
Logics for Coalitional Games

THOMAS ÅGOTNES, PAUL E. DUNNE, WIEBE VAN DER HOEK,
AND MICHAEL WOOLDRIDGE

ABSTRACT. We define formalisms to reason about Coalitional Games (CGs), in which one can express what coalitions of agents can achieve. We start with Quantified CGs (QCGs), in which each agent has some goals he wants to satisfy, which may change over time. Then we focus on CGs themselves. Although CGs can be well analysed in a formalism close to Pauly’s Coalition Logic, in QCGs, when having preferences, some differences become apparent.

1 Introduction

Recently, one has seen a shift in focus in the research of multi-agent systems from representing the cognitive structure of the agents, to logics that represent the strategic structure of multi-agent environments, and in particular, the powers that (groups of) agents have in such environments [Pau02]. Such logics have proved to have important applications, for example in the specification and verification of social choice mechanisms [Pau02]. One significant feature of these cooperation logics is that they have a close link with formal games: the semantic models underpinning Coalition Logic can be understood as extensive games of almost perfect information.

In this overview paper, we survey our work on logical characterisations of concepts from cooperative, or coalitional games. In a *Coalitional Game* [OR94, Part IV] each coalition C (i.e., set of agents) is assigned a value $v(C)$. Questions that naturally arise now are which coalitions will form, and whether such solutions are stable. *Qualitative Coalitional Games* were introduced in [WD04], as an abstract model of goal-oriented cooperative systems. In a QCG, each agent is assumed to have certain goals: an agent is “satisfied” with any outcome that accomplishes one of his goals, but is indifferent about which goal is satisfied.

This paper is a report on our following previous work, to which we will omit to refer in the next sections. After giving a brief introduction to Pauly’s Coalition Logic, in Section 2 we give a formal analysis of Quantified Coalitional Games (QCG’s). In Section 2.1 we follow [DvdHW07] and study

when coalitions can satisfy certain goals. In 2.2 (based on [ÅvdHW06b]) we shift the emphasis to the question when an agent is satisfied, and we add a temporal component to QCG's. In Section 3 then, we present a logic CGL for Coalitional Games, which derives from [ÅvdHW06a]. As a survey paper of our own work and due to space constraints, this paper is highly self-referential, and we know we don't do others justice in only a brief related work Section 4.

Coalition Logic. The logic we use as a starting point is known as *Coalition Logic* [Pau02]. It was introduced by Pauly as a framework for representing and reasoning about the powers of coalitions in game-like multi-agent encounters.

Informally, CL is a propositional modal logic, containing an indexed collection of unary modal operators $\langle C \rangle$, where C is a set of agents. The intended interpretation of a formula $\langle C \rangle \varphi$ is that the set of agents (coalition) C are *effective* for φ . That is, the agents C could cooperate to ensure that, in the next state of the environment, φ was true. We refer to an expression of the form $\langle C \rangle \varphi$ as a *coalition* or *cooperation* modality.

Syntactically, formulae φ of CL are defined over a set \mathcal{A} of agents and a set Φ_0 of atomic formulae by the Boolean connectives and the construct $\langle C \rangle \varphi$ with $C \subseteq \mathcal{A}$ a set of agents. Pauly [Pau02] uses $[C]$ where we use $\langle C \rangle$; here we use the latter notation for easier comparison.

Semantically, a *model*, \mathcal{M} , for CL is a quintuple: $\mathcal{M} = \langle \mathcal{A}, \mathcal{S}, \mathcal{E}, \Phi_0, v \rangle$, where:

- $\mathcal{A} = \{1, \dots, m\}$ is a finite, non-empty set of *agents* ;
- $\mathcal{S} = \{s_1, \dots, s_o\}$ is a finite, non-empty set of *states*;
- $\mathcal{E} : 2^{\mathcal{A}} \times \mathcal{S} \rightarrow 2^{2^{\mathcal{S}}}$ is an *effectivity function*, where $S \in \mathcal{E}(C, s)$ is intended to mean that from state s , the coalition C can cooperate to ensure that the next state will be a member of S ;
- Φ_0 is the set of propositional variables for \mathcal{M} ; and
- $v : \mathcal{S} \rightarrow 2^{\Phi_0}$ is a valuation function, which for every state $s \in \mathcal{S}$ gives the set $v(s)$ of propositional variables that are satisfied at s .

It is possible to define a number of constraints on effectivity functions. For the purposes of this paper, we shall assume just one property of effectivity functions: we require that the empty coalition has no power to do anything other than ensure that the model is closed, in the sense that the next state will be one of the defined possible states. Formally: $\mathcal{E}(\emptyset, s) = \{\mathcal{S}\}$, for all s .

An *interpretation* for CL is a pair \mathcal{M}, s , where \mathcal{M} is a model and s is a state in \mathcal{M} . The satisfaction relation " \models " for CL holds between interpretations and formulae of CL. The satisfaction relation has the following main clause: $\mathcal{M}, s \models \langle C \rangle \varphi$ iff $\exists S \in \mathcal{E}(C, s)$ such that $\forall s' \in S$, we have $\mathcal{M}, s' \models \varphi$.

Sometimes, when we fix the root of the interpretation, we also will write (\mathcal{M}, ρ) , in which cases it is implicitly assumed that $\rho \in \mathcal{S}$.

2 Qualitative Coalitional Games

We give a brief introduction to Qualitative Coalitional Games (QCGs): details may be found in [WD04]. A QCG contains a (non-empty, finite) set $\mathcal{A} = \{1, \dots, m\}$ of *agents*. Each agent $i \in \mathcal{A}$ is assumed to have associated with it a (finite) set \mathcal{G}_i of *goals*, drawn from a set of overall possible goals \mathcal{G} . The intended interpretation is that the members of \mathcal{G}_i represent all the individual rational outcomes for i – intuitively, the outcomes that give it “better than zero utility”. That is, agent i would be happy if *any* member of \mathcal{G}_i were achieved – then it has “gained something”. But, in QCGs, we are not concerned with preferences over individual goals. Thus, at this level of modelling, i is *indifferent* among the members of \mathcal{G}_i : it will be *satisfied* if *at least one* member of \mathcal{G}_i is achieved, and *unsatisfied* otherwise.

We assume that each possible coalition has available to it a set of possible *choices*, where each choice intuitively characterises the outcome of one way that the coalition could cooperate. We model the choices available to coalitions via a *characteristic function* with the signature $\mathcal{V} : 2^{\mathcal{A}} \rightarrow 2^{2^{\mathcal{G}}}$. Thus, in saying that $G \in \mathcal{V}(C)$ for some coalition $C \subseteq \mathcal{A}$, we are saying that one choice available to the coalition C is to bring about *exactly* the goals in G . At this point, the reader might expect to see some constraints placed on characteristic functions. For example, at first sight the following *monotonicity* constraint might seem natural: $C \subseteq C'$ implies $\mathcal{V}(C) \subseteq \mathcal{V}(C')$. Although such a constraint is entirely appropriate for many scenarios, there are cases where such a constraint is not appropriate¹.

