of the same type. Finally, we will assume that, at the moment of taking a decision, each player is aware of the other player's options, but need not be aware of his preferences.[2]

The game consists of a single exchanging round. During this round each player decides how many of its resources it wants to concede to the other player. Considering that at the beginning each player has two resources of the same type, each player has three available options: conceding *none*, *one* or *two* resources to the other. The players cannot negotiate during the exchange and both have to decide and exchange simultaneously.[3]

In this first example players are taken to be utility maximizers — an assumption that will be dropped later on in the paper — and, intuitively, their best strategy in the game is the decision of keeping their resources for themselves. In this case, considering that both players adopt the same strategy, there is no exchange of resources between the participants, leaving the starting situation unchanged. Notice however that the outcome of the game is not an optimal outcome, as players could be better off by exchanging some goods, rather than not. Notice as well that if players acted with no greed, conceding all their resources, the outcome of the game would be no good either.

Starting out from this observation, we can point out how normative statements imposing players to behave *extremely selfishly* — but a similar argument can be brought for *extremely altruistic* behavior — would lead the system to suboptimal outcomes. Instead, we would like to tell our players to behave sub-ideally, conceding some resources but not all, in such a way that even in the presence of sub-ideal decisions an ideal outcome can be reached.

Paper Structure The paper is structured as follows. In Section 2 we introduce the main technical notions coming from the literature on game theory and deontic logic. In Section 3 we describe the mathematical structures that we use to interpret our deontic language, and which is studied in Section 4. All the main definitions are illustrated by means of the abovedescribed example, which is comprehensively analyzed in Section 5. The conclusive section wraps up the paper and suggest ideas for future development.

[2] As will be clear later, even though we will model the example as a one-shot strategic game, we will not need to postulate demanding epistemic assumptions such as common knowledge of the game structure (players know each others' strategies and preferences) or common knowledge of rationality (i.e. players are rational and everyone knows this and everyone knows that everyone knows this etc.), but it will be enough to have players that know what their opponents can do.

[3] Simultaneity of events is a common feature in game-theoretical examples and should be thought as an expedient to model players' unawareness of the actual choices that their opponents have taken.

2 Preliminaries

This section introduces the mathematical preliminaries needed in the rest of the paper. We start with consequentialist models [10], a simplification of STIT models [2] studied by Horty, which represent the coalitional power of players in games; then we move on to treat preference relations and a notion of dominance among choices available to coalitions.

The present work deals with strategic interaction. Therefore the basic ingredients we will be working with are a finite set N, to be understood as a set of *players*, and a set W to be understood as a set of *alternatives*. Players are denoted $i, j, k, \ldots$ while sets of players, i.e. elements of 2^N, are denoted $C, C', C'', \ldots$ and are henceforth called *coalitions*. The coalition made by all players, i.e. the set N, will be referred to as the *grand coalition*, while the coalition made by the players not belonging to a coalition C will be denoted as $\overline{C}$ and referred to as the set of *opponents* of C. Alternatives are denoted $u, v, w, \ldots$ and are also called *outcomes*, *states* or *worlds*. Players are assumed to have preferences over the alternatives. Therefore, each player i is endowed with a preference order $(\succeq_i)_{i \in N}$, a total preorder on the set of alternatives, where $v \succeq_i w$ has the intuitive reading that outcome v is *at least as good* as outcome w for player i. The corresponding strict partial order is defined as expected: $v \succ_i w$ if, and only if, $v \succeq_i w$ and not $w \succeq_i v$, to mean that for player i outcome v is *strictly better* than outcome w. The notation $\prec_i, \preceq_i$ for the reverse relations will be used as well when no confusion can arise.

2.1 Consequentialist Models

The theory of agency adopted in this paper takes inspiration from the one presented by Horty in [9] to study deontic notions within a utilitarian perspective. There an interpretation of coalitional rationality is proposed, based on STIT models [2], a branching-time account of coalitional ability. For the present purposes, an adoption of the full-blown history-based models used by Horty would take this work far from its scope, and therefore we resort to the simpler *consequentialist* models, that share with Horty's models the local features that are necessary to treat one shot interactions. Consequentialist models have been used already as one-shot STIT counterpart by Kooi and Tamminga [10], who also present a model of coalitional rationality with classical utility functions, which has much in common with our account.

Here is the formal definition.

Definition 1 (Choice Structures). *A* choice structure *is a triple*

$$(W, N, Choice)$$

where W is a set of outcomes, N a finite set of players, and $Choice : 2^N \to 2^{2^W}$ a function defined as follows:

- *for each $i \in N$, $Choice(\{i\})$ is a partition of W;*

– *Let S be the set of functions $s : N \to 2^W$ such that for each $i \in N$, $s(i) \in Choice(\{i\})$. We have that for $C \subseteq N$:*
  - $\bigcap_{i \in C} s(i) \neq \emptyset$*, for every $s \in S$, i.e. the pairwise intersection of players' choices is nonempty.*
  - $Choice(C) = \{\bigcap_{i \in C} s(i) \mid \text{ for } s \in S\}$*, i.e. coalitional choices are constructed by taking the pairwise intersection of individual choices.*

– $|Choice(N)| = |W|$*, i.e. players together can force any available outcome.*[4]

The definition illustrates choice structures as a description of how groups of players (possibly empty, or made by one single player) are able to decide the future course of events. Choice structures model the possible decisions of coalitions by means of the following two key features:

– The choices available to a coalition are *a partition of the set of possible states*. The sets in this partition, i.e. the available choices, intersect nontrivially [5] with each set in the partition of the opposing coalition. In this view, choosing means deciding that the resulting outcome of the interaction will be contained in some set of worlds, leaving to the opponents the possibility of choosing within that set;

– The choices available to a coalition are *a combination of all the possible choices available to its members*, which is obtained by pairwise intersecting their choice structures. In this view, a coalition of players is assumed to be able to fully coordinate their members and to dispose of their collective choices.

