A Modal Representation of Strategic Reasoning

Jan Broersen

Rosja Mastop

John-Jules Meyer

Paolo Turrini

Technical Report UU-CS-2009-014
July 2009

Department of Information and Computing Sciences
Utrecht University, Utrecht, The Netherlands
www.cs.uu.nl

ISSN: 0924-3275

Department of Information and Computing Sciences
Utrecht University
P.O. Box 80.089
3508 TB Utrecht
The Netherlands

A Modal Representation of Strategic Reasoning

Jan Broersen, Rosja Mastop, John-Jules Meyer, Paolo Turrini

Utrecht University, paolo@cs.uu.nl

Abstract

Strategic reasoning representation is a key issue in the theoretical study and the implementation of Multi Agent System. To this purpose a language is developed that can explicitly describe the dynamics of strategies, the preferences and their interaction. The relation with Pauly's Coalition Logic is studied and a full reduction is shown, together with the characterization of gametheoretical concept such as undominated choice.

1 Introduction

In their latest foundational book on Multi Agent Systems [12] Shoham and Leyton-Brown address the study of the field's broad umbrella by means of five keywords: "coordination, competition, algorithms, game theory and logic". Stemming from this view it is natural to maintain that strategic reasoning representation - that is finding a suitable mathematical language to describe the inferences of interactive decision makers - is a key issue in the theoretical study and the implementation of Multi Agent System. Interactive decision makers that reason on the best action to select, taking into account the other players and the environment, are object of study of Game Theory, that constitutes a well founded model for Multi Agent Systems [12].

Game Theory deals with solution concepts, as for instance that of Nash Equilibrium [10], in which players reason on the possible reactions of their opponents and choose the best strategy given such reactions. To represent strategic reasoning a language is then needed that is able to describe the strategies and the preferences of the agents, together with their dynamics. However the present languages that talk about gametheoretical interaction such as ATL [2], Coalition Logic [11] or STIT [7] do not explicitly represent preferences and only allow to reason on what a coalition of agents can achieve independently of the moves of the other players [7, 8, 13, 14]. As pointed out in [14], p.1:

> Much of game theory is about the question whether strategic equilibria exist. But there are hardly any explicit languages for defining, comparing, or combining strategies as such - the way we have them for actions and plans, maybe the closest intuitive analogue to strategies. True, there are many current logics for describing game structure - but these tend to have existential quantifiers saying that "players have a strategy" for achieving some purpose, while descriptions of these strategies themselves are not part of the logical language.

To tackle this problem, we extend Coalition Logic, the relation of which with strategic games is clarified by Pauly Representation Theorem [11], with an operator that is interpreted

i \ j	C	D
C	$(3,3)$	$(0,4)$
D	$(4,0)$	$(1,1)$

Table 1: A Prisoner Dilemma

in the game restriction, or *update*, induced by the choice of a coalition. Players can select their best responses considering such restrictions. Both strategies and preferences are made explicit in the language, that can be used to represent strategic reasoning.

A further issue concerns implementation. Whenever a language is made for computational purposes the issue of complexity of representation need to be taken into account. A language that is proposed for being a suitable extension of already present languages should not be much more complicated to implement than the languages it is extending. We address this issue showing a full reduction of a language for strategic ability update to the language of coalitional interaction.

Updates are not new to the realm of logics. Formalizations of dynamics of information flow, like Dynamic Epistemic Logic [16] (DEL), reason about how agents' knowledge is updated after an epistemic event, for instance a public announcement, takes place. The idea of this paper is to extend the *update paradigm* of public announcements to account for the changes that moves in a game induce to players' strategic decisions.

1.1 Illustrative Example

For an intuition, let us consider a Prisoner Dilemma [10], that is an interactive situation in which the advantages of cooperation are overruled by the incentive for individual players to defect. In Table 1 a Prisoner Dilemma is described, where players i and j, that we assume to be rational, can choose between a cooperative move C and a defective move D, yielding an outcome (x_i, x_j), x_k being the payoff for each $k \in \{i, j\}$. If we focus on player i we can observe that, after the choice C by j, the choice D becomes preferable to the choice C - yielding $(4, 0)$ instead of $(3, 3)$ - and the same holds in case j moved D - yielding $(1, 1)$ instead of $(0, 4)$. Our rationality assumption warrants player i to reason on the updates of his own choices brought about by player j, and to select his best response in each such update. We call *undominated* a choice that remains a best response for all possible reactions of one's opponents.

Our aim is to formally capture the reasoning structure of players in strategic interaction. To do so we will provide a semantics for the notion of undominated choice, seen as an optimal solution in each game restriction induced by the moves of the players. We will work on cooperative structures, where players can form coalitions to achieve their goals [3]. The paper is structured as follows: The first part of the paper introduces Coalition Logic, that we use to model strategic ability; in the second part we introduce an operator to talk about the model transformations induced by the choices of coalitions: the subgame operator. In the third part we give a semantics to preferences and combine them to the subgame operator to characterize the notion of undominated choice. Finally we give reduction axioms for the full language and discuss its links with Public Announcement Logic.

2 Coalition Logic and Strategic Ability

In Game Theory players may be able to force the interaction to end up in an outcome satisfying certain properties. An abstract representation of this notion is given by the dynamic effectivity function, first described in [11], which we adopt to model strategic ability.

