Preferences in Qualitative Coalitional Games

Paul E. Dunne and Michael Wooldridge
Dept of Computer Science, University of Liverpool
Liverpool L69 7ZF, United Kingdom
{ped,mjw}@csc.liv.ac.uk

Abstract

*Qualitative coalitional games (*QCGs*) are a variation of conventional coalitional games in which each coalition may choose to cooperate in a number of different ways, with different choices resulting in potentially different sets of goals being achieved; each agent is associated with a set of goals, the intuition being that an agent is "satisfied" if any of its goals are achieved, but is indifferent between them. In this paper, we extend the framework of* QCGs *to incorporate* preferences *that agents have over their goals. In addition to establishing some basic properties of* QCGs *with Preferences (*QCGPs*), we investigate and characterise the complexity of six natural decision problems associated with* QCGPs*. For example, we prove that the problem of establishing Pareto optimality of a goal set with respect to some coalition is co-*NP*-complete. We end with some brief conclusions and a discussion of related work.*

1. Introduction

Coalitional games are abstract models of cooperative interaction between self-interested agents, and as such, they have been widely studied in the game theory literature [5, pp.255–312]. The success of such models has led to them being adapted for use within the multiagent systems community, where they have proved to be of great value in understanding the nature of coalitions and coalition formation [8, 9, 7]. Conventional coalitional games (with transferable payoff) assign to each potential coalition a numeric value, corresponding to utility that could be distributed among coalition members if the coalition chose to cooperate. Given such models, solution concepts such as the core attempt to predict which coalitions might form, by considering the potential payoff an agent could get from joining different coalitions: a coalition is viewed as being "stable" if the members of that coalition have no incentive to defect and join any other coalition [5, pp.258–259].

In [10], we introduced a variation of coalitional games called *Qualitative Coalitional Games* (QCGs). In a QCG, each agent is assumed to have a set of goals: an agent is "satisfied" with any outcome that accomplishes one of its goals, but is indifferent about *which* of its goals should be achieved – all are considered equally good (individual rational). Each potential coalition is then modelled as having a set of choices available, intutitively corresponding to the different ways in which they could choose to cooperate. Associated with each choice is a set of goals, which would be achieved if the coalition chose to cooperate in this way. QCGs seem an appropriate abstract framework within which to reason about goal-oriented multiagent systems, where numeric utility values are either inappropriate or else impossible to derive. In [10], we formulated and investigated the computational complexity of a wide range of solution concepts associated with QCGs: for example, we demonstrated that showing the core of a QCG is non-empty (i.e., that a coalition has a choice from which there is no incentive for any member of the coalition to deviate) is D^p-complete.

The aim of the present paper is to extend the basic framework of QCGs, by allowing for agents that have preferences over goals. We begin by formulating *Qualitative Coalitional Games with Preferences* (QCGPs), defining some concepts (such as the core and stable sets) associated with QCGPs, and establishing some fundamental properties of these. The remainder of the paper is then largely taken up with an investigation into the computational complexity of six decision problems associated with QCGPs. For example, we show that the problem of establishing Pareto optimality of a goal set with respect to some coalition is co-NP-complete. In general we will simply state results without presenting detailed proofs. The reader interested in such is referred to a more extensive version of this paper [3]. We end with a brief discussion of related work and some conclusions.

Notation: We use $\top$ and $\bot$ to denote the Boolean constants "true" and "false". For a propositional formula $\Phi(x_1, \ldots, x_n)$ over the variables $X_n = \langle x_1, \ldots, x_n \rangle$, given $Z \subseteq X_n$, we denote by $\Phi[Z]$ the result of evaluating Φ under the instantiation $x_i := \top$ if $x_i \in Z$, $x_i := \bot$ if $x_i \notin Z$. In addition to the standard Boolean operations of conjunction ($\wedge$), disjunction ($\vee$), implication ($\Rightarrow$), and negation ($\neg$), we use the binary exclusive-or operation, denoted by $\oplus$. We usually omit the conjunction symbol ($\wedge$) in formulae, for example writing $\varphi\psi$ to abbreviate $\varphi \wedge \psi$. Some familiarity with computational complexity theory is assumed [6].

2. QCGPs

As noted above, QCGs were introduced as an abstract model of scenarios in which an agent, having some set of goals of which it would like to realise at least one, may have to cooperate with other agents in order to bring this about [10]. The characteristic function of conventional coalitional games is replaced in QCGs by a function which allocates to every coalition a set of choices, where each choice intuitively corresponds to one way that the coalition could choose to cooperate. Each choice is then associated with a set of goals, the intuition being that if the coalition chose to cooperate in this way, then the associated set of goals would be achieved. Here, we extend the QCG model by assuming that agents have *preferences* over goals.