These features make choice structures mathematically equivalent (modulo coalitions) to strategy profiles in games.[6]

Example 2. The *left shoes & right shoes* game in Example 1 can be described as a choice structure. In Table 1, we show its matrix representation.

The columns of the table represent the possible choices of player i: K_0, K_1 and K_2. For each choice K_n, n represents the number of resources that the player i concedes to the opponent during the exchanging round. In the same way, the rows of the table represent the possible choices of player j.

The cells of the table represent the results of the exchange round. The top-right corner of each cell represents the resources that player i owns once that combination of choices is made. Similarly the bottom-left corner represents the

[4] This condition is sometimes referred to as *rectangularity* [1]. Rectangularity is not assumed in [10], but it is presupposed by all game-theoretical matrix representation of choice structures.

[5] The intersection is nonempty.

[6] Strategy profiles are classically modelled as bijective functions from the set of individual strategies to the set of outcomes — often represented by vectors of utilities — It is straightforward to notice that the strategy profiles *modulo a coalition* generate partitions of the set of outcomes that are obtained by pairwise intersecting the partitions assigned to the members of that coalition. This simple observation allows us to represent choice structures as matrixes, just like standard strategic game forms.

resources owned by player j after the exchange. The number of resources possessed by player $p \in \{i, j\}$ after the exchange is represented as follows: $x_p | y_p$ where x_p refers to the number of *right shoes* and y_p to the left shoes possessed.

$j \backslash i$	K_0	K_1	K_2
S_0	$2 \mid 0$ \ $0 \mid 2$	$2 \mid 1$ \ $0 \mid 1$	$2 \mid 2$ \ $0 \mid 0$
S_1	$1 \mid 0$ \ $1 \mid 2$	$1 \mid 1$ \ $1 \mid 1$	$1 \mid 2$ \ $1 \mid 0$
S_2	$0 \mid 0$ \ $2 \mid 2$	$0 \mid 1$ \ $2 \mid 1$	$0 \mid 2$ \ $2 \mid 0$

Table 1. Strategy outcomes for left shoes & right shoes in a choice structure

Utility calculation It can be noted that a preference relation can be assigned to individual players, once we are able to calculate the invidual utility for each outcome. In our case the utility u_p for each player $p \in \{i, j\}$ after the exchange round is calculated using the following evaluation function:

– $u_p = 2x_p + y_p$, whenever $y_p > x_p$;
– $u_p = 2y_p + x_p$, otherwise.

Intuitively the utility function is computed in two steps:

1. The first step counts the individual value of the single resources.
2. The second step attributes additional value to the outcome if there is at least one element per resource type, representing the additional value of having both of them.

The outcomes calculated by using the utility functions are represented in Table 2. In the cells of the table are shown the outcomes for the players i and j dependent on the choices made. On the top-right corner of each cell is represented the outcome for player i and in the bottom-left the outcome for player j.

Consequentialist models are obtained by adding to choice structures a valuation function, which gives a description of the relevant properties holding at each state.

Definition 2 (Consequentialist Models [10]). *A* consequentialist model *is a pair* (Γ, V) *where* Γ *is a choice structure and* V *is a* valuation function, *i.e. a function from the set of states* W *to the powerset of a countable set of propositions* $Prop$*, with the usual understanding that propositions that get assigned to a state should be understood as true at that state.*

The valuation function tells us what propositions correspond to what outcomes or worlds. By using them we will be able to reason logically on the properties of choice structures.

j \ i	K_0	K_1	K_2
S_0	2, 2	4, 1	6, 0
S_1	1, 4	3, 3	4, 1
S_2	0, 6	1, 4	2, 2

Table 2. Utilities for left shoes & right shoes added to its choice structure

2.2 Dominance

As for the notion of coalitional rationality, we employ the notion of dominance, which represents the comparison of choices at a coalition's disposal, taking the moves of the opponents into account. As the dominance relations compares sets of outcomes and the preference relations are formulated on individual outcomes, the following definition bridges the gap between the two notions.

Definition 3 (Lifting). *Let $X, Y \subseteq W$ be two sets of outcomes and let $i \in N$ be a player. X is* preferred to *Y by i — which we denote $X \trianglerighteq_i Y$ — whenever $w \succeq_i w'$ for all $w \in X$ and $w' \in Y$.*

The lifting we have just defined states that a choice X is better than a choice Y only if all elements in X are better than all elements in Y, relative to the preferences of some player. This type of lifting, often called *for all - for all* lifting, is fairly simple and yet particularly well-suited for characterizing standard solution concepts in games.[7] We are now ready to define the notion of dominance among coalitional choices.

Definition 4 (Dominance). *Let $K, K' \in Choice(C)$ and $\succeq_i \subseteq W \times W$ a preference relation over the outcomes for each player. K* dominates *K' if and only if for all $S \in Choice(\overline{C})$ we have that $K \cap S \trianglerighteq_i K' \cap S$ for all $i \in C$.*

Intuitively what the definition says is that, when a coalition C disposes of two choices K and K', K will be preferred to K' in case all worlds in $K \cap S$ are better than those in $K' \cap S$ for each member of C, for each possible choice S of the opposing coalition.[8]

Definition 4 notably simplifies what is to be found in the literature. Both in [9] and [10] a utility function is employed associating to each outcome (histories in Horty's framework) an element of a closed interval in the reals (positive reals in Horty's framework, the interval $[-5, 5]$ in Kooi and Tamminga's framework).

[7] For a discussion on its merits and the possible alternatives , we refer to [18].

[8] The notion of dominance clearly resembles that of *dominant strategy* typical of strategic games [12].

Example 3. In Example 1, based on the outcomes shown in Table 2, the dominance relation among coalitional choices can be constructed as follows.