Definition 1 (Dynamic Effectivity Function)
Given a finite set of agents Agt and a set of states W, a dynamic effectivity function *is a function $E : W \to (2^{Agt} \to 2^{2^W})$.*

Any subset of Agt will henceforth be called a *coalition*. The elements of W are called *states* or *worlds*; the sets of states $X \in E(w)(C)$ are called the *choices* of coalition C in state w. The set $E(w)(C)$ is called the *choice set* of C in w. The complement of a set X is indicated as $\overline{X}$ and calculated relative to the expected domain. A dynamic effectivity function can be seen as a "formal description of the power structure in a society" [1]; it assigns, in each world, to every coalition a set of sets of states that represents the strategic ability of that coalition. Intuitively, if $X \in E(w)(C)$, C is said to be able from w to *force* the interaction to end up in some member of X. Every effectivity function has the property of **outcome monotonicity**: for all $X \subseteq W, Y \subseteq W, w \in W, C \in 2^{Agt}$, if $X \in E(w)(C)$ and $X \subseteq Y$, then $Y \in E(w)(C)$. Said in other words, if a coalition is able to force the the interaction to end up in some member of X then is also able to force the interaction to end up in some member of any supersets of X. Together with outcome monotonicity we will assume the properties of **regularity**: if $X \in E(w)(C)$, then $\overline{X} \notin E(w)(\overline{C})$; and **closed-worldness**: $E(w)(\emptyset) = \{W\}$. Regularity means that disjoint coalitions do not make choices that contradict each other, while closed-worldness requires the empty coalition not to influence the interaction. For an in depth discussion on the desirability of these properties see the results in [5].

2.1 Models and Language

The models we refer to are structures of the form

$$\langle W, E, V \rangle$$

where W is a nonempty set of states, E an outcome monotonic, regular and closed-world effectivity function, $V : W \to 2^P$ a valuation function that assigns to each state a subset of a countable set of atomic propositions P, to be interpreted as true at that state. The formulas for the basic language are of the form

$$p|\neg\phi|\phi \wedge \psi|[C]\phi|A\phi$$

where p is any atomic proposition in P, $[C]\phi$ is the coalitional operator expressing the fact that coalition C can force or bring about the formula ϕ; $A\phi$ is the global modality, which talks about a formula that holds in every world in the model. Their interpretation is standard [11, 4, 13] and it is given as follows:

$$\begin{array}{rcl} M, w \models p & \text{iff} & p \in V(w) \\ M, w \models \neg\phi & \text{iff} & \text{not } M, w \models \phi \\ M, w \models \phi \wedge \psi & \text{iff} & M, w \models \phi \text{ and } M, w \models \psi \\ M, w \models [C]\phi & \text{iff} & \phi^M \in E(w)(C) \\ M, w \models A\phi & \text{iff} & M, v \models \phi, \text{ for all } v \in W \end{array}$$

where $\phi^M = \{w \in W | M, w \models \phi\}$ is the *truth set* of ϕ. We write $\langle C \rangle \phi$ as an abbreviation for $\neg[C]\neg\phi$.

What we can say in Coalition Logic The Prisoner Dilemma can intuitively be rewritten as a coalition model. Here coalition $\{i\}$ can force that $\{i\}$ defects and can force that $\{i\}$ cooperates, but $\{i\}$ cannot force that $\{j\}$ cooperates (and equally it cannot force that $\{j\}$ defects). In any world w, we have therefore that $PD, w \models [\{i\}](\text{ i defects }) \wedge \neg[\{i\}](\text{ j defects })$. On the other hand we cannot express what i can do given that j defects. This would mean that i would have a strategy forcing that i defects and j defects and a strategy forcing that i cooperates and j defects. This at the model level is $PD, w \models [\{i\}](\text{ i defects and j defects }) \wedge [\{i\}](\text{ i cooperates and j defects })$. By the property of outcome monotonicity, we would then get $PD, w \models [\{i\}](\text{ j defects })$, which is at odds with our initial statement. The reason of this limitation is to be found in the interpretation of the coalition logic operator, that expresses what a coalition can achieve *independently* of what its opponents do. Reasoning on how the strategic ability of a coalition is updated by the possible moves of its opponents requires a different interpretation, namely what a coalition can achieve *given* what its opponents do.

3 Strategic Ability Update

To model strategic ability update we construct an operator $[C \downarrow \psi]\phi$ the informal reading of which is: "after coalition C chooses ψ, ϕ holds". We define the dual $\langle C \downarrow \psi \rangle \phi$ as an abbreviation of $\neg[C \downarrow \psi]\neg\phi$. Intuitively what we do is to talk about the model *restrictions* that are caused by the possible move ψ of coalition C. For this reason it will be called *the subgame operator*. Its formal interpretation goes as follows:

$$M, w \models [C \downarrow \psi]\phi \Leftrightarrow \psi^M \in E(w)(C) \text{ implies } M \downarrow_{(C,\psi^M,w)}, w \models \phi$$

The operator has a conditional reading: if a coalition C has a certain choice ψ^M at w, then the state where this choice is actually executed makes a certain proposition ϕ true. The capacity of C to choose ψ^M is seen here as a precondition for C to actually execute ψ^M.

The restricted models $M \downarrow_{(C,\psi^M,w)}$ are so defined:

$$M \downarrow_{(C,\psi^M,w)} \doteq \langle W, E \downarrow_{(C,\psi^M,w)}, V \rangle$$

They inherit the domain and the valuation function from the original coalition model while they update the coalitional relation [1] $E \downarrow_{(C,\psi^M,w)}$ in the following way:

$$\begin{array}{ll} E \downarrow_{(C,\psi^M,w)} (w)(D) \doteq (\{\psi^M\})^{\text{sup}} & \text{for } D \cap C \neq \emptyset \\ E \downarrow_{(C,\psi^M,w)} (w)(D) \doteq (E(w)(D) \sqcap \psi^M)^{\text{sup}} & \text{for } D \cap C = \emptyset \text{ and } D \neq \emptyset \\ E \downarrow_{(C,\psi^M,w)} (w')(D) \doteq E(w')(D) & \text{for } w' \neq w \text{ or } D = \emptyset \end{array}$$

where for a set of sets $\mathcal{X}$, $(\mathcal{X})^{\text{sup}} = \{X \subseteq W | \text{ there is } Y \in \mathcal{X} \text{ and } Y \subseteq X \subseteq W\}$. In words, $()^{\text{sup}}$ is the superset closure of a set of sets. Moreover taken two sets of sets $\mathcal{X}, \mathcal{Y}$, $\mathcal{X} \sqcap \mathcal{Y} = \{X \cap Y | X \in \mathcal{X} \text{ and } Y \in \mathcal{Y}\}$.