Bringing these components together, a qualitative coalitional game (QCG) is a tuple: $\Gamma = \langle \mathcal{A}, \mathcal{G}, \mathcal{G}_1, \dots, \mathcal{G}_n, \mathcal{V} \rangle$ where

- \mathcal{A} is a finite, non-empty set of *agents*;
- \mathcal{G} is a finite, non-empty set of possible *goals*;
- $\mathcal{G}_i \subseteq \mathcal{G}$ is the set of goals for agent $i \in \mathcal{A}$; and
- $\mathcal{V} : 2^{\mathcal{A}} \rightarrow 2^{2^{\mathcal{G}}}$ is the characteristic function of the game.

EXAMPLE 1. Let Γ_1 be the following QCG for a collection of agents and a collection of goals $\{g_1, \dots\}$. Agent 1 is satisfied with g_1 and g_4 , and agent 2 is satisfied with g_2 and g_3 . The characteristic function is:

$$\begin{aligned} \mathcal{V}(C_1) &= \{ \{g_1, g_2\} \} & \mathcal{V}(C_2) &= \{ \{g_2, g_3\}, \{g_1\} \} \\ \mathcal{V}(C_3) &= \{ \{g_5, g_6\} \} & \mathcal{V}(C_4) &= \{ \{g_2, g_3\}, \{g_1\}, \{g_4\} \} \end{aligned}$$

¹For example, consider a legal scenario in which certain coalitions are forbidden by monopoly or anti-trust laws.

2.1 Goal Satisfaction

We now define a *correspondence* relation, “ \simeq ”, between QCGs and interpretations. The idea is that, for a QCG Γ and an interpretation \mathcal{M}, s , if $\Gamma \simeq \mathcal{M}, s$, then the QCG Γ and the interpretation \mathcal{M}, s are “equivalent” with respect to what they say about the way in which coalitions can cooperate. First, we say that QCG Γ and model \mathcal{M} are *comparable* iff:

1. The sets of agents in both structures are the same.
2. There is a propositional variable g in the model \mathcal{M} for every possible goal g in Γ , and \mathcal{M} contains no other propositional variables.

Hence, if a model $\mathcal{M} = \langle \mathcal{A}, \mathcal{S}, \mathcal{E}, \Phi_0, v \rangle$ and a game $\Gamma = \langle \mathcal{A}', \mathcal{G}, \mathcal{G}_1, \dots, \mathcal{G}_n, \mathcal{V} \rangle$ are comparable, then $\mathcal{A} = \mathcal{A}'$ and $\Phi_0 = \mathcal{G}$. As the reader may now be able to guess, the truth of a propositional variable g in a state s will be intended to mean that the corresponding goal g is achieved in state s .

In what follows, $G \subseteq \mathcal{G}$. Define

$$\pi_G^- \triangleq \bigwedge_{g \in G} \neg g, \quad \sigma_G^- \triangleq \bigvee_{g \in G} \neg g, \quad \pi_G^+ \triangleq \bigwedge_{g \in G} g, \quad \sigma_G^+ \triangleq \bigvee_{g \in G} g$$

So, if $\mathcal{M}, s \models \pi_G^-$, then this will mean that *no* goal in G is achieved in state s , whereas if $\mathcal{M}, s \models \pi_G^+$, then *every* goal in G is achieved in state s . In contrast, $\mathcal{M}, s \models \sigma_G^-$ means that *some* member of G is *not* achieved in s , while $\mathcal{M}, s \models \sigma_G^+$ will mean that some member of G *is* achieved in s .

Next, we define a formula that characterises exactly when a given set of goals is achieved in a given state: $\chi_G \triangleq \pi_G^+ \wedge \pi_{\mathcal{G} \setminus G}^-$. The following property is obvious. Let \mathcal{M} and Γ be a comparable model and game, and let $G \subseteq \mathcal{G}$; then:

$$(1) \quad \mathcal{M}, s \models \chi_G \quad \Leftrightarrow \quad v(s) = G$$

We can now define the correspondence relation. Let $\Gamma = \langle \mathcal{A}, \mathcal{G}, \mathcal{G}_1, \dots, \mathcal{G}_n, \mathcal{V} \rangle$ be a QCG game. We write $\Gamma \simeq (\mathcal{M}, \rho)$ iff:

1. \mathcal{M} and Γ are comparable; and
2. For all $C \subseteq \mathcal{A}$ and $G \subseteq \mathcal{G}$, we have:

$$\underbrace{G \in \mathcal{V}(C)}_{\text{QCG}} \quad \Leftrightarrow \quad \underbrace{\exists S \in \mathcal{E}(C, \rho) \text{ s.t. } \forall s \in S : v(s) = G}_{\text{interpretation}}$$

The first condition essentially says that the game and model contain the same agents and goals, while the second says that a game indicates that it is possible for a coalition to get some outcome iff the interpretation indicates this also.

Consider the CL predicate $\text{FEAS}(\dots)$, defined as follows:

$$\text{FEAS}(G, C) \hat{=} \langle C \rangle \chi_G$$

It is not hard to see that if \mathcal{M}, s is a CL interpretation that corresponds to some QCG Γ , then $\mathcal{M}, s \models \text{FEAS}(G, C)$ iff in Γ , $G \in \mathcal{V}(C)$. We say that $\text{FEAS}(G, C)$ corresponds to $G \in \mathcal{V}(C)$ and write $\text{FEAS}(G, C) \equiv G \in \mathcal{V}(C)$.

Some Correspondences in Qualitative Coalitional Games

First, we define a formula γ_C^+ such that γ_C^+ will be satisfied in a state s if *every* agent is satisfied in that state, i.e., if every agent has at least one of its goals satisfied in s . Similarly, γ_C^- will mean that *no* member of C is satisfied.

$$\gamma_C^+ \hat{=} \bigwedge_{i \in C} \sigma_{\mathcal{G}_i}^+ \quad \gamma_C^- \hat{=} \bigwedge_{i \in C} \pi_{\mathcal{G}_i}^-$$

Successful Coalitions. In many ways, the idea of a successful coalition incorporates the most basic question that is of interest with respect to any given QCG [WD04, p.47]. A coalition is *successful* if that coalition has a feasible choice satisfying all members of the coalition. Formally, given a QCG $\Gamma = \langle \mathcal{G}, \mathcal{A}, \mathcal{G}_1, \dots, \mathcal{G}_n, \mathcal{V} \rangle$ and a coalition $C \subseteq \mathcal{A}$, we say that C is successful iff:

$$\exists G \in \mathcal{V}(C) \text{ s.t. } \forall i \in C, \text{ we have } G \cap \mathcal{G}_i \neq \emptyset.$$

Given that a particular coalition is successful in this sense, we cannot be certain that this coalition *will* form; but we *can* be certain that an *unsuccessful* coalition will *not* form – because, by definition, the formation of such a coalition would leave at least one member unsatisfied. We can easily characterise successful coalitions, via the defined predicate $\text{SC}(C) \hat{=} \bigvee_{G \in \mathcal{V}(C)} \langle C \rangle (\chi_G \wedge \gamma_C^+)$.