Consider $\{j\}$, the coalition consisting only of player j. We have $Choice(\{j\}) = \{S_0, S_1, S_2\}$. According to Definition 4, choice S_0 dominates S_1 and S_2, because $\forall K_n \in Choice(\{i\})$ we have that $S_0 \cap K_n \trianglerighteq_j S_1 \cap K_n \trianglerighteq_j S_2 \cap K_n$.

Symmetrically, consider $\{i\}$. We have $Choice(\{i\}) = \{K_0, K_1, K_2\}$. According to Definition 4, choice K_0 dominates K_1 and K_2, because $\forall S_n \in Choice(\{j\})$ we have that $K_0 \cap S_n \trianglerighteq_i K_1 \cap S_n \trianglerighteq_i K_2 \cap S_n$.

As expected, from the point of view of each player, the *dominant* strategy for him is to play selfishly, or in other words to keep the resources for himself.

F-dominance As pointed out in the introduction several contributions have generalized Horty's notion of dominance. In the present paper we focus on the notion of dominance in the interest of some coalition, first studied by Kooi and Tamminga, although in a more involved formulation that we simplify as follows.

Definition 5 (F-dominance [10]). *Let $K, K' \in Choice(C)$, $F \subseteq N$ and $\succeq_i \subseteq W \times W$ a preference relation over the outcomes for each player. K* F-dominates *K' if and only if for all $S \in Choice(\overline{C})$ and for all $i \in F$ we have that $K \cap S \trianglerighteq_i K' \cap S$.*

Intuitively the definition provides a notion of dominance among choices of a coalition C looked at from the point of view of another coalition F. Obviously, when F and C coincide, F-dominance and dominance do, as well. As F-dominance is more general than dominance it can be used to analyze a wider class of situations when players do not necessarily behave selfishly. A similar stance is taken in [19] to reason on exchange of favours in deontic logic.

Example 4. In Example 1 the analysis of F-dominance among coalitional choices can be carried out as follows. Consider that each single player now takes a decision looking at the welfare of its opponent. With this in mind the coalition $\{i\}$ adopts a $\{j\}$-dominance, while $\{j\}$ adopts a $\{i\}$-dominance. The *F-dominance* in Definition 5, boiling down in our case to a $\overline{C}$-dominance for each coalition C, indicates as ideal the strategy of playing generously.

According to Definition 5, choice S_2 $\{i\}$-dominates S_1 and S_0 because $\forall K_n \in Choice(\{i\})$ we have that $S_2 \cap K_n \trianglerighteq_i S_1 \cap K_n \trianglerighteq_i S_0 \cap K_n$.

Likewise, if player i makes his choices based on the utility for player j (as shown in the bottom-left of the cells), choice K_2 $\{j\}$-dominates K_1 and K_0, because $\forall S_n \in Choice(\{j\})$ we have that $K_2 \cap S_n \trianglerighteq_j K_1 \cap S_n \trianglerighteq_j K_0 \cap S_n$.

3 Models

This section brings together the model-theoretic notions defined in the previous part of the paper and defines the structures on which to interpret our deontic language.

We will work with consequentialist models (Definition 2) endowed with a preference order $\succeq_i$ for each player $i \in N$, i.e. our structures will have the form

$$(\Gamma, V, (\succeq_i)_{i \in N})$$

where Γ and V are given as in Definition 2.

In the style of Kooi and Tamminga, we introduce a relation of choice equivalence in a consequentialist model. The idea is that, if two worlds w and w' are in this relation with respect to a coalition C, then the coalition cannot alone elect either outcome to be the final outcome of the game.

Definition 6 (C-choice equivalence). *Let $\Gamma = (W, N, Choice)$ be a choice structure. The relation $\sim_C \subseteq W \times W$ of C-choice equivalence is defined as follows:*

$$w \sim_C w' \textit{ if and only if } w \in K \textit{ implies that } w' \in K, \textit{ for some } K \in Choice(C)$$

Intuitively, if two worlds w, w' are in a relation $w \sim_C w'$, then *only* the opponents of coalition C can decide whether the outcome will be w or w' or some other outcome linked to them by the same relation.

Proposition 1. *Let $\Gamma = (W, N, Choice)$ be a choice structure. The set*

$$\{[w] \mid w' \in [w] \textit{ if and only if } w \sim_C w'\}$$

is a partition of W.

Proof. Straightforward. Notice on the fly that the set $\{[w] \mid w' \in [w] \text{ iff } w \sim_\emptyset w'\}$ has cardinality 1.

4 Language and Semantics

In this section we introduce the syntax of our language and the interpretation of its formulas in terms of the models provided in the previous section. We start out by defining the language $\mathcal{L}$, an extension of propositional logic with modalities to reason about obligations and coalitional choices.

Definition 7 (Syntax). *Let Prop be a countable set of atomic propositions. The formulas of the language $\mathcal{L}$ have the following grammar:*

$$p \mid \neg\varphi \mid \varphi \wedge \varphi \mid E\varphi \mid [C]\varphi \mid \bigcirc_C \varphi$$

where $p \in Prop$ and $C \subseteq N$. The informal reading of the modalities is "there exists a world satisfying φ", "coalition C achieves φ", "it is obligatory for coalition C to choose φ". Within this language the fact that a coalition *can achieve* a property is expressed by formulas such as $E[C]\varphi$, intuitively saying that there is a world where coalition C chooses φ. This reading of strategic ability is in line with the standard treatment of STIT-like logics in Kripke models [3].

Definition 8. *Let M be a consequentialist model with a set of outcomes W and let $w \in W$. The interpretation of the formulas in Definition 7 with respect to a tuple M, w is as follows:*

$$
\begin{array}{lll}
M, w \models p & \textit{iff} & w \in V(p) \\
M, w \models \neg\varphi & \textit{iff} & M, w \not\models \varphi \\
M, w \models \varphi \wedge \psi & \textit{iff} & M, w \models \varphi \textit{ and } M, w \models \psi \\
M, w \models [C]\varphi & \textit{iff} & M, w' \models \varphi \textit{ for all } w' \textit{ with } w \sim_C w' \\
M, w \models E\varphi & \textit{iff} & \textit{there exists } w' \in W \textit{ such that } M, w' \models \varphi \\
M, w \models \bigcirc_C \varphi & \textit{iff} & i)\ \varphi^M \in Choice(C) \textit{ and} \\
 & & ii)\ \varphi^M \textit{ dominates each } K \neq \varphi^M \textit{ with } K \in Choice(C)
\end{array}
$$

where $\varphi^M = \{w \mid M, w \models \varphi\}$ is called the *extension* or *truth set* of formula φ.
We omit the reference M when it is clear which model is intended.