[1] Here the word *functional relation* would be more appropriate. In fact the Effectivity Function behaves as a relation in a Neighbourhood model and our restriction uniquely associates to an Effectivity Function the restriction imposed by a coalitional choice.

The way the relation is updated deserves some comment. A distinction is made between the strategic ability update of the players who made a certain choice ϕ and all the other players. After coalition C has made a choice ϕ, all the coalitions involving agents belonging to C are given $(\phi^M)^{\text{sup}}$ as a choice set. This view maintains that a coalition comprising players that have moved cannot further influence the outcome of the game. The models of reference are strategic games, in which strategies are decided in the beginning once and for all [10]. The other (nonempty) coalitions instead *truly update* their choice set having it restricted by the choice of C. Restriction is implemented in this case by intersecting the effectivity function with the move that has been carried out. If for instance C chooses ψ and $\overline{C}$ was able to choose ξ, given the choice by C, $\overline{C}$ is able to choose $\xi \wedge \psi$. The coalitional relation at worlds different from the one where the choice is made remains instead unchanged. This means that the update is local. Again, the references are strategic games, where the sequential structure of strategies is substantially ignored. Notice that by the last condition the empty coalition never gains power. In sum the strategic ability update is governed by three principles: the **irrelevance of hybrid coalitions**, that does not allow members of the coalition that moved to further influence the interaction, the **restriction of opponents' choices**, that truly updates the effectivity function of the coalitions opposing the one that moved, and the **locality of the update**, that leaves the coalitional power at different worlds untouched.

The following relevant fact can be easily verified:

Proposition 1 *For every C,w, $\psi^M \in E(w)(C)$, we have that $E\downarrow_{(C,\psi^M,w)}$ is outcome monotonic, regular and closed-world.*

The proposition represents the basis for our reduction results. Whatever update is carried out a model is obtained that obeys the properties that have been assumed for coalition models.

Even though the interpretation of the update operator may look complex, its structural behaviour is rather simple. The validities in Table 2 allow us to translate every sentence where the operator is occurring to a sentence where the operator is not occurring, provided an appropriate law for substitution of equivalent formulas (as $R5$ in the proof system). The resemblance to Public Announcement Logic is no coincidence. The axioms in fact reduce the update operator to the global modality and the coalition logic operator [2]. So the operator adds no expressivity to the language and completeness of the language with the update operator follows from the completeness of the language without it. A completeness proof for Closed-World coalition logic, where the global modality interacts with the coalition logic modality by means of the axiom $[\emptyset]\phi \leftrightarrow A\phi$ is provided in [5].

3.1 Back to the game

With the new operator it becomes possible to formalize the conditional aspect of strategic reasoning. In the structure PD we have that $PD, w \models [\{i\} \downarrow \text{ i defects }]([\{j\}](\text{ j defects and i defects }) \wedge [\{j\}](\text{ j cooperates and i defects }))$. Nothing changes at the level of grand coalition, since $PD \models [\emptyset \downarrow \phi][Agt]\psi \leftrightarrow [Agt]\psi$.

[2]Axiom A11 shows that the same holds when a preference operator is added to the language, as needed in the coming sections.

	Axioms
	Regularity
A1	$[C]\phi \to \neg[\overline{C}]\neg\phi$
	Closed-Worldness
A2	$[\emptyset]\phi \leftrightarrow A\phi$
	Global Modality Axioms
A3	$\phi \to E\phi$
A4	$EE\phi \to E\phi$
A5	$\phi \to AE\phi$
A6	$A(\phi \to \psi) \to (A\phi \to A\psi)$
	Strategic Ability Update Axioms
A7	$[C \downarrow \xi]p \leftrightarrow ([C]\xi \to p)$
A8	$[C \downarrow \xi]\neg\phi \leftrightarrow ([C]\xi \to \neg[C \downarrow \xi]\phi)$
A9	$[C \downarrow \xi](\phi \wedge \psi) \leftrightarrow ([C \downarrow \xi]\phi \wedge [C \downarrow \xi]\psi)$
A10	$[C \downarrow \xi]A\phi \leftrightarrow ([C]\xi \to A\phi)$
A11	$[C \downarrow \xi]\Box_i^{\leq}\phi \leftrightarrow ([C]\xi \to \Box_i^{\leq}\phi)$
A12	$[C \downarrow \xi][D]\phi \leftrightarrow ([C]\xi \to [D](\xi \to \phi))$ (for $D \cap C = \emptyset$ and $D \neq \emptyset$)
A13	$[C \downarrow \xi][D]\phi \leftrightarrow A(\xi \to \phi)$ (for $D \cap C \neq \emptyset$)
A14	$[C \downarrow \xi][D]\phi \leftrightarrow ([C]\xi \to [D]\phi)$ (for $D = \emptyset$)
	Preference Axioms
A15	$\phi \to \Diamond_i^{\leq}\phi$
A16	$\Diamond_i^{\leq}\Diamond_i^{\leq}\phi \to \Diamond_i^{\leq}\phi$
A17	$\Diamond_i^{\leq}\phi \to E\phi$
A18	$(\phi \wedge \Box_i^{\leq}\psi) \to A(\psi \vee \phi \vee \Diamond_i^{\leq}\phi)$
	Rules
R1	$\phi \wedge (\phi \to \psi) \Rightarrow \psi$
R2	$\phi \to \psi \Rightarrow [C]\phi \to [C]\psi$
R3	$\phi \Rightarrow A\phi$
R4	$\phi \Rightarrow [C \downarrow \xi]\phi$
R5	$\phi \leftrightarrow \psi \Rightarrow [C \downarrow \xi]\chi \leftrightarrow [C \downarrow \xi]\chi[\phi/\psi]$

Table 2: Proof System

4 Characterizing Undomination

Recall that in the Prisoner Dilemma, D is an undominated choice for each agent because, whatever its opponent does, D remains the best possible choice. In the previous part of the paper we have given an explicit representation of the expression *whatever the opponent does* by introducing the subgame operator. However to fully model undominated choices we need as well to give a precise semantics to the notion of *best possible choice.*

In line with a well established gametheoretical framework [10], we will assume a preference ordering $(\geq_i)_{i\in Agt}$ to be a weak linear order (reflexive, transitive, trichotomous) [3] $\geq_i \subseteq W \times W$ for each $i \in Agt$. When two states v, w are in the relation $v \geq_i w$ we say that v is 'at least as nice' as w for agent i. The corresponding strict order is defined as usual: $v >_i w$ if, and only if, $v \geq_i w$ and not $w \geq_i v$. The duals $w \leq_i v$, $w <_i v$ are defined in the expected way.