Definition 1 *A* Qualitative Coalitional Game with Preferences *(*QCGP*) is a $2n+3$-tuple:*

$$\Gamma = \langle G, Ag, G_1, \ldots, G_n, \Psi, \triangleright_1, \ldots, \triangleright_n \rangle, \quad \textit{where:}$$

- $G = \{g_1, \ldots, g_m\}$ *is a (finite, non-empty) set of* goals*;*
- $Ag = \{1, \ldots, n\}$ *is a (finite, non-empty) set of* agents*;*
- $G_i \subseteq G$ *is the set of* acceptable goals *for agent $i \in Ag$;*
- Ψ *is a propositional logic formula defined over the variables $Ag = \langle a_1, \ldots, a_n \rangle$ and $G = \langle g_1, \ldots, g_m \rangle$, representing the* characteristic function *of the game: for every coalition $C \subseteq Ag$ and goal set $G' \subseteq G$, we have $\Psi[C, G'] = \top$ if one of the choices for coalition C is goal set G';*
- $\triangleright_i \subseteq G_i \times G_i$ *is a partial order over G_i representing i's* preference relation*, so that $g_1 \triangleright_i g_2$ indicates that i would rather have goal g_1 satisfied than goal g_2.*

We assume that $\triangleright_i$ is presented as a directed acyclic graph D_i, so that forming the transitive closure of this graph's adjacency matrix will yield a matrix D_i^* for which $D_i^*[g_1, g_2] = \top$ if and only if $g_1 \triangleright_i g_2$.

A set of goals G' *satisfies* agent i if $G' \cap G_i \neq \emptyset$. We say G' is *feasible* for coalition C if this goal set corresponds to one of the choices of C, i.e., if $\Psi[C, G'] = \top$. For $C \subseteq Ag$, let $\mathcal{X}(C)$ denote the set of subsets of G that are both feasible for C and satisfy every member of C:

$$\mathcal{X}(C) \; = \; \{G' \subseteq G \mid \Psi[C, G'] = \top \wedge \bigwedge_{i \in C} G_i \cap G' \neq \emptyset\}$$

If $\mathcal{X}(C) \neq \emptyset$, then we say that C is successful (since it has at least one feasible choice that would result in an individual rational outcome for each of its members). Let $\mathcal{X}^\Gamma$ denote the set of all goals sets which both satisfy and are feasible for some coalition in Γ: $\mathcal{X}^\Gamma = \cup_{C \subseteq Ag} \mathcal{X}(C)$.

We first comment on some aspects of these definitions, in particular the relationship between the agents, goal sets and feasibility function represented by the propositional formula Ψ. It should be noted that, in principle, f_Ψ the *function* described by $\Psi(Ag, G)$ could be *any* propositional function, e.g. we do not assume *a priori* that monotonicity conditions hold such as $f_\Psi[C, G'] = \top$ implying $f_\Psi[C', G'] = \top$ whenever $C \subseteq C'$. To allow such generality may well seem to be rather excessive, however, there are a number of reasons why we proceed in this way. Firstly, it is certainly the case that such latitude does not preclude any subsequent consideration of restrictions such as monotonicity, i.e. our formalism does not *lose* any expressive power by being general. A second reason concerns rather subtle technical issues that would need to be addressed were only feasibility functions meeting given criteria to be employed. A full elaboration of this point would be out of place in the context of the present paper and so we merely comment that these arise in describing instances of QCGPs in decision problems: if we wish to consider *only* those $f_\Psi(Ag, G)$ that satisfy some property then it is reasonable to insist that $\Psi(Ag, G)$ – the representing formula – can be easily *validated* as defining such functions. For a more detailed (and rather technical) discussion of such representation issues with respect to monotone feasibility functions, we refer the interested reader to [10, pp. 13–17][1]. Finally, as regards the specific case of monotone feasibility functions, it is interesting to note there are strong indications suggesting that the decision problems we consider for QCGPs would not become any easier if such a restriction were imposed: of the nine QCG related decision problems in [10] for which monotone variants are examined, in all but two cases the complexities of the general and monotone forms are identical, i.e. problems NP-complete, D^p-complete, etc. in the general case remain so in the monotone variant. In only one instance does monotonicity result in a polynomial time decision process for a problem whose unrestricted form is intractable.

Of course, in practice "almost all" of the $2^{2^{n+m}}$ possible definitions of $f(Ag, G)$ will be of no interest either by reason of their being unlikely to occur in a realistic context or because the shortest formula Ψ equivalent to $f(Ag, G)$ has excessive length. We also note that the manner in which Ψ arises may vary considerably depending on the exact scenario being modelled within an associated QCGP. Thus, sometimes it may happen that the relationship between agents and their associated goals suffices in itself to define Ψ; in others external factors may influence how Ψ (or more accurately f_Ψ the associated propositional function) is formed. In order better to appreciate this distinction, consider the following two examples.

1 The page numbers are for the technical report version of [10], not the journal version which is in press at the the time of writing.