The interpretation of the obligation modality $\bigcirc_C$ deserves some comment. The evaluation rule says that the formula φ in its scope is obligatory for coalition C if the proposition that φ expresses:

- is an available choice for C,
- it dominates every other available choice for C.

Obligation modalities in this language boild down to succints statements on how coalitions *should rationally play*, comparing each available choice against the possible reactions of their opponents. Different kind of obligations — 'altruistic' ones — will be introduced later on in the paper and will make use of the notion of F-dominance (Definition 5).

Notice that by the definition of Choice Structure (Definition 1) we can formulate a universal modality such as E looking at what the empty coalition can achieve.

Proposition 2. *For each consequentialist model M and each $w \in W$ we have that $M, w \models E\varphi \leftrightarrow \neg[\emptyset]\neg\varphi$.*

Proof. Direct consequence of Definition 1.

4.1 Contrary-to-Duty Obligations

In this section we expand the language with formulas of the type $\bigcirc_C(\psi/\varphi)$, to express what coalition C should do if some state of affairs φ is already the case. The idea is that, if φ represents a violation of a main obligation, then $\bigcirc_C(\psi/\varphi)$ is a *contrary-to-duty* obligation.

Definition 9 (CTD). *Let M be a consequentialist model and w a state in its domain. The interpretation of formulas of the type $\bigcirc_C(\psi/\phi)$ is as follows:*

$$
\begin{array}{ll}
M, w \models \bigcirc_C(\psi/\varphi) \textit{ iff} & i)\ \psi^M \in Choice(C) \textit{ and} \\
 & ii)\ \psi^M \textit{ dominates each } K \neq \psi^M \in Choice(C) \setminus \overline{\varphi^M}
\end{array}
$$

where $\overline{\varphi^M}$ is the complement of φ^M, and $\setminus$ is set-theoretic difference.

Definition 9 behaves similarly to the evaluation rule for the monadic $\bigcirc_C$. The only difference is that the set of $\neg\varphi$-worlds (whose corresponding formula φ, when in the scope of $\bigcirc_C$, corresponds to the best option) is no longer available in the choice set. The intuition behind formulas of the type $\bigcirc_C(\psi/\varphi)$ is that once the best option $\neg\varphi$ is ruled out, coalition C is left with ψ as best alternative. To quote Hansson, the role of CTDs is to "make the best out of the sad circumstances" [7]. One might refer to $Choice(C) \setminus \overline{\varphi^M}$ as a *zoom-in* operation. Its main effect is to rule out outcomes that are no longer relevant in the comparison process.

Technically, the game-theoretical account of CTDs can be described as a combination of the Lewis/Hansson preference-based account of conditional obligation (see [7] and [11]) with so-called neighborhood semantics. To see this, Table 3 below gives a typical model of $\{\bigcirc_j\neg p, \bigcirc_j(q/p)\}$ – from now on we will drop curly brackets for singleton C. The numbers show the preference order for j. $\bigcirc_j\neg p$ holds, because $S_3 = \neg p^M$ dominates S_2 and S_1. When evaluating

j \ i	K_1	K_2
S_1	$w_1 : p, \neg q$ 0	$w_2 : p, \neg q$ 0
S_2	$w_3 : p, q, r$ 0.5	$w_4 : p, q$ 0.5
S_3	$w_5 : \neg p, \neg q,$ 1	$w_6 : \neg p, \neg q,$ 1

Table 3. A model of $\{\bigcirc_j\neg p, \bigcirc_j(q/p)\}$

$\bigcirc_j(q/p)$, S_3 is taken out of from the choice set. So, $\bigcirc_j(q/p)$ holds, because S_2 dominates S_1. The resemblance with the Hansson/Lewis account should be obvious to the reader, as the latter also interprets the obligatory worlds as the best worlds according to a given preference order. If the preference orders coincide, the obligations returned are the same. For instance an Hansson/Lewis account would also yield $\bigcirc_j\neg p$, because the best worlds are all $\neg p$-worlds. And it would also give $\bigcirc_j(q/p)$, because the best p-worlds are all q-worlds.

4.2 Some valid/invalid formulas

We first note that the monadic $\bigcirc_C$ can be defined in terms of the dyadic $\bigcirc_C(-/-)$ in the usual way.

Proposition 3. *For each consequentialist model M and each $w \in W$ we have that $M, w \models \bigcirc_C\varphi \leftrightarrow \bigcirc_C(\varphi/\top)$.*

Proof. $Choice(C) \setminus \overline{\top^M} = Choice(C)$.

Even if the game-theoretical account of CTD bears resemblance with the Hansson/Lewis one, there is an importance difference between the two. The operator $\bigcirc_C(-/-)$ does not satisfy the principle known as "Weakening the Consequent", and neither does it satisfy the law "S" below, named after Shoham [16] who first discussed it in the context of the study of non-monotonic reasoning. This is because the evaluation rule also incorporates some aspects of so-called neighborood semantics. This is condition i) in Definition 9. For $\bigcirc_C(\psi/\varphi)$ to be true, the truth set ψ^M must be part of the choice set of coalition C. As a result of this, the principle that 'ought implies can' holds in the following form: $\bigcirc_C\varphi \to E[C]\varphi$.