PROPOSITION 2. $\text{SC}(C) \equiv \text{coalition } C \text{ is successful.}$

At first sight, the reader may suspect that the definition of $\text{SC}(\dots)$ is over engineered: would the following, simpler definition not suffice to characterise successful coalitions? $\text{SC}?(C) \hat{=} \langle C \rangle \gamma_C^+$. The answer is no. To see why, consider model \mathcal{M}, s , where $\mathcal{E}(i)(s) = \{\{s_1, s_2, s_3\}\}$ and $v(s_i) = \{g_i, g\}$. Also, in Γ we have $\mathcal{G}_i = \{g_1, g_2, g_3\}$ and $\mathcal{V}(i) = \{\{g\}\}$. Then clearly, according to the definition of $\text{SC}?(C)$, we would have that i is successful, since $\mathcal{M}, s \models \langle i \rangle \gamma_i^+$. But this does not imply that any non-empty subset of

$\{g_1, g_2, g_3\}$ represents a feasible choice for i in Γ ; in fact, we have, for all $G \in \mathcal{V}(i)$, $G \cap \mathcal{G}_i = \emptyset$. We come back to this example in Section 2.2.

Goal Realisability. The idea of goal realisability is somewhat related to that of selfish successful coalitions. We say a set of goals G is realisable if there is *any* coalition for which G is both feasible and satisfies every member [WD04, p.50]. Thus, the fact that a set of goals is realisable implies that there is at least some chance of this goal set being achieved, as it would satisfy at least one coalition. Of course, it does not imply that this goal set will be the *actual* choice of any coalition. Thus realisability is a necessary condition for the achievement of any set of goals – although it is of course not sufficient. We characterise realisability via the predicate $\text{GR}(G) \doteq \bigvee_{C \subseteq \mathcal{A}} \bigvee_{G' \subseteq \mathcal{G}} \langle C \rangle (\chi_{G'} \wedge \pi_G^+ \wedge \gamma_C^+)$. We have:

PROPOSITION 3. $\text{GR}(G) \equiv$ *goal set G is realisable.*

Minimal Coalitions. We say a coalition is *minimal* if no strict subset of this coalition is successful. The notion of minimality is important because it implies a kind of *internal stability* for a coalition (cf. [OR94, p.281]). That is, in a minimal coalition, there is no incentive for subsets of the coalition to defect away from the coalition, as, by definition, such sub-coalitions cannot be successful. Formally, a coalition C is minimal iff $\forall C' \subset C$, $\forall G \subseteq \mathcal{G}$, if $\forall i \in C'$, $G \cap \mathcal{G}_i \neq \emptyset$, then $G \notin \mathcal{V}(C')$. Minimality is easily captured in the predicate $\text{MC}(C) \doteq \bigwedge_{C' \subset C} \neg \text{SC}(C')$.

PROPOSITION 4. $\text{MC}(C) \equiv$ *coalition C is minimal.*

Core Membership and Core Non-emptiness. Perhaps the most widely studied issue in cooperative game theory is that of coalitional stability, and the tool used most widely to analyse this issue is the core [OR94, pp.257–274]. Intuitively, the core of a coalition is the set of feasible choices for that coalition from which the members of that coalition have no incentive to deviate. In the QCG setting, a parallel notion was introduced in [WD04, p.54]. Formally, we say a set of goals G is in the core of a coalition C iff: (i) C is minimal; (ii) G is feasible for C ; and in addition (iii) G satisfies every member of C . Formally, G is in the core if (i) $G \in \mathcal{V}(C)$; (ii) $\forall i \in C$ $\mathcal{G}_i \cap G \neq \emptyset$; and (iii) $\forall C' \subset C$, $\forall G' \subseteq \mathcal{G}$ if $\forall i \in C'$, $\mathcal{G}_i \cap G' \neq \emptyset$ then $G' \notin \mathcal{V}(C)$. We define the predicate $\text{CNE}(\dots)$ to capture core membership.

$$\text{CM}(G, C) \doteq \text{MC}(C) \wedge \langle C \rangle (\chi_G \wedge \gamma_C^+)$$

The correspondence result is now obvious.

PROPOSITION 5. $\text{CM}(G, C) \equiv$ *goal set G is in the core of C .*

The core of a coalition will thus be non-empty if that coalition is both minimal and successful, which easily leads to the following predicate defini-

tion.

$$\text{CNE}(C) \hat{=} \text{MC}(C) \wedge \text{SC}(C)$$

PROPOSITION 6. $\text{CNE}(C) \equiv$ *the core of C is non-empty.*

Veto Players. The notion of a veto player is generally defined in cooperative game theory with respect to *simple* coalitional games: those where every coalition simply either wins or loses. A veto player is said to be one that is a member of every winning coalition. Veto players are important because their cooperation is essential for every coalition that aspires to win: by definition, without their support, a coalition cannot win. In our framework, we can generalise the concept of a veto player to more general conditions. We say i is a veto player for φ (where φ is a formula which characterises some state of affairs) if i is a member of every coalition that can choose φ .

$$\text{VETO}(i, \varphi) \hat{=} \bigwedge_{C \subseteq \mathcal{A}} (([C]\varphi) \rightarrow \neg[C \setminus \{i\}]\varphi)$$

Note that i being a veto player for φ does *not* imply that i can bring about φ , and thus $\text{VETO}(i, \varphi) \rightarrow \langle i \rangle \varphi$ is not a valid formula scheme.

Let us now return to QCGs. In [WD04, p56–57], a notion of veto player was defined that generalised that of conventional coalitional games [OR94, p.261]. This definition related to the circumstances under which one agent is a veto player for another agent: that is, whether one agent i is a member of every coalition that is capable of satisfying j .