Preferences can be dealt with in modal logic. The standard operator (for a discussion see for instance [9]) is interpreted as follows:

$$M, w \models \Diamond_i^{\leq}\phi \quad \text{iff} \quad M, w' \models \phi, \text{ for some } w' \text{ with } w \leq_i w'$$

The dual $\Box_i^{\leq}\phi$ is an abbreviation for $\neg\Diamond_i^{\leq}\neg\phi$. $\Diamond_i^{\leq}\phi$ tells that from a given situation there is a world that is at least as nice as the present one for agent i and that makes ϕ true.

Even though preferences are defined as an ordering on states, in interactive situations agents are confronted with choices, that are here modelled as sets of states. Preferences over choices can be retrieved from the preference over states. To this purpose, we lift the ordering on states to an ordering on sets of states by means of the following principle:

$$X \geq_i Y \text{ iff for all } x \in X, y \in Y, x \geq_i y$$

For the strict ordering $>_i$ we obtain the lifting substituting every occurrence of $\geq_i$ in the previous definition with $>_i$. The idea is that if an agent were ever confronted with two choices X, Y he would not choose X over Y whenever $Y >_i X$.

In our example both $\{(\text{i defects and j cooperates})\}^{PD}$, i.e. all the worlds in which i defects and j cooperates, and $\{(\text{i defects and j defects})\}^{PD}$, i.e. all the worlds in which i and j defect, belong to $E\downarrow_{(i,\{(\text{i defects})\}^{PD},w)}(w)(j)$, the alternatives left to j once i decides to defect. Because of the preferences of j, we have that $\{(\text{i defects and j defects})\}^{PD} >_j \{(\text{i defects and j cooperates})\}^{PD}$, that is given the defective move by i, j strictly prefers a defective move to a cooperative move. However $\{(\text{j defects})\}^{PD} >_j \{(\text{j cooperates})\}^{PD}$ (abbreviated as $D(j) >_j C(j)$) is not true in general.

In Game Theory, to talk about situations that are preferable to any other situation, the notion of Pareto Optimality is often used. Pareto Optimality selects the maxima in a given ordering of states. A state x is Pareto Optimal iff for no state y, $y >_i x$ for all agents i [4]. In the same spirit of the lifting from states to sets of states, we generalize this definition to what we call Pareto Optimal Choice (in short POC), that selects the maxima in a given ordering of choices.

[3] A relation R is trichotomous (or weakly connected) if for all elements x, y it holds that $(xRy \vee yRx \vee y = x)$. Notice that if R is reflexive and trichotomous then R is connected, that is for all elements x, y it holds that $xRy \vee yRx$. Trichotomy is definable by a global modality, that we denote with A as usual [4] by means of the schema $(p \wedge \Box q) \rightarrow A(q \vee p \vee \Diamond p)$.

[4] For the sake of precision, the present definition is known in the literature as *Weak Pareto Optimality* [10], whilst the *Strong Pareto Optimality* holds iff for no state y, $y >_i x$ for all agents i, and $y \geq_i x$ for some. We make only use of the weak version and we call it Pareto Optimality for simplicity.

Definition 2 (Pareto Optimal Choice) *Given a choice set $E(w)(C)$, $X \in E(w)(C)$ is* Pareto Optimal *for coalition C (abbr. POC_C) in w if, and only if, for no $Y \in E(w)(C)$, $Y >_i X$ for all $i \in C$.*

Pareto Optimal Choices can be characterized combining the coalition logic and the preference operator.

Proposition 2 *ϕ^M is Pareto Optimal Choice for C in w iff $M, w \models [C]\phi \wedge \langle C \rangle \bigvee_{i \in C} \Diamond_i^{\leq} \phi$*

($\Rightarrow$) ϕ^M is Pareto Optimal Choice for C in w whenever for no $X \in E(w)(C)$, $X >_i \phi^M$ for all $i \in C$ and that $\phi^M \in E(w)(C)$. $X >_i \phi^M$ means that for all $x \in X$ and $y \in \phi^M$ we have that $x \geq_i y$ and not $y \geq_i x$ for all $i \in Agt$. The whole sentence means that $\exists x \in X, \exists y \in \phi^M$ such that it is not the case that $x \geq_i y$ or it is the case that $x \leq_i y$. As we pointed out before, trichotomy and reflexivity imply connectedness. So we can safely conclude that $x \leq_i y$. That is no set $X \in E(w)(C)$ is such that $X \subseteq \neg \Diamond_i^{\leq} \phi$ for some agent $i \in C$. So $M, w \models [C]\phi \wedge \langle C \rangle \bigvee_{i \in C} \Diamond_i^{\leq} \phi$.

($\Leftarrow$) $M, w \models [C]\phi \wedge \langle C \rangle \bigvee_{i \in C} \Diamond_i^{\leq} \phi$ means that $\phi^M \in E(w)(C)$ and $(\neg \bigvee_{i \in C} \Diamond_i^{\leq} \phi)^M \notin E(w)(C)$. So, by outcome monotonicity, every $X \in E(w)(C)$ has a world x such that $x \models \bigvee_{i \in C} \Diamond_i^{\leq} \phi$, so that $x \leq_i y$ for some $y \in \phi^M$ and $i \in C$. So for no $X \in E(w)(C)$, $X >_i \phi^M$ for all $i \in C$. Q.E.D.