Example 1 *A set of students have to select project topics from a given collection. Each project can be allocated to at most one student and students have preferences over the subset of projects they would be prepared to undertake.*

Example 2 *We have a set of students and collection of project topics as in Example 1, and the same model of preferences. The allocation control is, however, complicated by two additional factors: each project may be undertaken by* several *students, e.g. as a team working exercise; there are external considerations governing whether an individual student (or team of students) is seen as suited to undertake particular projects, e.g. some may require background in certain specialist fields, some may be considered "too easy" or "too demanding" for specific classes of student.*

In both examples we have $Ag = \{s_1, \ldots, s_n\}$ corresponding to the student set and $G = \{p_1, \ldots, p_m\}$ the pool of available projects so that G_i and the associated partial order $\triangleright_i$ defines the subset which s_i is prepared to undertake and preferences over these. Now in Example 1 there are a number of ways in which one could define $\Psi(Ag, G)$: for example, by $\Psi[C, G'] = \top$ if and only if there is an injective mapping $\beta : C \to G'$ with which $\beta(a_i) \in G_i$ for each $a_i \in C$, i.e. G' is feasible if it allows each s_i to be allocated a project in their target set. It is clearly the case that such Ψ depend solely on the specified sets G_i. In contrast, with Example 2, knowing G_i may not suffice to allow Ψ to be given: the presence of the additional factors colours which G' may be feasible for a given C. It is worth noting that Example 2 gives rise to a number of non-trivial strategic issues: suppose a final allocation of projects to students is to be generated by requiring individuals, s_i, to specify a (linear) preference ordering over some small number of projects, P_i. Now it cannot be *known* if $P_i \subseteq G_i$ (although, presumably, $P_i \cap G_i \neq \emptyset$), and thus if individuals know how the allocation algorithm operates, coalitions may be able to form that ensure those within it are given their most preferred choice by including "false preferences" in their returns P_i. One aim underlying our formulation of QCGPs is to provide methods by which such possibilities can be analysed.

Given a QCGP Γ, we can extend the concept of an agent's preferences to that of a *coalition*'s. We present two definitions and discuss how these are related. It is, of course, the case that a number of methods for defining preferences between *sets* from an underlying basis of partial order relation have been considered before, e.g. [4] gives an overview of such extensions in one context, and we make no claim to have originated the methods below.

Definition 2 *For $C \subseteq Ag$ and $G', G'' \in \mathcal{X}^\Gamma$, we say that C* strongly prefers *the goal set G' to G'', written $G' \sqsupset_C G''$ if*

1. $G' \in \mathcal{X}(C)$.
2. $\forall i \in C,\ \exists r_i \in G' \cap G_i,\ \forall s_i \in G'' \cap G_i :\ r_i \triangleright_i s_i$.

We say that C weakly prefers *the goal set G' to G'', written $G' \succ_C G''$ if*

1. $G' \in \mathcal{X}(C)$.
2. $\forall i \in C,\ \forall s_i \in G'' \cap G_i,\ \exists r_i \in G' \cap G_i :\ r_i \triangleright_i s_i$.

If $G' \sqsupset_C G''$ this indicates that for each $i \in C$ there is a *single goal*, $r_i \in G' \cap G_i$ that i ranks higher than *every one* of the goals, s_i, that it could have satisfied within G''. If $G' \succ_C G''$ this indicates that for each $i \in C$ no matter which goal $s_i \in G'' \cap G_i$ that i may have the possibility of realising, it can identify some goal $r_i \in G' \cap G_i$ that it would rather achieve. We note that the goal set G'' is *not* required to be in the set $\mathcal{X}(C)$, although as a member of $\mathcal{X}^\Gamma$ it must belong to $\mathcal{X}(S)$ for at least one coalition $S \subseteq Ag$.

Our formulations of G' being preferred to G'' by C capture the idea that G' is a set of a goals that are feasible for and satisfy C, i.e., $G' \in \mathcal{X}(C)$, whereas for each member i of C, G'' either fails to satisfy i at all or any goal $s_i \in G_i$ that can be realised within G'' can be outranked by a goal $r_i \in G' \cap G_i$: the difference between *strong* and *weak* preference is that in the former case at least one r_i must be present in $G' \cap G_i$ that is preferred to *every* goal $s_i \in G'' \cap G$, whereas in the latter case different r_i can be used depending on the goal chosen within G''.

In terms of the scenario outlined in Example 2, for a set of students, S, having feasible choices, P and P', $P \sqsupset_C P'$, indicates that for each student in S, there is some *single* project in the pool P' that is considered to be preferable to any option that would satisfy them in in P. In contrast, $P \succ_C P'$, indicates that no matter which project s is given within P, the student will be able to identify (at least) one project in P' they would rather be allocated.