Formally, i is a veto player for j iff for all $C \subseteq \mathcal{A}$ and $G \in \mathcal{V}(C)$, if $G \cap \mathcal{G}_j \neq \emptyset$ then $i \in C$. It should be noted that j need not be a member of C .

$$\text{VP}(i, j) \hat{=} \bigwedge_{C \subseteq \mathcal{A}} \bigwedge_{G \subseteq \mathcal{G}} (([C](\chi_G \wedge \gamma_j^+)) \rightarrow \neg[C \setminus \{i\}]\chi_G)$$

PROPOSITION 7. $\text{VP}(i, j) \equiv$ *agent i is a veto player for agent j .*

2.2 Agent Satisfaction

On a finer level of granularity, the situation in the counterexample of page 5 demonstrates an interesting difference between QCG-games and CL-interpretations, i.e., that in the latter, is it possible to express that a coalition can achieve a goal, *without* having to specify which set of goals it *exactly* can bring about. We now consider a language that is more in line with CL. It is defined in two parts: \mathcal{L}_c is the *satisfaction language*, and is used to express properties of choices made by agents. The basic constructs in this language are of the form sat_i , meaning “agent i is satisfied”. The overall language

$\mathcal{L}(QCG)$ is used for expressing properties of QCGs themselves. The main construct in this language is of the form $\langle C \rangle \varphi$, where φ is a formula of the satisfaction language, and means that C have a choice such that this choice makes φ true. For example, $\langle 3 \rangle (sat_1 \wedge sat_4)$ will mean that 3 has a choice that simultaneously satisfies agents 1 and 4.

Formally, the grammar φ_c defines the *satisfaction language* \mathcal{L}_c , while φ_q defines the *QCG language* $\mathcal{L}(QCG)$.

$$\varphi_c ::= sat_i \mid \neg \varphi_c \mid \varphi_c \vee \varphi_c \quad \varphi_q ::= \langle C \rangle \varphi_c \mid \neg \varphi_q \mid \varphi_q \vee \varphi_q$$

where $i \in \mathcal{A}$ and $C \subseteq \mathcal{A}$.

Write $[C]\varphi$ to abbreviate $\neg \langle C \rangle \neg \varphi$. The formula $[C]\varphi$ will be defined to be true exactly when φ is a *necessary* consequence of the coalition C making a choice; φ will be true no matter *which* choice the coalition makes.

When $\Gamma = \langle \mathcal{A}, \mathcal{G}, \mathcal{G}_1, \dots, \mathcal{G}_n, \mathcal{V} \rangle$ is a QCG, $G \subseteq \mathcal{G}$ and $\varphi \in \mathcal{L}_c$, $\Gamma, G \models_Q sat_i$ is defined as follows:

$$\Gamma, G \models_Q sat_i \text{ iff } \mathcal{G}_i \cap G \neq \emptyset$$

When $\Gamma = \langle \mathcal{A}, \mathcal{G}, \mathcal{G}_1, \dots, \mathcal{G}_n, \mathcal{V} \rangle$ is a QCG then $\Gamma \models_Q \langle C \rangle \psi$ is defined as:

$$\Gamma \models_Q \langle C \rangle \psi \text{ iff there is a } G \in \mathcal{V}(C) \text{ such that } \Gamma, G \models_Q \psi$$

EXAMPLE 8. Let Γ_1 be as in Example 1. Then:

$$\begin{aligned} \Gamma_1 &\models_Q \langle C_1 \rangle (sat_1 \wedge sat_2) \\ \Gamma_1 &\models_Q (\langle C_2 \rangle sat_1 \wedge \langle C_2 \rangle sat_2) \wedge \neg (\langle C_2 \rangle (sat_1 \wedge sat_2)) \\ \Gamma_1 &\models_Q \neg (\langle C_3 \rangle sat_1 \vee \langle C_3 \rangle sat_2) \end{aligned}$$

Summarising, the satisfaction of agents is evaluated against a set of goals, while Boolean combinations of expressions referring to choices of coalitions are evaluated on a QCG Game Γ .

Expressive Power of $\mathcal{L}(QCG)$ and Axiomatisation

We look at the properties of QCGs which are definable in our language. It is clear from our language definition that what $\mathcal{L}(QCG)$ can express is which coalition can satisfy which set of agents concurrently. Note that we are not interested in *how* the coalitions make certain sets of agents satisfied, nor *why* an agent is satisfied (i.e., which goal satisfied him). We will now demonstrate that the properties of QCGs we can express in the language $\mathcal{L}(QCG)$ are exactly the properties closed under a notion of *QCG-simulation*. In other words, the language can not differentiate two games Γ and Γ' iff they *QCG-simulate* each other.

Obviously, equivalence of models transcends mere isomorphism. In particular, the semantics of performing a *choice* seem to depend only on which agents are satisfied by the choice. For example, one could imagine a mapping between “equivalent” goals of two models, maybe collapsing two goals of one model into one goal of the other. However, such a relation between models does not capture all instances of equivalent models. What is needed is a relation between *sets* of goals. This motivates the following definition of a QCG-simulation as a *relation* between two models. It is only necessary to relate goals which can actually be chosen by some coalition. Furthermore, it only makes sense to relate models which are defined over the same set of agents.

A relation

$$Z \subseteq \bigcup_{C \subseteq \mathcal{A}} (\mathcal{V}(C) \times \mathcal{V}'(C))$$

is a *QCG-simulation* between two QCGs $\Gamma = \langle \mathcal{A}, \mathcal{G}, \mathcal{G}_1, \dots, \mathcal{G}_n, \mathcal{V} \rangle$ and $\Gamma' = \langle \mathcal{A}, \mathcal{G}', \mathcal{G}'_1, \dots, \mathcal{G}'_n, \mathcal{V}' \rangle$ iff the following conditions hold for all coalitions C .

1. If GZG' then $G \cap \mathcal{G}_i = \emptyset$ iff $G' \cap \mathcal{G}'_i = \emptyset$, for all i (the *satisfaction condition*)
2. For every $G \in \mathcal{V}(C)$ there is a $G' \in \mathcal{V}'(C)$ such that GZG' (Z is *total*)
3. For every $G' \in \mathcal{V}'$ there is a $G \in \mathcal{V}(C)$ such that GZG' (Z is *surjective*)

If there exist a QCG-simulation between two games Γ and Γ' , we write $\Gamma \rightleftharpoons \Gamma'$. If $\Gamma \rightleftharpoons \Gamma'$, we can simulate any choice in one model with a choice in the other, and vice versa. This notion of simulation is somewhat similar to the notion of “alternating simulation” between alternating transition systems in [AHKV98].

EXAMPLE 9. Let Γ_2 be the QCG with the same agents as in Γ_1 (Example 1), goals f_1, f_2, \dots such that agent 1 is satisfied in f_1 and f_3 and agent 2 is satisfied in f_2, f_3 and f_4 , and the following characteristic function:

$$\begin{aligned} \mathcal{V}(C_1) &= \{ \{f_3\} \} & \mathcal{V}(C_2) &= \{ \{f_2\}, \{f_1\} \} \\ \mathcal{V}(C_3) &= \{ \{f_5\} \} & \mathcal{V}(C_4) &= \{ \{f_1\}, \{f_2\}, \{f_4\} \} \end{aligned}$$

Then $\Gamma_1 \rightleftharpoons \Gamma_2$.