Pareto Optimal Choices always exist. Trivially, being $W \in E(w)(C)$, for every w, C, we can never have that $X >_i W$, for $i \in C$ and $X \subseteq W$. Nevertheless, often there are various Pareto Optimal Choices, apart from W. In the Prisoner Dilemma, both $D(k)$ and $C(k)$ are Pareto Optimal Choices for each player $k \in \{i, j\}$. This suggests that the mere use of such concept cannot provide a good characterization of undominated choice. A further reason is that Pareto Optimal Choice does not make reference to the preferences or the abilities of the opponents. In this sense it lacks a strategic dimension. However the combination of subgame operator and POC can define the notion of undomination.

4.1 Undomination

We call a choice *undominated* if it is Pareto Optimal no matter what the others decide to do. This is the formal definition:

Definition 3 (Undomination) *Given an effectivity function E, ϕ^M is* undominated *for C in w (abbr. $\phi^M \rhd_{C,w}$) iff $\phi^M \in E(w)(C)$ and for all $\psi^M \in E(w)(\overline{C})$, $(\phi^M \cap \psi^M)$ is Pareto Optimal in $E^{\psi^M}(w)(C)$ for C.*

$E^{\psi^M}(w)(C)$ is an abbreviation for $E(w)(C) \sqcap \psi^M$. Pareto Optimality in $E^{\psi^M}(w)(C)$ is defined in the expected way.

As emphasized before, $\{$(i defects and j defects)$\}^{PD}$ is Pareto Optimal for i in $E^{D(j)}(w)(i)$ and $\{$(i defects and j cooperates)$\}^{PD}$ is Pareto Optimal for i in $E^{C(j)}(w)(i)$. This means that D is an undominated choice for i in w. This is not true for C, since $\{$(i cooperates and j defects)$\}^{PD}$ is *not* Pareto Optimal for i in $E^{D(j)}(w)(i)$.

Undomination can be characterized within the structures we have so far defined. The first characterization will be carried out assuming that *every coalition has only a finite amount of choices that are the truth set of some proposition.* We will remove this assumption when studying undomination as a frame condition.

Proposition 3 *Given* $\{\psi_1, ..., \psi_n\} = E(w)(\overline{C})$,
$\phi^M \rhd_{C,w} \Leftrightarrow M, w \models \bigwedge_{\psi_i \in \{\psi_1,...,\psi_n\}} [\overline{C} \downarrow \psi_i] POC_C(\phi \wedge \psi_i)$

($\Leftarrow$) Trivial.

($\Rightarrow$) $\phi^M \rhd_{C,w}$ means that $\phi^M \in E(w)(C)$ and for all $\psi^M \in E(w)(\overline{C})$, $\phi^M \cap \psi^M$ is Pareto Optimal in E^{ψ^M} for C in w. Let us now observe that given a subgame $E^Y(w)(C)$ and a choice $X \in E^Y(w)(C)$, X is Pareto Optimal in $E^Y(w)(C)$ iff it is Pareto Optimal in $E^Y(w)(C)$ closed under supersets. So $\phi^M \cap \psi^M$ is Pareto Optimal in E^{ψ^M} for C in w closed under supersets. Notice that E^{ψ^M} for C in w closed under supersets is equivalent with $E \downarrow_{(\overline{C}, \psi^M, w)} (w)$ for C. By assumption we can finitely enumerate the choices of $\overline{C}$ that have a propositional form, that we call $\psi_1, ..., \psi_n$. We can conclude that $\phi^M \rhd_{C,w}$ means that for every $\psi_i \in \{\psi_1, ..., \psi_n\}$, $(\phi \wedge \psi_i)^M$ is Pareto Optimal in $E \downarrow_{(\overline{C}, \psi_i^M, w)} (w)$. This means $M, w \models \bigwedge_{\psi_i \in \{\psi_1,...,\psi_n\}} [\overline{C} \downarrow \psi_i] POC_C(\phi \wedge \psi_i)$. Q.E.D.

The following proposition holds when undomination is taken to be a frame condition.

Proposition 4 *For the class* $\mathcal{C}$ *of all frames based on the models described above, the axiom* $[C]\phi \rightarrow [\overline{C} \downarrow \xi] POC_C(\phi \wedge \xi)$ *determines the following condition:* $X \in E(w)(C)$ *implies that* $X \cap Y$ *is Pareto Optimal in* $E^Y(w)(C)$.

The proof is in the appendix.

The proposition allows for interesting observations. First of all, since we are characterizing undomination as a property of the frames, we do not need any restriction on the choices of coalitions. Another advantage of this characterization is that we can characterize a much finer notion of Undomination and Pareto Optimality: we can talk about all sets in an effectivity function, and not only those that are truth set of some proposition.

4.2 Back to the Game

The new language formalizes agents' reasoning in the Prisoner Dilemma. Its structure is a model PD with

$W = \{(3,3), (0,4), (4,0), (1,1)\}$,

$E(w)(\{i\}) = (\{\{(3,3)(0,4)\}, \{(4,0), (1,1)\}\})^{\text{sup}}$,

$E(w)(\{j\}) = (\{\{(3,3)(4,0)\}, \{(0,4), (1,1)\}\})^{\text{sup}}$,

$E(w)(\{i,j\}) = (\{(3,3), (0,4), (4,0), (1,1)\})^{\text{sup}}$ for $w \in W$. With a bit of sloppiness, we identify a formula with its truth set, not distinguishing $\{(3,3), (4,0)\}$ and ϕ^M s.t. $\phi^M = \{(3,3), (4,0)\}$. We have for every $w \in W$ that

$PD, w \models [j \downarrow \{(3,3)(4,0)\}] POC_i\{(4,0), (1,1)\}$ and
$PD, w \models [j \downarrow \{(0,4)(1,1)\}] POC_i\{(4,0), (1,1)\}$. By the previous characterizations we can conclude that $\{(4,0), (1,1)\} \rhd_{\{i\},w}$. For player j the situation is symmetric, while we have that $\{3,3\} \rhd_{\{i,j\},w}$: cooperation is socially rational but individually overruled by the incentive to defect.