Proposition 4. *Let A be the universal modality, defined as the dual of E. Non-validities include*

$$(\bigcirc_C(\psi/\varphi) \wedge A(\psi \to \xi)) \to \bigcirc_C(\xi/\varphi) \qquad \text{(WC)}$$

$$(\bigcirc_C(\psi/\varphi \wedge \varphi') \to \bigcirc_C(\varphi' \to \psi/\varphi) \qquad \text{(S)}$$

Proof. Table 4 depicts a typical countermodel to $\bigcirc_j(p \wedge q) \to \bigcirc_j p$, which is a special case of WC.

j \ i	K_1	K_2
S_1	$\neg p, \neg q$ 0.25	$\neg p, \neg q$ 0.25
S_2	$\neg p, q$ 0.5	$\neg p, q$ 0.5
S_3	$p, \neg q$ 0.75	$p, \neg q$ 0.75
S_4	p, q 1	p, q 1

Table 4. Failure of WC

Since $S_4 = p \wedge q^M$ dominates S_1, S_2, and S_3, the obligation $\bigcirc_j(p \wedge q)$ holds. But $p^M = S_3 \cup S_4$, and p^M is not in the choice set of j. Therefore $\bigcirc_j p$ fails.

Table 5 shows a countermodel to S.
Since $S_3 = r^M$ dominates S_2, the obligation $\bigcirc_j(r/p \wedge q)$ holds. But $q \to r^M = S_1 \cup S_3 \notin \mathit{Choice}(j)$, and thus $\bigcirc_j(q \to r/p)$ fails. □

The same pattern is involved in the failure of the principle CTD below, which is weaker than WC:

$$\bigcirc_C\, \varphi \wedge \bigcirc_C(\psi/\neg\varphi) \to \bigcirc_C(\varphi \vee \psi) \qquad \text{(CTD)}$$

Failure of these laws might be considered bad news. Especially (CTD) seems constitutive of the notion of CTD. Roughly speaking, it says that the obligation

$j \backslash i$	K_1	K_2
S_1	$p, \neg q, \neg r$ 0.05	$p, \neg q, \neg r$ 0.05
S_2	$p, q, \neg r$ 0.25	$p, q, \neg r$ 0.25
S_3	p, q, r 0.5	p, q, r 0.5

Table 5. Failure of S

of φ together with the obligation of ψ given $\neg\phi$ entail that the agent is under the obligation of φ (the best) or ψ (the second-best).

(WC) has good intuitive support too. As Sergot and Prakken observe, "someone who is told not to kill must surely be able to infer that he or she ought not to kill by strangly, say" [14, p.224].

We would however like to point out that the abovementioned failures are *all* to be attributed to condition i) in Definition 9, which does not impose monotonicity of coalitional action (being able to choose ϕ does not imply being able to choose $\phi \vee \psi$). We will see, at the end of this section, that a slightly more liberal definition on choice sets, requiring closure under supersets, allows one to validate these laws.

In spite of the above, Proposition 5 shows that the logic is not as weak as one might think at first sight, listing a number of inference patterns that are validated. Note that Hansson's *official* system DSDL3 supports these laws either as they stand or in a slightly modified form.

Proposition 5 (Validities). *For each consequentialist model M and each $w \in W$ we have that*

$$M, w \models A(\varphi \leftrightarrow \varphi') \to (\bigcirc_C(\psi/\varphi) \leftrightarrow \bigcirc_C(\psi/\varphi')) \quad \text{(Equivalence)}$$

$$M, w \models \bigcirc_C(\psi/\varphi) \wedge \bigcirc_C(\psi'/\varphi) \wedge E(\psi \wedge \psi') \to \bigcirc_C(\psi \wedge \psi'/\varphi) \quad \text{(Consistent And)}$$

$$M, w \models \bigcirc_C(\psi/\varphi) \wedge \bigcirc_C(\psi/\varphi') \wedge E(\neg\varphi \wedge \neg\varphi') \wedge E[C]\neg\varphi \wedge E[C](\neg\varphi') \to \bigcirc_C(\psi/\varphi \vee \varphi') \quad \text{(Consistent OR)}$$

$$M, w \models \bigcirc(\psi/\varphi) \wedge \bigcirc(\xi/\psi) \to \bigcirc(\xi/\varphi) \quad \text{(DD)}$$

$$M, w \models \bigcirc_C(\psi/\varphi) \wedge A\varphi \to \bigcirc_C\psi \quad \text{(SFD)}$$

The labels DD, and SFD stand for "Deontic Detachment", and "Strong Factual Detachment", respectively.

The law 'Equivalence' permits the replacement of equivalent sentences in the antecedent of deontic conditionals. 'SFD' is a principle of modus-ponens (or detachment) for obligations. It tells us when a conditional obligation can be deconditionalized: the antecedent must be settled as true.[9] The two appear as they stand in DSDL3.

The law 'Consistent And' restricts aggregation to those obligations whose consequents are jointly possible. Similarly for the law 'Consistent OR'. It allows to reason by cases if some consistency proviso is satisfied. In DSDL3, the principles of aggregation and reasoning by cases hold unrestrictively.

'DD' is a principle of chaining for obligations. DSDL3 supports a weaker version known as 'cut': from $\bigcirc(\psi/\varphi)$ and $\bigcirc(\xi/\varphi \wedge \psi)$ infer $\bigcirc(\xi/\varphi)$.

Proof. For Equivalence, this is routine check.

The reason why Consistent And holds can be seen as follows. Suppose both $\bigcirc_C(\psi/\varphi)$ and $\bigcirc_C(\psi'/\varphi)$ hold at w. Then, $\|\psi\|^M, \|\psi'\|^M \in Choice(C)$, and both $\|\psi\|^M$ and $\|\psi'\|^M$ dominate any $S \in Choice(C) \setminus \overline{\|\varphi\|^M}$. By the fact that $M, w \models E(\psi \wedge \psi')$ we have that $\|\psi\|^M = \|\psi'\|^M$, as $Choice(C)$ is a partition. Hence $\|\psi \wedge \psi'\|^M \in Choice(C)$, and $\|\psi \wedge \psi'\|^M$ dominates any $S \in Choice(C) \setminus \overline{\|\varphi\|^M}$. This shows that $\bigcirc_C(\psi \wedge \psi'/\varphi)$ holds at w too.