Before proceeding with some further properties of these relationships we extend our preference concepts to apply to *sets* of coalitions. While our principal interest lies in the case $\mathcal{R} = 2^C$ for $C \subseteq Ag$ since this forms the basis for our concepts of stability and core, the definition below is given in terms of arbitrary (non-empty) sets $\mathcal{R} \subseteq 2^{Ag}$.

To avoid excessive repetition we use the relational symbol $\rhd$ to indicate either $\sqsupset$ or $\succ$.

Definition 3 *Given a* set *of coalitions $\mathcal{R} \subseteq 2^{Ag}$ the binary relationship $\rhd_\mathcal{R}$ over $\mathcal{X}(\mathcal{R}) = \cup_{C \in \mathcal{R}} \mathcal{X}(C)$ is defined as $G' \rhd_\mathcal{R} G''$ if for some coalition $C \in \mathcal{R}$ it holds that $G' \rhd_C G''$. If $G' \rhd_\mathcal{R} G''$ we say that G' strongly/weakly* dominates *G'' with respect to $\mathcal{R}$, noting that $G' \rhd_\mathcal{R} G''$ implies there is some $C \in \mathcal{R}$ for which G' strongly/weakly dominates G'' with respect to C.*

The following Lemma summarises some key properties of the relations $\succ_C$ and $\sqsupset_C$.

Lemma 1

a. *If $G' \sqsupset_{\mathcal{R}} G''$ then $G' \succ_{\mathcal{R}} G''$; the converse, however, does not always hold.*

b. *If $\mathcal{R} = \{C\}$, (i.e., contains exactly one coalition), then $\rhd_{\mathcal{R}}$ induces a partial order over $\mathcal{X}(C)$.*

c. *If $\mathcal{R}$ satisfies the property that for every pair $\{C, D\}$ of coalitions in $\mathcal{R}$, $C \cap D \neq \emptyset$ then for all $\{G', G''\} \in \mathcal{X}(\mathcal{R})$ at most one of $G' \rhd_{\mathcal{R}} G''$ and $G'' \rhd_{\mathcal{R}} G'$ hold, i.e., $\rhd_{\mathcal{R}}$ is asymmetric.*

Proof: We omit all but the proof of (a), concerning which it is immediate from the definitions that $G' \sqsupset_{\mathcal{R}} G''$ implies $G' \succ_{\mathcal{R}} G''$. To see that the converse may fail to hold, consider the following QCGP:

$$\begin{array}{rcl} Ag & = & \{a_1\} \\ G & = & \{g_1, g_2, g_3, g_4\} \\ G_1 & = & G \\ \Psi & = & a_1 \Leftrightarrow ((g_1 \wedge g_3) \oplus (g_2 \wedge g_4)) \\ \rhd_1 & = & \{g_1 \rhd_1 g_2,\ g_3 \rhd_1 g_4\} \end{array}$$

Then $\{g_1, g_3\} \succ_{\{a_1\}} \{g_2, g_4\}$ since g_1 outranks g_2 and g_3 outranks g_4. It is not the case, however, that $\{g_1, g_3\} \sqsupset_{\{a_1\}} \{g_2, g_4\}$ since g_1 is not strictly preferred to g_4 and g_3 is not strictly preferred to g_2. ∎

From Lemma 1(b), the following subsets of $\mathcal{X}(C)$ are well-defined for any coalition.

Definition 4 *For any $C \subseteq Ag$, the* maximal strongly preferred goal sets with respect to C, *denoted $\mu^{\sqsupset}(C)$, are defined through*

$$\mu^{\sqsupset}(C) = \{G' \in \mathcal{X}(C) \mid \forall G'' \in \mathcal{X}(C),\ G'' \not\sqsupset_C G'\}.$$

For any $C \subseteq Ag$, the maximal weakly preferred goal sets with respect to C, *denoted $\mu^{\succ}(C)$, are defined through*

$$\mu^{\succ}(C) = \{G' \in \mathcal{X}(C) \mid \forall G'' \in \mathcal{X}(C),\ G'' \not\succ_C G'\}.$$

In the event of $\mu^{\sqsupset}(C) = \mu^{\succ}(C)$ we write simply $\mu(C)$.

We note from this definition that $G' \in \mathcal{X}(C) \setminus \mu^{\rhd}(C)$ implies there is some $G'' \in \mu^{\rhd}(C)$ such that $G'' \rhd_C G'$.

In informal terms, the sets of goal sets $\mu^{\sqsupset}(C)$ and $\mu^{\succ}(C)$ describe the *optimal* outcomes that could be realised by a coalition, C, with respect to each of the orderings $\sqsupset$ and $\succ$: if $G' \in \mu^{\rhd}(C)$ then not only is G' a feasible and satisfying choice for C (by virtue of $\mu^{\rhd}(C)$ being a subset of $\mathcal{X}(C)$), in addition, G' cannot be outranked by any other feasibly satisfying choice for C. Certainly, in the event that C *could* succeed, i.e. $\mathcal{X}(C) \neq \emptyset$, one would expect it to seek to bring about one of its optimal outcomes, i.e. some element $G' \in \mu^{\rhd}(C)$. From such a perspective the issue of whether a given $G' \in \mathcal{X}(C)$ also belongs to $\mu^{\rhd}(C)$ becomes a decision question of some interest and is one whose complexity we address in Section 4.