5 Discussion: Choices as Announcements

Public Announcement Logic formalizes the effect of the announcement of a true formula in each agent's a epistemic relation $R(a)$, defined as a partition on a domain W. The standard operator $[\phi]\psi$ says that ψ holds after ϕ is announced. Its semantics is given as follows:

	Axioms
	Public Announcement Axioms
A1	$[\phi]p \leftrightarrow (\phi \rightarrow p)$
A2	$[\phi]\neg\psi \leftrightarrow (\phi \rightarrow \neg[\phi]\psi)$
A3	$[\phi](\xi \wedge \psi) \leftrightarrow ([\phi]\xi \wedge [\phi]\psi)$
A4	$[\phi]\Box_a\psi \leftrightarrow (\phi \rightarrow \Box_a[\phi]\psi)$
	Rules
R1	$\xi \wedge (\xi \rightarrow \psi) \Rightarrow \psi$
R2	$\xi \Rightarrow [\phi]\xi$

Table 3: Proof System for Public Announcement Logic

$$M, w \models [\phi]\psi \Leftrightarrow M, w \models \phi \text{ implies } M|\phi, w \models \psi$$

where $M|\phi = (W', R'(a), V')$ takes these values:

- $W' = \phi^M$
- $R'(a) = R(a) \cap (W \times \phi^M)$
- $V'(p) = V(p) \cap \phi^M$

The model restriction of public announcement consists in restricting the domain of worlds. Alternatively, as shown for instance in [15], the same effect can be achieved by restricting the epistemic relation. A reduction can be shown in which every sentence from the modal language with the $S5$ knowledge relation and the public announcement operator can be translated into a sentence from the same language without the public announcement operator occurring in it. We report the reduction axioms in Table 3.

If we compare the public announcement operator to the subgame operator, we can observe the structure of the two axiom systems is very similar in the atomic and boolean case, but very different in the modal case. The appendix will make it clear that the similarity applies to the proof techniques as well, that are at least for the basic cases identical to those of Public Announcement Logic [16]. The specific differences are given, once again, by the way the coalitional relation is updated.

6 Conclusion and future work

We have built a logic for strategic ability update, where we can represent the effects of a coalitional choice on the players' strategic ability, that, combined with a standard logic of preferences, allows for the characterization of gametheoretical notions like undominated choice. In the spirit of the well known update operators from Dynamic Epistemic Logic, our framework explicitly expresses how a coalitional move modifies the ability of all the players involved in the interaction. In the same fashion a reduction has been proved from the language with an update operator to the language without an update operator, providing a useful benchmark for the complexity of implementation. Our results are limited to Coalition Logic. Further study is needed to analyze whether the same characterizations are possible in different frameworks for strategic ability, for instance the Consequentialist-STIT framework, ATL and the full Game Logic. Further work can also be done in characterizing within this framework a number of other gametheoretical concepts like Nash Equilibrium and the Core for Cooperative Games without transferable utility.

References

[1] Joseph Abdou and Hans Keiding. *Effectivity Functions in Social Choice*. Kluwer Academic Publishers, 1991.

[2] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. In *FOCS '97: Proceedings of the 38th Annual Symposium on Foundations of Computer Science*, page 100, Washington, DC, USA, 1997. IEEE Computer Society.

[3] Robert J. Aumann and Bezalel Peleg. Von Neumann-Morgenstern solutions to cooperative games without side payments. *Bulletin of the American Mathematical Society*, 66:173–179, 1960.

[4] Patrick Blackburn, Martin de Rijke, and Yde Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science, 2001.

[5] Jan Broersen, Rosja Mastop, John-Jules Ch. Meyer, and Paolo Turrini. A Logic for Closed-World Interaction. In *Proceedings of the 11th European Conference on Logics in Artificial Intelligence (JELIA)*, pages 89–99, 2008.

[6] Brian Chellas. *Modal Logic: an Introduction*. Cambridge University Press, 1980.

[7] J. Horty. *Deontic Logic and Agency*. Oxford University Press, 2001.

[8] B. Kooi and A.Tamminga. Conflicting obligations in multi-agent deontic logic. In Lou Goble and John-Jules Ch. Meyer, editors, *8th International Workshop on Deontic Logic in Computer Science (DEON 2006)*, pages 175–186. LNCS 4048, 2006.

[9] F. Liu. *Changing for the Better: Preference Dynamics and Agent Diversity*. ILLC Dissertation Series, 2008.

[10] Martin Osborne and Ariel Rubinstein. *A course in Game Theory*. The MIT Press, 1994.

[11] Marc Pauly. *Logic for Social Software*. ILLC Dissertation Series, 2001.

[12] Yoav Shoham and Kevin Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic and Logical Foundations*. Cambridge University Press, 2008.

[13] Johan van Benthem. Where is logic going, and should it? In E.Bencivenga, editor, *What is to be Done in Philosophy? Topoi*, pages 117–122.

[14] Johan van Benthem. In praise of strategies. Research Report, http://www.illc.uva.nl/Publications/ResearchReports/PP-2008-03.text.pdf, 2007.

[15] Johan van Benthem and Fenrong Liu. Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logics*, 14, 2004.

[16] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Synthese Library, 2007.

A Proofs for Reduction Axioms

Atomic and Boolean Cases

$$[C \downarrow \xi]p \leftrightarrow ([C]\xi \rightarrow p)$$

Take arbitrary M, w. $M, w \models [C \downarrow \xi]p \Leftrightarrow M, w \models [C]\xi$ implies that $M \downarrow_{(C, \xi^M, w)}, w \models p \Leftrightarrow M, w \models [C]\xi$ implies that $M, w \models p \Leftrightarrow M, w \models [C]\xi \rightarrow p$. Q.E.D.

$$[C \downarrow \xi]\neg\phi \leftrightarrow ([C]\xi \rightarrow \neg[C \downarrow \xi]\phi)$$