The picture below shows why Consistent OR is valid. Here the convention is that worlds on a lower level dominates all those on an upper level. To avoid cluster, we list the propositions that are made true, and omit those made false.

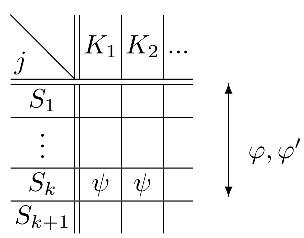

Put $C = \{j\}$. φ and φ' are true everywhere from S_1 downwards up to S_k. ψ is true at S_k only. $\bigcirc_C(\psi/\varphi)$ and $\bigcirc_C(\psi/\varphi')$ hold because $S_k = \|\psi\|^M$ dominates any of $S_1, ...,$ and S_{k-1}. In S_{k+1}, neither φ nor φ' is true. Therefore, when evaluating $\bigcirc_C(\psi/\varphi \vee \varphi')$, the choice set does not change, and thus $\bigcirc_C(\psi/\varphi \vee \varphi')$ also holds. That this holds in general is a consequence of the truth of $E(\neg\varphi \wedge \neg\varphi') \wedge E[C]\neg\varphi \wedge E[C]\neg\varphi'$.

For DD, the argument is similar. The diagram below should provide enough information to convince the reader of the soundness of the inference pattern. It says that, given transitivity of $\succeq_j$, if ψ dominates the φ-zone, and ξ in turn dominates the ψ-zone, then ξ dominates the φ-zone.

[9] The qualifier 'strong' is commonly used to avoid any confusion with the principle obtained by replacing $A\varphi$ with φ. This other principle is usually referred to as simply "Factual Detachment". For a good discussion of SFD, see Prakken and Sergot [14].

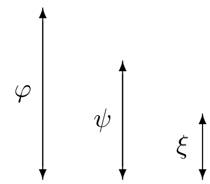

The fact that ξ^M is in the choice set follows from the assumption that $\bigcirc_C(\xi/\psi)$ holds.

The validity of SFD is straightforward. If $A\varphi$ holds, then $Choice(C) \setminus \overline{\varphi^M} = Choice(C)$. □

4.3 Monotonic obligation

In this section we generalize the account described in the previous sections. The idea is to incorporate a monotonicity condition in the semantics to secure the validity of such laws as (WC) and (CTD), which we think are desirable. For the unconditional operator, we adopt the following.

Definition 10 (Monotonic Obligations).

$$M, w \models \bigcirc^{\uparrow}_C \varphi \textit{ iff } \quad \begin{array}{l} i)\, \exists X \subseteq \varphi^M \textit{such that } X \in Choice(C) \textit{ and} \\ ii)\, X \textit{ dominates each } K \neq X \textit{ with } K \in Choice(C) \end{array}$$

Definition 10 is much alike the evaluation rule for in $\bigcirc_C$ in Definition 8. The former is obtained from the latter, by changing "=" into "⊆" in clause i). Thus, it is no longer required that exactly the set of worlds where φ is true is amongst those the agent can choose. It could be that he or she can choose a subset of them only.

With CTD a similar stance can be taken. But care must be taken in the formulation of the zoom-in operation appearing in clause ii) of Definition 9. The definition looks a bit more involved, but it has similar effects.

Definition 11 (Monotonic CTD obligation).

$$M, w \models \bigcirc^{\uparrow}_C(\psi/\varphi) \textit{ iff } \quad \begin{array}{l} i)\, \exists X \subseteq \psi^M \textit{ with } X \in Choice(C) \textit{ and} \\ ii)\, X \textit{ dominates each } K \neq X \in Choice(C) \setminus \bigcup_{Y \in Choice(C)} Y \subseteq \overline{\varphi^M} \end{array}$$

Below we argue that such refinements are also needed to deal with some of the typical CTDs scenarios discussed in the deontic logic literature.

Requirements There are known requirements that any satisfactory account of CTDs is expected to meet. These are discussed in depth by Carmo and Jones [5]. For present purposes, suffice it to consider the two basic ones. These are:

- the logical representation of the premises set should be consistent;

– the logical formulas used to represent the scenario should be independent from each other.

It is not difficult see that, if the above two definitions are used, then these two most basic requirements are met for an important class of CTDs scenarios, like the Chisholm scenario and the dog-and-sign scenario, among others. These two are structurally identical. Below we focus on the former.

Chisholm scenario. The premises set is $\{\bigcirc^{\uparrow}_{C} h, \bigcirc^{\uparrow}_{C}(t/h), \bigcirc^{\uparrow}_{C}(\neg t/\neg h), \neg h\}$, where h and t are for *helping* and *telling*, respectively. To show that the set is consistent amounts to showing that it is satisfiable in a model. Table 6 gives one such model.

$j \backslash i$	K_1
S_1	$w_1 : \neg h, t$ 0.25
S_2	$w_2 : \neg h, \neg t$ 0.5
S_3	$w_3 : h, \neg t$ 0.75
S_4	$w_4 : h, t$ 1

Table 6. Consistency of the Chisholm set

All the sentences are true at e.g. w_2. $\neg h$ holds there, and so does $\bigcirc_j^{\uparrow} h$. For $S_4 \subset h^M$, and S_4 dominates the other elements in the choice set of j. For the other sentences, the argument is similar.

The proof of independence is by showing that each formula can be falsified in a model that satisfies the other three formulae. For the propositional formula $\neg h$, this is just a matter of changing the world at which all the sentences are evaluated. For the normative sentences, this is just a matter of modifying the ranking in a suitable way, like in a usual preference-based semantics. For instance, table 7 makes $\bigcirc_j^{\uparrow} h$ false while making the other three formulae true at e.g. w_2. This demonstrates the independence of $\bigcirc_j^{\uparrow} h$.