3. The Core and Stability

In this section we introduce some solution concepts in respect of the preferred goal sets. These solution concepts are closely based on the corresponding concepts from cooperative game theory [5].

Definition 5 *Let Γ be a* QCG *$G = \langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$ with preference relations $\langle \rhd_1, \ldots, \rhd_n\rangle$. For a coalition $C \subseteq Ag$, the* strong core *of C, denoted $\kappa^{\sqsupset}(C)$ is the set*

$$\{G' \in \mu^{\sqsupset}(C) \mid \forall C' \subset C, \forall G'' \in \mathcal{X}(C'),\ G'' \not\sqsupset_{C'} G'\}.$$

For a coalition $C \subseteq Ag$, the weak core *of C, denoted $\kappa^{\succ}(C)$ is the set*

$$\{G' \in \mu^{\succ}(C) \mid \forall C' \subset C, \forall G'' \in \mathcal{X}(C'),\ G'' \not\succ_{C'} G'\}.$$

Again, if it is the case that $\mu^{\sqsupset}(D) = \mu^{\succ}(D)$ for all $D \in 2^C$ we refer simply to the core *of C denoting this $\kappa(C)$.*

These notions of coalitional core describe one motivation for a coalition, C, to remain intact in order to bring about a given $G' \in \mu^{\rhd}(C)$. For suppose it were the case that G' did *not* belong to the set $\kappa^{\rhd}(C)$, i.e. was not an element of the core as we have defined it above. Certainly, C as it stands cannot do better than to bring about the set G' since G' is one of its optimal outcome choices, however, the fact that $G' \notin \kappa^{\rhd}(C)$, indicates that there is some *strict* subset of C' of C that has good reason to secede: C' can realise some choice G'' that its members prefer to those goals that can be achieved within G'. Thus, by analogy with the classical quantitative view of the core, we might say that the "pay-off" that the members of C' achieve by bringing out G'' is better than they would receive as part of C in bringing about G'. Our formulation of coalitional core above, not only gives rise to the obvious decision question for a coalition C and feasibly satisfying goal set G' of whether $G' \in \kappa^{\rhd}(C)$, but also motivates a rather more subtle issue: that of whether C *can* succeed and with *every* feasibly satisfying outcome for C in its core, i.e. whether $\mathcal{X}(C) = \kappa^{\rhd}(C)$ with $\mathcal{X}(C) \neq \emptyset$. For QCGPs within which C has the latter property, C may safely bring about *any* of its feasibly satisfying outcomes, G', being sure that the coalition intact cannot do "any better" than achieve G' *and* that no strict subset C' can realise outcomes it would prefer, and thence has no "rational" incentive to break away. We address both of these questions in Section 4.

A coalition attempting to realise some $G' \in \mu^{\rhd}(C) \setminus \kappa^{\rhd}(C)$ may be undermined via some $C' \subset C$ forming to bring about that goal set G'' which attests to the non-membership of G' in $\kappa^{\rhd}(C)$. Our next definitions introduce another class of methods through which C may be "at risk" in attempting to realise $G' \in \mathcal{X}(C)$. In these we consider some *set* of coalitions $\mathcal{R}$ in order to define notions of a collection of feasible outcome, $\mathcal{Y}$, being *stable*.

Definition 6 *For $\mathcal{R} \subseteq 2^{Ag}$ and $\mathcal{Y} \subseteq \mathcal{X}(\mathcal{R})$, we say that $\mathcal{Y}$ is* internally stable with respect to the set of coalitions $\mathcal{R}$ *if*

$$\forall G', G'' \in \mathcal{Y},\ \neg\exists C \in \mathcal{R} \text{ for which } G' \rhd_C G'' \text{ or } G'' \rhd_C G'.$$

The set $\mathcal{Y}$ is externally stable with respect to the set of coalitions $\mathcal{R}$ *if*

$$\forall G' \in \mathcal{X}(\mathcal{R}) \setminus \mathcal{Y}, \exists G'' \in \mathcal{Y} \text{ and } C \in \mathcal{R} \text{ s.t. } G'' \rhd_C G'.$$

The set $\mathcal{Y}$ is stable with respect to the set of coalitions $\mathcal{R}$ *if it is both internally and externally stable with respect to $\mathcal{R}$.*