Take arbitrary M, w. $M, w \models [C \downarrow \xi]\neg\phi \Leftrightarrow M, w \models [C]\xi$ implies that $M \downarrow_{(C,\xi^M,w)}, w \models \neg\phi \Leftrightarrow M, w \models [C]\xi$ implies that $(M, w \models [C]\xi$ and $M \downarrow_{(C,\xi^M,w)}, w \models \neg\phi) \Leftrightarrow M, w \models [C]\xi$ implies that not$(M, w \models [C]\xi$ implies $M \downarrow_{(C,\xi^M,w)}, w$ not $\models \neg\phi) \Leftrightarrow \models [C]\xi$ implies that not$(M, w \models [C]\xi$ implies $M \downarrow_{(C,\xi^M,w)}, w \models \phi) \Leftrightarrow M, w \models [C]\xi$ implies that M, w not $\models [C \downarrow \xi]\phi \Leftrightarrow M, w \models [C]\xi \rightarrow \neg[C \downarrow \xi]\phi$ Q.E.D.

$$[C \downarrow \xi](\phi \wedge \psi) \leftrightarrow ([C \downarrow \xi]\phi \wedge [C \downarrow \xi]\psi)$$

Take arbitrary M, w. $M, w \models [C \downarrow \xi](\phi \wedge \psi) \Leftrightarrow M, w \models [C]\xi$ implies that $M \downarrow_{(C,\xi^M,w)}, w \models \phi \wedge \psi \Leftrightarrow M, w \models [C]\xi$ implies that $(M \downarrow_{(C,\xi^M,w)}, w \models \phi$ and $M \downarrow_{(C,\xi^M,w)}, w \models \psi) \Leftrightarrow (M, w \models [C]\xi$ implies that $M \downarrow_{(C,\xi^M,w)}, w \models \phi)$ and $(M, w \models [C]\xi$ implies that $M \downarrow_{(C,\xi^M,w)}, w \models \psi) \Leftrightarrow (M, w \models [C \downarrow \xi]\phi)$ and $(M, w \models [C \downarrow \xi]\psi) \Leftrightarrow M, w \models ([C \downarrow \xi]\phi \wedge [C \downarrow \xi]\psi)$ Q.E.D.

Interaction with Global Modality

$$[C \downarrow \xi]A\phi \leftrightarrow ([C]\xi \rightarrow A\phi)$$

Take an arbitrary M, w. $M, w \models [C \downarrow \xi]A\phi \Leftrightarrow M, w \models [C]\xi$ implies that $M \downarrow_{(C,\xi^M,w)}, w \models A\phi \Leftrightarrow M, w \models [C]\xi$ implies that $M \downarrow_{(C,\xi^M,w)}, w \models [\emptyset]\phi \Leftrightarrow M, w \models [C]\xi$ implies that $M, w \models [\emptyset]\phi \Leftrightarrow M, w \models [C]\xi$ implies that $M, w \models A\phi \Leftrightarrow M, w \models [C]\xi \rightarrow A\phi$. Q.E.D.

Interaction with Preference Modality

$$[C \downarrow \xi]\Box_i^{\leq}\phi \leftrightarrow ([C]\xi \rightarrow \Box_i^{\leq}\phi)$$

Take an arbitrary M, w. $M, w \models [C \downarrow \xi]\Box_i^{\leq}\phi \Leftrightarrow M, w \models [C]\xi$ implies that $M \downarrow_{(C,\xi^M,w)}, w \models \Box_i^{\leq}\phi \Leftrightarrow M, w \models [C]\xi$ implies that $M \downarrow_{(C,\xi^M,w)}, v \models \phi$ for every v such that $w \leq_{i \in C} v \Leftrightarrow M, w \models [C]\xi$ implies that $M, v \models \phi$ for every v such that $w \leq_i v \Leftrightarrow M, w \models [C]\xi$ implies that $M, w \models \Box_i^{\leq}\phi \Leftrightarrow M, w \models ([C]\xi \rightarrow \Box_i^{\leq}\phi)$. Q.E.D.

Interaction with Coalition Modality

$$[C \downarrow \xi][D]\phi \leftrightarrow ([C]\xi \rightarrow [D](\xi \rightarrow \phi))(\text{ for } D \cap C = \emptyset \text{ and } D \neq \emptyset)$$

Proof by contraposition.

$\Leftarrow$: Suppose, for some $D \neq \emptyset$, that $M, w \models [C]\xi \rightarrow [D](\xi \rightarrow \phi)$ and M, w not $\models [C \downarrow \xi][D]\phi$ for some C such that $(C \cap D) = \emptyset$. The semantic clauses then tell us that if $\xi^M \in E(w)(C)$ then $(\xi \rightarrow \phi)^M \in E(w)(D)$ and that $\xi^M \in E(w)(C)$ and $\phi^M \notin E'(w)(D)$. [We write E' for $E \downarrow_{(C,\xi^M,w)}$.] By modus ponens we have that $\phi^M \notin E'(w)(D)$.

By the definition of update, $E'(w)(D) = (E(w)(D) \sqcap \xi^M)^{\text{sup}}$. So, $((\xi \rightarrow \phi)^M \cap \xi^M) \in E'(w)(D)$. By elementary set theory this just says that $\phi^M \in E'(w)(D)$. Contradiction.

$\Rightarrow$: Suppose, for some $D \neq \emptyset$, that $M, w \models [C \downarrow \xi][D]\phi$ and M, w not $\models [C]\xi \rightarrow [D](\xi \rightarrow \phi)$ for some C such that $(C \cap D) = \emptyset$. The semantic clauses then tell us that if $\xi^M \in E(w)(C)$

then $\phi^M \in E'(w)(D)$, and that $\xi^M \in E(w)(C)$ and $(\xi \to \phi)^M \notin E(w)(D)$. By modus ponens we have that $\phi^M \in E'(w)(D)$.

By the definition of update, $E'(w)(D) = (E(w)(D) \sqcap \xi^M)^{\text{sup}}$. Because $\phi^M \in E'(w)(D)$, there must be some $X \in E(w)(D)$, such that $(X \cap \xi^M) \subseteq \phi^M$. By elementary set theory we have that $(X \cap \xi^M) \cup (X \cap (\neg\xi)^M) \subseteq \phi^M \cup (X \cap (\neg\xi)^M) \subseteq \phi^M \cup (X \cap (\neg\xi)^M) \cup (\overline{X} \cap (\neg\xi)^M) = (\phi \vee \neg\xi)^M = (\xi \to \phi)^M$.