$j \setminus i$	K_1
S_1	$w_1 : \neg h, t$ 0.75
S_2	$w_2 : \neg h, \neg t$ 1
S_3	$w_3 : h, \neg t$ 0.25
S_4	$w_4 : h, t$ 0.5

Table 7. Independence of $\bigcirc_j^{\uparrow} h$

For the independence of $\bigcirc_j^{\uparrow}(t/h)$ and $\bigcirc_j^{\uparrow}(\neg t/\neg h)$, the argument is similar, and left to the reader.

4.4 Obligations for someone else

The obligation operators in Definitions 7 and 9 can also be naturally generalized to obligations for someone else by replacing the notion of dominance in their interpretation with that of F-dominance of Definition [10].

Definition 12. *Let M be a consequentialist model and w a state in its domain. The interpretation of the formulas in Definitions 7 and 9 with respect to a tuple 'M, w' is as follows:*

$M, w \models \bigcirc_C^F \varphi$ *iff* *i)* $\varphi^M \in Choice(C)$ *and*
ii) φ^M *F-dominates each* $K \neq \varphi^M$ *with* $K \in Choice(C)$

$M, w \models \bigcirc_C^F (\psi/\varphi)$ *iff* *i)* $\psi^M \in Choice(C)$ *and*
ii) ψ^M *F-dominates each* $K \neq \psi^M \in Choice(C) \setminus \overline{\varphi^M}$

5 Back to the Example

Now we can revisit the *left & right shoes* game described in Example 1. First, we show that CTD obligations can be used to encode in the syntax the second-best decision for each player as specified by his own standard (selfishness, or altruism). Next, we show that, if the players go for the second-best, then the outcome of the game turns out to be the best one, according to their standard again. Paradoxical as it may seem, the players would serve their principles better by accepting to compromise them: the second-best is best.

Let p_0, p_1 and p_2 denote the propositions that j concedes none, one and two resources to i, respectively. Let q_0, q_1 and q_2 denote the propositions that i concedes none, one and two resources to j, respectively. Table 8 below recapitulates the moves available to the players along with the associated utilities.

j \ i	K_0	K_1	K_2
S_0	2 $w_0 :\ p_0, q_0$ 2	1 $w_1 : p_0, q_1$ 4	0 $w_2 : p_0, q_2$ 6
S_1	4 $w_3 : p_1, q_0$ 1	3 $w_4 : p_1, q_1$ 3	1 $w_5 : p_1, q_2$ 4
S_2	6 $w_6 : p_2, q_0$ 0	4 $w_7 : p_2, q_1$ 1	2 $w_8 : p_2, q_2$ 2

Table 8. Moves available with corresponding utilities

Selfishness

We first illustrate how selfishness on both sides leads the system to a suboptimal outcome.

Let us start with j. We have $p_0^M = S_0 = \{w_0, w_1, w_2\}$. As explained in Example 3, S_0 dominates S_1 and S_2. Therefore, according to Definition 8, the obligation $\bigcirc_j p_0$ holds; that is, from his own point of view j should concede nothing. The same goes for i. We have $q_0^M = K_0 = \{w_0, w_3, w_6\}$. As also explained in Example 3, K_0 dominates K_1 and K_2. Therefore, according to Definition 8, the obligation $\bigcirc_i q_0$ holds too; that is, from his own point of view i should concede nothing either.

However, if the players behave selfishly, and if i and j comply with $\bigcirc_i q_0$ and $\bigcirc_j p_0$, respectively, then the outcome of the game is w_0. This is clearly suboptimal, because in w_0 each player gets 2 only.

Now let us see what CTD obligations hold, starting with j. We have $p_1^M \in Choice(\{j\}) \setminus p_0^M = \{S_1, S_2\}$. Furthermore, S_1 dominates S_2. Therefore, according to Definition 9, the obligation $\bigcirc_j(p_1/\neg p_0)$ holds. Intuitively, the obligation says that, if j concedes something, then (from his own point of view) he should concede one only.

The same goes for i. On the one hand, $q_1^M \in Choice(\{i\}) \setminus q_0^M = \{K_1, K_2\}$. On the other hand, K_1 dominates K_2. So, the obligation $\bigcirc_i(q_1/\neg q_0)$ also holds. Intuitively, the obligation says that, if i concedes something, then (from his own point of view) he should concede one only.

Hence, we can see that each CTD obligation, in combination with a primary obligation, encodes in the syntax the second-best choice available to each player. But we can also see that, if the players go for it – in other words, if they accept to compromise their principles – then the outcome of the game is w_4, which turns out to be the optimal one. By choosing w_4, they both get the maximum, namely 3.

Altruism

A similar point can be made about altruistic behavior.

We have $\bigcirc_j^i p_2$ because $p_2^M = S_2 \in Choice(\{j\})$ and (as explained in Example 4) S_2 $\{i\}$-dominates S_1 and S_0. And we have $\bigcirc_i^j q_2$ because $q_2^M = K_2 \in Choice(\{i\})$ and (as also explained in Example 4) K_2 $\{j\}$-dominates K_1 and K_0. Hence, if each player is motivated by the interests of the other, then they each should concede 2. However, if both behave altruistically, the outcome of the game is w_8, and thus it is sub-optimal. They get 2 only.

Now let us see what CTD obligations hold. First, $\bigcirc_j^i(p_1/\neg p_2)$ holds, because $p_1^M = S_1 \in Choice(\{j\}) \setminus p_2^M = \{S_0, S_1\}$ and (as easily verified) S_1 $\{i\}$-dominates S_0. Next, $\bigcirc_i^j(q_1/\neg q_2)$ holds, because $q_1^M = K_1 \in Choice(\{i\}) \setminus q_2^M = \{K_0, K_1\}$ and (as easily verified too) K_1 $\{j\}$-dominates K_0. Again, these CTD obligations encode in the syntax the second-best choices available to the players. The first obligation says that, if j does not concede two resources, then (from i's point of view) he should concede one resource. The second obligation says that, if i does not concede two resources, then (from j's point of view) he should concede one resource.