THEOREM 10. *Satisfaction is invariant under QCG-simulation:*

$$\Gamma \rightleftharpoons \Gamma' \quad \Rightarrow \quad \forall \varphi \in \mathcal{L}(\text{QCG}) [\Gamma \models_Q \varphi \Leftrightarrow \Gamma' \models_Q \varphi]$$

Elsewhere, we showed that all that one needs to axiomatise validity in QCG’s is the modal logic \mathbf{K} , we omit the details here.

Solution Concepts. It should be clear that many of the solution concepts of Section 2.1 can be characterised via formulae of $\mathcal{L}(QCG)$. For example, we can characterise successful coalitions as $succ(C) \equiv \langle C \rangle (\bigwedge_{i \in C} sat_i)$. Similarly, the notion of a minimal coalition may be captured as by $min(C) \equiv \bigwedge_{C' \subsetneq C} \neg succ(C')$. Thus, the core of a coalition being non-empty may be captured as $cne(C) \equiv (succ(C) \wedge min(C))$.

Apart from an agent vetoing an outcome, he can also veto an other player [WD04, p.57]: $veto(i, j) \equiv \bigwedge_{C \subseteq \mathcal{A}} (\langle C \rangle sat_j \rightarrow \neg \langle C \setminus \{i\} \rangle sat_j)$. Finally, the idea of a coalition being mutually dependent [WD04, p.58] is then is: $md(C) \equiv \bigwedge_{i \neq j \in C} veto(i, j)$.

2.3 Temporal QCGs

In principle there are many ways to temporalise QCGs. As a first investigation, we assume a linear time model, in which, at each time point, a (possibly different) QCG Γ is played. A *temporal qualitative coalitional game* (TQCG) is then a triple

$$M = \langle S, \sigma, Q \rangle \text{ where:}$$

- S is a set of *states*;
- $\sigma : \mathbb{N} \rightarrow S$ associates a state $\sigma(u)$ with every point $u \in \mathbb{N}$; and
- $Q : S \rightarrow \mathbf{Q}$, where \mathbf{Q} is the class of all QCGs, associates a qualitative coalitional game $Q(s) = \langle \mathcal{A}^s, \mathcal{G}^s, \mathcal{G}_1^s, \dots, \mathcal{G}_n^s, \mathcal{V}^s \rangle$ with every state s .

We will make just one requirement of TQCGs: that the set of agents and overall goals remains the same in all states. Formally, $\forall s, t \in S: \mathcal{A}^s = \mathcal{A}^t = \mathcal{A}$ and $\mathcal{G}^s = \mathcal{G}^t = \mathcal{G}$. This does not mean that an agent's goals must remain fixed, however: we allow for the possibility that an agent has different goals in different states.

A Logic for TQCGs

To express properties of TQCGs, we extend the QCG language $\mathcal{L}(QCG)$ with the standard temporal operators of linear-time temporal logic: \bigcirc – “next”, \diamond – “eventually”, \square – “always in the future”, and \mathcal{U} – “until”. Formally, the language $\mathcal{L}(TQCG)$ is defined by the grammar φ_t .

$$\varphi_t ::= \langle C \rangle \varphi_c \mid \neg \varphi_t \mid \varphi_t \vee \varphi_t \mid \varphi_t \mathcal{U} \varphi_t \mid \bigcirc \varphi_t$$

We again assume the usual derived propositional connectives, in addition to $\diamond \varphi$ for $\top \mathcal{U} \varphi$ and $\square \varphi$ for $\neg \diamond \neg \varphi$. Moreover, we define $\square^* \varphi$ as $(\varphi \wedge \square \varphi)$ (φ is true now and always in the future), and $\diamond^* \varphi = \neg \square^* \neg \varphi$ (φ is true now or sometime in the future).

When $M = (S, \sigma, Q)$ is a TQCG, $u \in \mathbb{N}$, and φ is a $\mathcal{L}(TQCG)$ formula, the satisfaction relation $M, u \models_T \varphi$ is defined as follows (the cases for negation and disjunction are defined as usual):

$$\begin{aligned} M, u \models_T \varphi \text{ iff } Q(\sigma(u)) \models_Q \varphi, \text{ when } \varphi \in \mathcal{L}(QCG) \\ M, u \models_T \bigcirc \psi \text{ iff } M, u + 1 \models_T \psi \\ M, u \models_T \psi_1 \mathcal{U} \psi_2 \text{ iff there is some } i \text{ such that } M, u + i \models_T \psi_2 \text{ and for} \\ \text{all } 0 < j < i, M, u + j \models_T \psi_1 \end{aligned}$$

For instance, the following formula of $\mathcal{L}(TQCG)$ means that eventually, agent 3 can always choose to satisfy agents 1 and 4 simultaneously: $\diamond \square \langle 3 \rangle (sat_1 \wedge sat_4)$.

We will henceforth use $\mathcal{L}(TQCG)$ to refer to both the language, and the logic we have defined over this language.

Properties of TQCGs

The notion of simulation for QCGs (Section 2.2) can be naturally lifted to the temporal case. When $M = (S, \sigma, Q)$ and $M' = (S', \sigma', Q')$ are TQCGs and $k \geq 0$, we define

$$\begin{aligned} M, k \rightleftharpoons_T M', k &\Leftrightarrow Q(\sigma(k)) \rightleftharpoons Q'(\sigma'(k)) \\ M \rightleftharpoons_T M' &\Leftrightarrow \forall n \geq 0, M, n \rightleftharpoons_T M', n \end{aligned}$$

The notion of elementary equivalence for TQCGs over the language $\mathcal{L}(TQCG)$ can be defined as follows. $M, k \equiv M', k$ iff, for every $\varphi \in \mathcal{L}(TQCG)$, $M, k \models_T \varphi$ iff $M', k \models_T \varphi$. $M \equiv M'$ iff $M, k \equiv M', k$ for every $k \geq 0$.

THEOREM 11. *For all TQCGs M, M' : $M \rightleftharpoons_T M' \Leftrightarrow M \equiv M'$*

The *satisfiability* problem for $\mathcal{L}(TQCG)$ is as follows: given a formula $\varphi \in \mathcal{L}(TQCG)$, does there exist a TQCG M and $u \in \mathbb{N}$ such that $M, u \models \varphi$?

THEOREM 12. *The sat. probl. for $\mathcal{L}(TQCG)$ is PSPACE-complete.*

Characterizing TQCGs

In this section, we investigate the axiomatic characterisation of various classes of TQCG. As usual, in saying that a formula scheme φ characterises a property P of models, we mean that φ is valid in a model M iff M has property P ; if only the right-to-left part of this biconditional holds, then we say property P implies φ . Also note that for an $\mathcal{L}(TQCG)$ formula φ , to say that φ is valid in a class of models, is the same as saying that $\square^* \varphi$ is valid in that class.