As $(X \cap \xi^M) \cup (X \cap (\neg\xi)^M) = X$, we have by outcome monotonicity of E that $(\xi \to \phi)^M \in E(w)(D)$. Contradiction. Q.E.D.

$$[C \downarrow \xi]([D]\phi \leftrightarrow A(\xi \to \phi))(\text{ for } D \cap C \neq \emptyset)$$

Proof. Take arbitrary M, w, and arbitrary $\xi^M \in E(w)(C)$. Consider a coalition D with $D \cap C \neq \emptyset$. We have that $E \downarrow_{(C,\xi^M,w)} (w)(D) = (\xi^M)^{\text{sup}}$ by semantics. This means that $\xi^M \subseteq \phi^M$ iff $\phi^M \in E \downarrow_{(C,\xi^M,w)} (w)(D)$. We conclude that $M, w \models [C \downarrow \xi]([D]\phi \leftrightarrow A(\xi \to \phi))$. Notice that this also means $M, w \models [C \downarrow \xi][D]\phi \leftrightarrow A(\xi \to \phi)$. Q.E.D.

$$[C \downarrow \xi][D]\phi \leftrightarrow ([C]\xi \to [D]\phi)(\text{ for } D = \emptyset)$$

It follows directly from the semantics of the update operator for the case of $D = \emptyset$. Q.E.D.

B Other Proofs

Proof of Proposition 4.

To prove this proposition we need to introduce the canonical model for the logic.

Definition 4 (Canonical Model) *A Canonical Model for our logic*

$$\mathcal{K} = \langle \mathcal{W}, \mathcal{E}, \preceq_i, \mathcal{R}_\forall, \mathcal{E} \downarrow_{(C,\psi^{\mathcal{K}},\omega)}, \mathcal{V} \rangle$$

consists of a coalitional relation $\mathcal{E}$, a preference relation $\preceq_i$, a global relation $\mathcal{R}_\forall$, the strategic ability update relation $\mathcal{E} \downarrow_{(C,\psi^{\mathcal{K}},\omega)}$ and a valuation function.

The domain $\mathcal{W}$ is made by all maximally consistent sets of formulas ω, where ω is a collection of formulas of the language, such that for any formula ϕ of the form p,$\neg\psi$,$[C]\psi$(for $C \in 2^{Agt}$), $\Diamond^{\leq_i}\psi$ (for $i \in Agt$), $A\psi$, $[D \downarrow \psi]$,(for $D \in 2^{Agt}$), either ϕ or $\neg\phi$ belongs to ω. ω is closed under the proof system depicted in Table 2.

For clarity reasons we explicitly introduce in the canonical model the global relation and the strategic ability update relation, the latest even if established to be definable from the previous relations. This is constructed associating to every maximal consistent set ω a maximal consistent set $\mathcal{K} \downarrow_{(C,|\phi|_{\mathcal{K}},\omega)}, \omega$, where the formulas of ω are updated by $[C \downarrow \phi]$.

For the scope of the proof it is convenient to make use of the smallest canonical model $\mathcal{K}$, that holds the following constraints

$$\mathcal{E}(\omega)(C) = \{|\phi|_{\mathcal{K}} | [C]\phi \in \omega\}$$
$$\mathcal{E} \downarrow_{(\overline{C},\psi^{\mathcal{K}},\omega)} (\omega)(C) = \{|\phi|_{\mathcal{K}} | [\overline{C} \downarrow \psi][C]\phi \in \omega\}$$

where $|\phi|_{\mathcal{K}}$ is the set of maximally consistent sets that contain ϕ.

Proof of Proposition 4.

($\Leftarrow$) Trivial.

(⇒) Take $X \in \mathcal{E}(\omega)(C)$ for some maximal consistent set ω of the smallest canonical model $\mathcal{K}$. By definition of $\mathcal{K}$, X is the proof set of some formula ϕ. So $|\phi|_{\mathcal{K}} \in \mathcal{E}(\omega)(C)$ and by the definition of the canonical relation $[C]\phi \in \omega$. By the Truth Lemma for Coalition Logic [6] [11] $\omega \models [C]\phi$. Since $\omega \models [C]\phi \rightarrow [\overline{C} \downarrow \xi]POC_C(\phi \wedge \xi)$, we can infer that $\omega \models [\overline{C} \downarrow \xi]POC_C(\phi \wedge \xi)$. Notice that this formula captures every choice Y belonging to $\mathcal{E}(\omega)(\overline{C})$, because Y can be described as the proof set of some proposition. Now given the axioms of the update operator we can construct a world $\mathcal{K} \downarrow_{(\overline{C},|\psi|_{\mathcal{K}},\omega)}, \omega$ such that $\mathcal{K} \downarrow_{(\overline{C},|\psi|_{\mathcal{K}},\omega)}, \omega \models POC_C(\phi \wedge \xi)$. This can be rewritten as $\mathcal{K} \downarrow_{(\overline{C},|\psi|_{\mathcal{K}},\omega)}, \omega \models [C](\phi \wedge \xi) \wedge \langle C \rangle \bigvee_{i \in C} \Diamond_i^{\leq}(\phi \wedge \xi)$. Given the axioms of the preference modality and that the fact that X, Y can be written as $X = |\phi|_{\mathcal{K}}$ and $Y = |\xi|_{\mathcal{K}}$, we have that for all $|\sigma|_{\mathcal{K}} \in \mathcal{E} \downarrow_{(\overline{C},|\psi|_{\mathcal{K}},\omega)} (\omega)(C)$ there is a world $\omega' \in |\sigma|_{\mathcal{K}}$ and a world $q \in X \cap Y$ such that $\omega' \leq_i q$ for some agent $i \in C$. So $X \cap Y$ is Pareto Optimal in $\mathcal{E} \downarrow_{(\overline{C},|\psi|_{\mathcal{K}},\omega)} (\omega)(C)$. Q.E.D.