Like in the selfish case, if the players go for the second-best, and accept to relax their altruistic principles, then the outcome of the game becomes the optimal one, namely w_4.

6 Conclusion

In this paper we have maintained that assigning a game-theoretical semantics to contrary-to-duty obligations considerably enriches the span of possible applications of their logics. Specifically, we have seen how reasoning on contrary-to-duty obligations can be seen as reasoning on second best choices in interaction. This seems to make perfect sense when players are confronted with coordination problems when fully individualistic or fully altruistic solutions fail. In these scenarios intermediate concepts should be sought, and contrary-to-duty logics offer a flexible framework to carry it out. Our approach is fully in line with the utilitarian treatment of deontic operators as logical notions that can be used to reason about rational choices in interaction. The finding that the second-best is sometimes better than the best departs from the conventional wisdom. In this respect, our approach is somewhat orthogonal with the usual treatments of CTDs. We believe such a finding is a fruitful avenue for future research.

Several directions for future work can be taken. Above all, an axiomatization of the logic would be desirable. Techniques have been developed in order to resolve the axiomatization problem of dyadic deontic logic [17, 6, 13]. It remains to investigate whether such techniques can be adapted to a game-theoretical setting such as the one proposed here.

Afterword

Thanks to Marek Sergot for his leading role in the DEON community. We acknowledge his influence on our work in general, and this paper in particular. His work with Henry Prakken on contrary-to-duty obligations in defeasible deontic logic determined the PhD research questions of the third author in the mid nineties. Marek's interest in contrary-to-duty obligations goes back to his joint work with Andrew Jones, since from the first DEON workshop in 1991 they have emphasized the distinction between what ought to be the case and what is the case, or as they call it, between the actual and the ideal. This has become the standard criterion for deciding whether deontic logic can or should be used in computer science. The present paper comments on this dogma by raising the fundamental question what it means in a game-theoretic context to be sub-ideal.

Acknowledgements

Silvano Colombo Tosatto is supported by the National Research Fund, Luxembourg. Paolo Turrini acknowledges the support of the National Research Fund of Luxembourg for the Trust Games project (1196394), cofunded under the Marie Curie Actions of the European Commission (FP7-COFUND).

References

1. J. Abdou. Rectangularity and tightness: A normal form characterization of perfect information extensive game forms. *Mathematics of Operations Research*, 3(23):553–567, 1998.
2. N. Belnap, M. Perloff, and M. Xu. *Facing The Future: Agents And Choices In Our Indeterminist World.* Oxford University Press, Usa, 2001.

3. J. Broersen, A. Herzig, and N. Troquard. A normal simulation of coalition logic and an epistemic extension. In Dov Samet, editor, *Proceedings Theoretical Aspects Rationality and Knowledge (TARK XI), Brussels*, pages 92–101. ACM Digital Library, 2007.
4. J. Broersen, R. Mastop, J.J. Ch. Meyer, and P. Turrini. A deontic logic for socially optimal norms. In Ron van der Meyden and Leendert van der Torre, editors, *DEON*, volume 5076 of *Lecture Notes in Computer Science*, pages 218–232. Springer, 2008.
5. J. Carmo and A.J.I. Jones. Deontic logic and contrary-to-duties. In D. M. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume 8, pages 265–344. Kluwer Academic Publishers, Dordrecht, Holland, 2nd edition, 2002.
6. L. Goble. Preference semantics for deontic logics. Part I: Simple models. *Logique & Analyse*, 46(183-184):383–418, 2003.
7. B. Hansson. An analysis of some deontic logics. *Noûs*, 3:373–398, 1969. Reprinted in [8, pp 121-147].
8. R. Hilpinen, editor. *Deontic Logic: Introductory and Systematic Readings*. Reidel, Dordrecht, 1971.
9. J. Horty. *Agency and Deontic Logic*. Oxford University Press, 2001.
10. B. Kooi and A. Tamminga. Moral conflicts between groups of agents. *Journal of Philosophical Logic*, 37(1):1–21, 2008.
11. D.K. Lewis. *Counterfactuals*. Blackwell, Oxford, 1973.
12. M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
13. X. Parent. On the strong completeness of Åqvist's dyadic deontic logic G. In R. van der Meyden and L. van der Torre, editors, *Deontic Logic in Computer Science* 9th International Conference, DEON 2008, Luxembourg, July 15-18, 2008. Proceedings, number 5076 in LNAI, pages 189–202. Springer, Berlin Heidelberg, 2008.
14. H. Prakken and M. Sergot. Dyadic deontic logic and contrary-to-duty obligation. In D. Nute, editor, *Defeasible Deontic Logic*, pages 223–262. Kluwer Academic Publishers, Dordrecht, 1997.
15. H. Prakken and M. Sergotk. Contrary-to-duty obligations. *Studia Logica: An International Journal for Symbolic Logic*, 57(1):pp. 91–115, 1996.
16. Y. Shoham. *Reaoning about changes*. MIT, 1988.
17. W. Spohn. An analysis of Hansson's dyadic deontic logic. *Journal of Philosophical Logic*, 4(2):237–252, 1975.
18. P. Turrini. *Strategic Reasoning in Interdependence: Logical and Game-theoretical Investigations*. SIKS Dissertation Series, 2011. PhD Thesis.
19. P. Turrini, D. Grossi, J. Broersen, and J.-J. Ch. Meyer. Forbidding undesirable agreements: A dependence-based approach to the regulation of multi-agent systems. In Guido Governatori and Giovanni Sartor, editors, *DEON*, volume 6181 of *Lecture Notes in Computer Science*, pages 306–322. Springer, 2010.