Basic Correspondences.

Let $h^s(C)$ denote the set of all agents that could possibly be satisfied (not necessarily jointly) by coalition C in state s :

$$h^s(C) = \{i : i \in \mathcal{A} \ \& \ \exists G \in \mathcal{V}^s(C), \mathcal{G}_i^s \cap G \neq \emptyset\}$$

The “ h ” here is for “happiness”: we think of $h^s(C)$ as all the agents that C could possibly make happy in s . Thus the semantic property $i \in h^s(C)$ is a counterpart to the syntactic expression $\langle C \rangle sat_i$.

Consider the following two constraints. The first, EH , says that eventually, C will be able to make i happy.

$$\exists u \in \mathbb{N}, (i \in h^{\sigma(u)}(C)) \quad (EH)$$

Notice that in the terminology of reactive systems, this is a *fairness* or *response* property: it implies that something (i being made happy) can happen infinitely often. (Of course, the fact that C *can* make i happy infinitely often does not mean they will do so.) Note that $\diamond^* \langle C \rangle sat_i$ characterises EH .

Now consider a *safety* property. The constraint AH says that C can always make i happy, while the constraint AU says that C can never make i happy.

$$\forall s \in S, (i \in h^s(C)) \quad (AH) \quad \forall s \in S, (i \notin h^s(C)) \quad (AU)$$

The reader will see that $\langle C \rangle sat_i$ characterises AH , and $\neg \langle C \rangle sat_i$ char. AU .

There are several properties we can investigate with respect to goal sets. First, suppose that agent i 's goal set is guaranteed to *monotonically decrease* over time. Suppose we impose this condition *strict*, so that an agent i is guaranteed to get strictly harder to satisfy over time. This condition is defined by the following further constraint, in addition to $MDGS$.

$$\forall u \in \mathbb{N} \quad \begin{array}{l} (\mathcal{G}_i^{\sigma(u)} = \emptyset) \vee \\ (\exists v \in \mathbb{N} : (v > u) \wedge (\mathcal{G}_i^{\sigma(v)} \subset \mathcal{G}_i^{\sigma(u)})) \end{array} \quad (SMDGS)$$

We get the following. $SMDGS$ is characterised by $\varphi = \neg sat_i \vee \diamond \square \neg sat_i$. Note that our language is too weak to distinguish $SMDGS$ from the following property, which is also characterised by φ : $\forall u \in \mathbb{N} (\mathcal{G}_i^{\sigma(u)} = \emptyset) \vee (\exists v \in \mathbb{N} : (v > u) \wedge (\mathcal{G}_i^{\sigma(v)} = \emptyset))$.

Solution Concepts. How might our solution concepts be extended into the temporal dimension of TQCGs and $\mathcal{L}(TQCG)$? It should first be clear that each concept has four different temporal versions, corresponding to prefixing the formula characterising it with one of the following four, increasingly powerful temporal operators:

$$\diamond \quad \square \diamond \quad \diamond \square \quad \square$$

Thus, for example, $\Box \Diamond succ(C)$ means that coalition C are successful *infinitely often* – no matter which time point we pick, there will be a subsequent time point at which C are successful. (Using the terminology of reactive systems, we might then say that C are hence *fairly successful*.) Similarly, a *temporally strong* form of coalitional stability is captured by the formula $\Box cne(\mathcal{A})$: if this formula is satisfied in a TQCG, then, it can be argued, the only coalition that will ever form is the grand coalition.

It is potentially more interesting, however, to study a richer interplay between temporal and QCG dimensions. For example, from agent i 's point of view, perhaps the only really interesting issue is whether at every time point there is some stable coalition, containing this agent.

$$tstable(i) \equiv \Box \bigvee_{C \subseteq \mathcal{A}: i \in C} cne(C)$$

From the point of view of a coalition C , which seeks to form, the notion of a *stable government* seems relevant: a stable government is a coalition that can always satisfy its “electorate”.

$$sg(C) \equiv \Box \langle C \rangle \left(\bigwedge_{i \in \mathcal{A}} sat_i \right)$$

This can of course be strengthened, requiring C to in addition be an internally stable coalition.

$$sg'(C) \equiv \Box (cne(C) \wedge \langle C \rangle \left(\bigwedge_{i \in \mathcal{A}} sat_i \right))$$

With respect to mutual dependence, one possibility, captured by the formula $\Box md(C)$, is that a coalition is *always* mutually dependent. However, we can capture a weaker type of mutual dependence as follows:

$$wmd(C) \equiv \bigwedge_{i \neq j \in C} \Diamond veto(i, j)$$

We draw two conclusions. The first is that the language $\mathcal{L}(TQCG)$ is well suited to capturing such solution concepts. The second is that extending QCGs into the temporal dimension adds an entirely new level of richness to their structure, which, as these examples suggest, demands further study.

3 Coalitional Games

A *coalitional game* (without transferable payoff) is an $(m+3)$ -tuple [OR94, p.268]: $\Gamma = \langle \mathcal{A}, \Omega, \preceq_1, \dots, \preceq_m, V \rangle$ where $, \preceq_i \subseteq \Omega \times \Omega$ is a complete, reflexive, and transitive *preference relation*, for each agent $i \in \mathcal{A}$. Its language is

defined in two parts. First, given a set of outcome symbols Ω (we will blur the difference between the semantic objects and the symbols that denote them), we have an *outcome language* \mathcal{L}_o , defined by the grammar φ_o , below, which expresses the properties of outcomes. The outcome symbols themselves are the main constructs of this language; a formula such as $\omega_1 \vee \omega_2$ means that the outcome corresponds to either ω_1 or ω_2 . Next, given a set of agent symbols \mathcal{A} and a set of coalition symbols Σ_C , we have a *cooperation language* \mathcal{L}_c , for expressing the properties of coalitional cooperation, and the preferences that agents have over possible outcomes. This language is generated by the grammar φ_c below. \mathcal{L}_c has two main constructs. First, $\omega_1 \succeq_i \omega_2$ expresses the fact that agent i either prefers outcome ω_1 over outcome ω_2 , or is indifferent between the two. Second, $\langle C \rangle \varphi$ (where $C \in \Sigma_C$) says that C can choose an outcome in which the formula φ will be true. This construct may seem syntactically similar to counterparts Coalition Logic, but it stands here for a fundamentally different concept due to the semantic differences mentioned above.

$$\begin{aligned} \varphi_o & ::= \omega \mid \neg \varphi_o \mid \varphi_o \vee \varphi_o \\ \varphi_c & ::= (\omega \succeq_i \omega') \mid \langle C \rangle \varphi_o \mid \neg \varphi_c \mid \varphi_c \vee \varphi_c \end{aligned}$$

where i is an agent symbol, C is a coalition symbol, and ω, ω' are outcome symbols.