Definition 7 *For $\mathcal{R} \subseteq 2^{Ag}$ and $\mathcal{Y} \subseteq \mathcal{X}(\mathcal{R})$, we say that a subset $\mathcal{Y}$ of $\mathcal{X}(\mathcal{R})$ is an* admissible goal set with respect to $\mathcal{R}$ *if $\mathcal{Y}$ is internally stable with respect to $\mathcal{R}$ and for every $G^p \in \mathcal{X}(\mathcal{R}) \setminus \mathcal{Y}$, if there is a coalition $C \in \mathcal{R}$ for which $G^p \rhd_C G^q$ for $G^q \in \mathcal{Y}$ then there is some $G^r \in \mathcal{Y}$ and coalition $D \in \mathcal{R}$ for which $G^r \rhd_D G^p$. An admissible goal set $\mathcal{Y}$ is maximal with respect to $\mathcal{R}$ if no strict superset of $\mathcal{Y}$ is admissible.*

Notice that every *stable set* with respect to $\mathcal{R}$ is also a *maximal* admissible goal set with respect to $\mathcal{R}$; however, the converse does not always hold. For example, consider $\mathcal{R} = \{C, D, E\}$, whose members are pairwise disjoint; furthermore suppose we have three goal sets $G^{(C)} \in \mathcal{X}(C)$, $G^{(D)} \in \mathcal{X}(D)$ and $G^{(E)} \in \mathcal{X}(E)$. It may be the case that

$$G^{(C)} \rhd_C G^{(D)} \rhd_D G^{(E)} \rhd_E G^{(C)}.$$

If $\mathcal{X}(\mathcal{R}) = \{G^{(C)}, G^{(D)}, G^{(E)}\}$, then the maximal admissible subset is the empty set; however, this is not externally stable: in fact this system has no stable set. This possibility motivates the following.

Definition 8 *The set of coalitions $\mathcal{R} \subseteq 2^{Ag}$ is* coherent *in the* QCGP Γ, *if every $\mathcal{Y} \subseteq \mathcal{X}(\mathcal{R})$ that defines a maximal admissible set with respect to $\mathcal{R}$ is a stable set with respect to $\mathcal{R}$.*

The sets $\mathcal{X}(\mathcal{R})$ and concepts of core and stability introduced in Definitions 5, 6 are analogous to the idea of *imputation*, core, and stability in classical coalitional games, cf. [5, Section 14.2].

As an aside, we note that the definition of the core of a succesful coalition differs from the (non-preference model) view of [10] in which "core" of a coalition, C, is non-empty if and only if $\mathcal{X}(C) \neq \emptyset$ *and* for all strict subsets C' of C, we have $\mathcal{X}(C') = \emptyset$. Suppose, however, we allow definitions of $\rhd_C$ to be *independent* of the preference relations $\langle \rhd_1, \ldots, \rhd_n \rangle$: then defining $\kappa(C)$ with respect to these we can capture the interpretation of [10], as shown in the following easy lemma.

Lemma 2 *Let Γ be a* QCGP *and define a preference relation for* coalitions $C \subseteq Ag$ *by*

$$G' \rhd_C G'' \text{ if } G' \in \mathcal{X}(C) \text{ and } \exists C' \supset C \text{ for which } G'' \in \mathcal{X}(C')$$

so that $\kappa(C)$ – the core of C – is

$$\{G' \in \mathcal{X}(C) \mid \forall C' \subset C,\ \forall G'' \in \mathcal{X}(C'),\ G'' \not\rhd_{C'} G'\}.$$

Then $\kappa(C) \neq \emptyset$ if and only if C is both minimal and successful.

Proof: If $\kappa(C) \neq \emptyset$ then it is certainly the case that C is successful since $\kappa(C) \subseteq \mathcal{X}(C)$. It must also be the case, however, that C is minimal: for otherwise we have some $C' \subset C$ with $\mathcal{X}(C') \neq \emptyset$ so that for any $G'' \in \mathcal{X}(C')$ we have $G'' \rhd_{C'} G'$ for every $G' \in \mathcal{X}(C)$ contradicting $\kappa(C) \neq \emptyset$. Similarly if C is both successful and minimal then the former yields $\mathcal{X}(C) \neq \emptyset$ while the latter indicates that for every $C' \subset C$, we have $\mathcal{X}(C') = \emptyset$ hence $\kappa(C) = \mathcal{X}(C) \neq \emptyset$ as required. ∎

We note that we could restrict notions of stability and admissibility to the case where the underlying set of coalitions $\mathcal{R}$ is simply 2^{Ag}, i.e., the set of all possible coalitions. There are, however, some disadvantages of this. Although it is certainly the case any set $\mathcal{Y} \subseteq \mathcal{X}^\Gamma$ that is internally stable with respect to 2^{Ag} will also have this property with respect to any subset of the set of coalitions $\mathcal{R} = \{C \mid \mathcal{Y} \cap \mathcal{X}(C) \neq \emptyset\}$, it may be the case that we wish to regard some goal sets as internally stable (with respect to a given set of coalitions $\mathcal{R}$) that could not be considered as such in terms of the set 2^{Ag}, i.e., if $\mathcal{Y}$ is internally stable for 2^C with some $C \subset Ag$, it does not necessarily have this property with respect to any $D \not\subseteq C$. Similarly, as regards our definitions of external stability and admissibility, using only $\mathcal{R} = 2^{Ag}$, will view some $\mathcal{Y}$ is externally stable (or admissible) that do not have this property with respect to a given coalition C or even 2^C. For example, we may have $\mathcal{Y} \subseteq \mathcal{X}(2^C)$ but with some $G'' \in \mathcal{X}(2^C) \setminus \mathcal{Y}$ for which no goal set $G' \in \mathcal{Y}$ and coalition $D \in 2^C$ gives $G' \rhd_D G''$ (hence $\mathcal{Y}$ is not externally stable with respect to the set of coalitions 2^C) even though $G' \rhd_{2^{Ag}} G'$: the latter preference being exhibited by some $D \not\subseteq C$.

In total, the choice of $\mathcal{R} = 2^{Ag}$ may be too restrictive sensibly to consider *internal* stability, but rather too general to use as a basis for *external* stability.

We introduce some further notation prior to proving some basic properties of these structures. For $\mathcal{R} \subseteq 2^{Ag}$, we define the sets

$$\begin{array}{lcl} \iota^\rhd(\mathcal{R}) & \hat{=} & \{\mathcal{Y} \subseteq \mathcal{X}(\mathcal{R}) \mid \mathcal{Y} \text{ is internally stable w.r.t. } \mathcal{R}\} \\ \eta^\rhd(\mathcal{R}) & \hat{=} & \{\mathcal{Y} \subseteq \mathcal{X}(\mathcal{R}) \mid \mathcal{Y} \text{ is externally stable w.r.t. } \mathcal{R}\} \\ \sigma^\rhd(\mathcal{R}) & \hat{=} & \iota^\rhd(\mathcal{R}) \cap \eta^\rhd(\mathcal{R}) \\ \delta^\rhd(\mathcal{R}) & \hat{=} & \{\mathcal{Y} \subseteq \mathcal{X}(\mathcal{R}) \mid \mathcal{Y} \text{ is maximally admissible w.r.t. } \mathcal{R}\} \end{array}$$

Thus, $\sigma^\rhd(\mathcal{R})$ defines the set of all stable goal sets with respect to $\mathcal{R}$.

3.1. Properties of Stable Sets and the Core

The relationships between these sets are summarised in the next lemma.

Lemma 3 *For any $\mathcal{R} \subseteq 2^{Ag}$,*

a. $\iota^{\succ}(\mathcal{R}) \subseteq \iota^{\sqsupset}(\mathcal{R})$

b. $\eta^{\sqsupset}(\mathcal{R}) \subseteq \eta^{\succ}(\mathcal{R})$

c. $\iota^{\succ}(\mathcal{R}) \cap \eta^{\sqsupset}(\mathcal{R}) \subseteq \sigma^{\rhd}(\mathcal{R}) \subseteq \iota^{\sqsupset}(\mathcal{R}) \cap \eta^{\succ}(\mathcal{R})$.

Proof: The proofs of (a) and (b) follow directly from Lemma 1(a): if $\mathcal{Y} \in \iota^{\succ}(\mathcal{R})$ then for any G', $G'' \in \mathcal{Y}$ and $D \in \mathcal{R}$, by definition, neither $G' \succ_D G''$ nor $G'' \succ_D G'$ hold, thus neither $G' \sqsupset_D G''$ nor $G'' \sqsupset_D G'$, i.e., $\mathcal{Y} \in \iota^{\sqsupset}(\mathcal{R})$. Similarly $\mathcal{Y} \in \eta^{\sqsupset}(\mathcal{R})$ indicates that for each $G' \in \mathcal{X}(\mathcal{R}) \backslash \mathcal{Y}$ we have some $G'' \in \mathcal{Y}$ and $D \in \mathcal{R}$ for which $G'' \sqsupset_D G'$, which gives $G'' \succ_D G'$ and thus $\mathcal{Y} \in \eta^{\succ}(\mathcal{R})$. Finally, (c) is immediate from (a), (b), and the definition of $\sigma^{\rhd}(\mathcal{R})$. ∎

Regarding $\kappa^{\rhd}(C)$ for $C \subseteq Ag$, we observe that parts (b–d) of our next result can be seen as an analogous result to [5, Proposition 279.2].

Lemma 4 *For any $C \subseteq Ag$,*

a. $\kappa^{\succ}(C) \subseteq \kappa^{\sqsupset}(C) \subseteq \mu^{\sqsupset}(C) \subseteq \mu^{\succ}(C)$.

b. For every $\mathcal{Y} \in \sigma^{\rhd}(2^C)$, $\kappa^{\rhd}(C) \subseteq \mathcal{Y}$.

c. If $\mathcal{Y}$ and $\mathcal{Z}$ are distinct sets in $\sigma^{\rhd}(2^C)$ then $\mathcal{Y} \not\subseteq \mathcal{Z}$ and $\mathcal{Z} \not\subseteq \mathcal{Y}$.

d. If $\kappa^{\rhd}(C) \in \sigma^{\rhd}(2^C)$ then $\sigma^{\rhd}(2^C) = \{\kappa^{\rhd}(C)\}$.