An \mathcal{L}_c formula γ is interpreted in a coalitional game Γ as follows. First, we define the satisfaction of a \mathcal{L}_o formula α in an outcome ω of a coalitional game Γ , written $\Gamma, \omega \models \alpha$:

$$\Gamma, \omega \models \omega' \text{ iff } \omega = \omega'$$

Satisfaction of γ in Γ is then defined as follows:

$$\Gamma \models (\omega_1 \succeq_i \omega_2) \text{ iff } (\omega_1 \supseteq_i \omega_2)$$

$$\Gamma \models \langle C \rangle \varphi \text{ iff } \exists \omega \in V(C) \text{ such that } \Gamma, \omega \models \varphi$$

Note that $\langle C \rangle \top$ iff C can at least bring about something: $V(C) \neq \emptyset$. $[C] \varphi$ means $\neg \langle C \rangle \neg \varphi$, i.e., every choice of C must involve φ . As an example, suppose $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ and $V(C) = \{\omega_1, \omega_2\}$. Then:

$$\langle C \rangle \omega_1 \wedge \langle C \rangle (\omega_1 \vee \omega_3) \wedge \neg \langle C \rangle \omega_3 \wedge [C] (\omega_1 \vee \omega_2) \wedge \neg [C] (\omega_1 \vee \omega_3)$$

Note that if $\omega_1 \neq \omega_2$, then we can have $\langle C \rangle \omega_1 \wedge \langle C \rangle \omega_2$, but the formula $\langle C \rangle (\omega_1 \wedge \omega_2)$ can never be true.

Let us, for any coalition C and set of outcome symbols Δ , suggestively write $\langle [C] \rangle \Delta$ for $\bigwedge_{\delta \in \Delta} \langle C \rangle \delta \wedge [C] \bigvee_{\delta \in \Delta} \delta$. A formula of this form is said to

fully describe C 's choices. It is easy to see that we have the following. Let $\Delta \subseteq \Omega$.

$$\Gamma \models \langle [C] \rangle \Delta \quad \text{iff} \quad V(C) = \Delta$$

Exclusive disjunctions $\varphi \nabla \psi$ play an important role in our proofs. Moreover, if Φ is a set of formulas, then we define $\nabla_{\varphi \in \Phi} \varphi$ to be true iff exactly one of the φ 's is true. Note that $\Gamma \models [C](\omega_i \vee \omega_j) \leftrightarrow [C](\omega_i \nabla \omega_j)$ when $i \neq j$: using the definition of $[C]$ and contraposition this is the same as $\Gamma \models \langle C \rangle \neg(\omega_i \nabla \omega_j) \leftrightarrow \langle C \rangle \neg(\omega_i \vee \omega_j)$. Now, syntactically, $\langle C \rangle \neg(\omega_i \nabla \omega_j)$ is equivalent to $\langle C \rangle ((\omega_i \wedge \omega_j) \vee \neg(\omega_i \vee \omega_j))$. But, inspecting the truth-definition of $\langle C \rangle$, this is again equivalent to $\langle C \rangle \neg(\omega_i \vee \omega_j)$ since the \mathcal{L}_o formula $\omega_i \wedge \omega_j$ is never true.

So, which properties can be expressed with our cooperation language of coalition game logic (CGL)? The answer, given by the following theorem, is “all”, when we restrict the possible outcomes of a game to a finite set.

THEOREM 13. *The logic CGL is expressively complete with respect to finite coalitional games. That is, for any two finite coalitional games Γ_1, Γ_2 such that $\Gamma_1 \neq \Gamma_2$, there exists a CGL formula ζ such that $\Gamma_1 \models \zeta$ and $\Gamma_2 \not\models \zeta$.*

3.1 Properties of CGL

Elsewhere we present an axiomatic system for the language \mathcal{L}_c , and proved its soundness and completeness with respect to the class of all finite coalitional games without transferable payoff. Of course, this contains axioms guaranteeing that \sqsubseteq_i is a complete, reflexive and transitive order. On top of that, there are the modal principles for $\Box C$ and the property $[C](\nabla_{\omega \in \Omega} \omega)$ which says that whatever a coalition chooses, must be a unique alternative from Ω .

It is trivial to see that the *model checking problem* for CGL (i.e., the problem of determining, for any given game Γ and φ , whether or not $\Gamma \models \varphi$) may be solved in deterministic polynomial time: an obvious recursive algorithm for this problem can be directly extracted from the semantic rules of the language. The *satisfiability problem* is the problem of checking whether or not, for any given φ there exists a game Γ such that $\Gamma \models \varphi$. For most modal logics, the corresponding satisfiability problem has a trivial NP-hard lower bound, since such logics subsume propositional logic, for which satisfiability is the defining NP-complete problem [BDRV01, p.374]. However, our logic is specialised for reasoning about coalitional games, and it is not so obvious that it subsumes propositional logic, since we do not have primitive propositions. NP-hardness must therefore be proven from first principles. We only give the result:

THEOREM 14. *The satisfiability problem for CGL formulae is NP-complete, even for CGL formulae φ such that $|ag(\varphi)| = 1$.*

3.2 Solution Concepts

Elsewhere, we characterised three solution concepts from the theory of coalitional games, viz. *the core*, *stable sets* and *the bargaining set* in CGL. We used the formulations of these solution concepts in [OR94]; there the two latter solution concepts are however defined only for games with real numbered payoffs and transferable utility and below we translate the definitions to the more general games with preference relations over general outcomes and non-transferable utility. We here demonstate the first two concepts. Henceforth, a *C-feasible outcome* is an outcome which can be chosen by the coalition C and a *feasible outcome* is an \mathcal{A} -feasible outcome. We start by looking at the core, which is a, possibly empty, set of outcomes.

DEFINITION 15 (Core). The core of a coalitional game is the feasible outcomes ω for which there is no coalition C with a C -feasible outcome ω' such that $\omega' \succ_i \omega$ for all $i \in C$.

We write $CM(\omega)$ to mean that ω is in the core.

$$CM(\omega) \equiv \langle \mathcal{A} \rangle \omega \wedge \neg \left[\bigvee_{C \subseteq \mathcal{A}} \bigvee_{\omega' \in \Omega} (\langle C \rangle \omega') \wedge \bigwedge_{i \in C} (\omega' \succ_i \omega) \right]$$

CNE will then mean that the core is non-empty: $CNE \equiv \bigvee_{\omega \in \Omega} CM(\omega)$

THEOREM 16. *The core of a finite coalitional game Γ is non-empty iff $\Gamma \models CNE$.*

A *stable set* is a set of outcomes. A coalitional game may have several stable sets, but must not necessarily have any. We characterize stable sets in terms of *imputations* and *objections*. An imputation is a feasible outcome that for each agent i is as least as good as any outcome the singleton coalition $\{i\}$ can choose on his own. The CGL formula $IMP(\omega)$ is true whenever ω is an imputation:

$$IMP(\omega) \equiv \langle \mathcal{A} \rangle \omega \wedge \bigwedge_{\omega' \in \Omega} \bigwedge_{i \in \mathcal{A}} (\langle \{i\} \rangle \omega' \rightarrow \omega \succeq_i \omega')$$

An imputation ω is a C -objection to an imputation ω' if every agent in C prefers ω over ω' and the coalition C can choose an outcome which for every agent in C is as least as good as ω . ω is an objection to ω' if ω is a C -objection to ω' for some coalition C . Next, $OBJ(\omega, \omega', C)$ expresses that outcome ω is an C -objection to outcome ω' , when both ω and ω' are imputations:

$$OBJ(\omega, \omega', C) \equiv \left(\bigwedge_{i \in C} \omega \succ_i \omega' \right) \wedge \bigvee_{\omega'' \in \Omega} (\langle C \rangle \omega'' \wedge \bigwedge_{i \in C} \omega'' \succeq_i \omega)$$

DEFINITION 17 (Stable Set). A set of imputations Y is a stable set if it satisfies: (Internal stability) If $\omega \in Y$, there is no objection to ω in Y , and (External stability) If $\omega \notin Y$, there is an objection to ω in Y .

Now consider

$$\begin{aligned} STABLE(Y) \equiv & \bigwedge_{\omega \in Y} IMP(\omega) \\ & \wedge \left(\bigwedge_{\omega \in Y} \bigwedge_{C \subseteq \mathcal{A}} \bigwedge_{\omega' \in Y} \neg OBJ(\omega', \omega, C) \right) \\ & \wedge \left(\bigwedge_{\omega \in \Omega \setminus Y} IMP(\omega) \rightarrow \left(\bigvee_{C \subseteq \mathcal{A}} \bigvee_{\omega' \in Y} OBJ(\omega', \omega, C) \right) \right) \end{aligned}$$

THEOREM 18. Y is a stable set of a finite coalitional game Γ iff $\Gamma \models STABLE(Y)$.

3.3 Relation to Coalition Logic

As we noted in section 3, it is rather tempting to believe that the outcomes of coalitional games can be interpreted as states, and that the characteristic function can be interpreted as an effectivity function, and that as a consequence CL could be interpreted directly in coalitional games. We now argue that in fact there is a fundamental difference between the two approaches. We say that a coalitional game Γ and a pointed coalition model M, t are *outcome-equivalent* if $\mathcal{S} = \Omega \cup \{t\}$, and Γ and (M, t) agree on \mathcal{L}_c formulae and (Γ, ω) and (M, ω) agree on \mathcal{L}_o formulae for any outcome ω . Consider the class of *limited games* where $V(C) = \{\omega\}$ for all coalitions $C \neq \mathcal{A}$, for some fixed outcome $\omega \in \Omega$.

THEOREM 19. *No non-limited coalitional game with more than one player has an outcome-equivalent coalition model.*

Thus, in general, a coalitional game is *not* simply a coalition model with outcomes as states. Even though the language of Coalition Logic is similar to the language of our logic, it follows from Theorem 19 that we cannot use the semantic rules of Coalition Logic “directly” to say whether a formula is true or not in a coalitional game. The main reason is that a difference between outcomes in coalitional games and states in coalition models is that an outcome is *local* to the coalition which chooses it, while states are *global*. As a consequence, while it is perfectly possible in a coalitional game that both a coalition C can choose outcome ω ($\omega \in V(C)$) and a coalition C' , C' and C disjoint, can choose outcome ω' ($\omega' \in V(C')$) when $\omega' \neq \omega$, it is not possible in a coalition model that both C is effective for $\{\omega\}$ and C' is effective for $\{\omega'\}$.

4 Related Work

Recently, game theory has come to be seen as an attractive foundation upon which to develop logic-based knowledge representation formalisms for multi-agent systems. It has been recognised for several decades that there are close links between modal logics of rational agency and the formal theory of games: see, for example, Ladner and Reif’s Church/Turing-like thesis for distributed computing, and the conclusions they draw from this [LR86, pp.208–209]. Recently, a number of formalisms have been proposed which attempt to synthesise logical and game-theoretic approaches in a single system, in which the links between the game and the logic are explicitly defined (for a survey, see [vdHP06]). Explorations have been undertaken by van Benthem, whose starting point is that the labelled transition systems/Kripke structures, which are canonically used to give a semantics to modal logics, can be interpreted as extensive form games, and that as a consequence modal operators of various kinds can be used to express properties of games [Ben02]. However, to the best of our knowledge, our work is the first to present a systematic logical characterisation of concepts from cooperative games.

Acknowledgements This research was supported by the EPSRC under grants GR/R60836/01 and GR/S62727/01.

BIBLIOGRAPHY

- [AHKV98] R. Alur, T. Henzinger, O. Kupferman, and M. Vardi. Alternating refinement relations. In *International Conference on Concurrency Theory*, pages 163–178, 1998.
- [ÅvdHW06a] T. Ågotnes, W. van der Hoek, and M. Wooldridge. On the logic of coalitional games. In *Proc. AAMAS-2006*, Hakodate, Japan, 2006.
- [ÅvdHW06b] T. Ågotnes, W. van der Hoek, and M. Wooldridge. Temporal qualitative coalitional games. In *Proc. AAMAS-2006*, Hakodate, Japan, 2006.
- [BdRV01] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press: Cambridge, England, 2001.
- [Ben02] J. van Benthem. Extensive games as process models. *Journal of Logic, Language, and Information*, 11(3):289–313, 2002.
- [DvdHW07] P. Dunne, W. van der Hoek, and M. Wooldridge. A logical characterisation of qualitative coalitional games. Accepted *Journal of Applied Non-Classical Logics*, 2007.
- [LR86] R. E. Ladner and J. H. Reif. The logic of distributed protocols: preliminary report. In *TARK 1986*, pages 207–222. Morgan Kaufmann Publishers: San Mateo, CA, 1986.
- [OR94] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press: Cambridge, MA, 1994.
- [Pau02] M. Pauly. A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1):149–166, 2002.
- [vdHP06] W. van der Hoek and M. Pauly. Modal logic for games and information. In P. Blackburn, J. van Benthem, and F. Wolter, editors, *Handbook of Modal Logic*, pages 1077–1148. Elsevier, Amsterdam, 2006.
- [WD04] M. Wooldridge and P. E. Dunne. On the computational complexity of qualitative coalitional games. *Artificial Intelligence*, 158(1):27–73, 2004.