Proof: Omitted. ∎

Lemma 5 $\sigma^{\rhd}(\{C\}) = \{\mu^{\rhd}(C)\}$.

Proof: That $\mu^{\rhd}(C)$ is a stable set w.r.t C is immediate from the definition of $\mu^{\rhd}(C)$: if $\{G', G''\} \in \mu^{\rhd}(C)$ then certainly neither $G' \rhd_C G''$ nor $G'' \rhd_C G'$ hold and, thus, $\mu^{\rhd}(C)$ is internally stable w.r.t. C. In addition if we consider any $G' \in \mathcal{X}(C) \setminus \mu^{\rhd}(C)$ then since $G' \notin \mu^{\rhd}(C)$ there must be some $G'' \in \mu^{\rhd}(C)$ for which $G'' \rhd_C G'$ establishing that $\mu^{\rhd}(C)$ is externally stable w.r.t. C. To see that $\mu^{\rhd}(C)$ is the unique stable set w.r.t. C, it suffices to observe that from external stability $\mu^{\rhd}(C) \subseteq \mathcal{Y}$ for any $\mathcal{Y} \in \sigma^{\rhd}(\{C\})$, which suffices to ensure uniqueness. ∎

We note that Lemma 5 indicates that *every* QCGP has *at least one subset of* 2^G which is stable. This set may, of course, simply be the empty set: in the event $C \subseteq Ag$ having no satisfying and feasible goal set, i.e., $\mathcal{X}(C) = \emptyset$, then $\mu^{\rhd}(C) = \emptyset$ which is stable. Since $\mu^{\rhd}(C)$ is the unique maximal admissible set with respect to C, it follows that in every QCGP the set of coalitions $\mathcal{R} = \{C\}$ is *coherent* for every $C \subseteq Ag$. Another interesting consequence of Lemma 5 is the following, which asserts that there is a *non-empty* stable set of goal sets within the QCGP Γ if and only if some coalition has a feasible and satisfying goal set.

Theorem 1 *Let Γ be a* QCGP

$$\Gamma = \langle G, Ag, G_1, \ldots, G_n, \Psi, \rhd_1, \ldots, \rhd_n \rangle.$$

There exists $\mathcal{R} \subseteq 2^{Ag}$ for which $\sigma^{\rhd}(\mathcal{R}) \neq \{\emptyset\}$ if and only if there exists some coalition $C \subseteq Ag$ for which $\mathcal{X}(C) \neq \emptyset$.

Proof: ($\Rightarrow$) Suppose $C \subseteq Ag$ is such that $\mathcal{X}(C) \neq \emptyset$. From Lemma 1(b), the set $\mu^{\rhd}(C)$ is well-defined and non-empty. From Lemma 5, $\mu^{\rhd}(C)$ is a stable set with respect to C. Thus if for some $C \subseteq Ag$, it holds that $\mathcal{X}(C) \neq \emptyset$ then we identify $\mathcal{R} \subseteq 2^{Ag}$, i.e., $\mathcal{R} = \{C\}$, for which $\sigma^{\rhd}(\mathcal{R}) \neq \emptyset$. ($\Leftarrow$) Suppose $\mathcal{R} \subseteq 2^{Ag}$ is such that $\sigma^{\rhd}(\mathcal{R}) \neq \{\emptyset\}$, and let $\mathcal{Y}$ be a stable set with respect to the set of coalitions $\mathcal{R}$. By definition, $\mathcal{Y} \subseteq \mathcal{X}(\mathcal{R}) = \bigcup_{C \in \mathcal{R}} \mathcal{X}(C)$, and thus $\mathcal{Y} \neq \emptyset$ implies that for some $C \subseteq Ag$, $\mathcal{X}(C) \neq \emptyset$. ∎

One of the main points of interest regarding Theorem 1 is that it allows the computational complexity of the following decision problem to be determined exactly.

> <u>NON-TRIVIAL STABLE SET</u>: (NTSS)
> *Instance*: QCGP $\langle G, Ag, G_1, \ldots, G_n, \Psi, \rhd_1, \ldots, \rhd_n \rangle$.
> *Question*: Does there exists $\mathcal{R} \subseteq 2^{Ag}$ and $\mathcal{Y} \subseteq \mathcal{X}(\mathcal{R})$ such that $\mathcal{Y} \neq \emptyset$ and $\mathcal{Y} \in \sigma^{\rhd}(\mathcal{R})$?

Corollary 1 NTSS *is* NP*-complete.*

Proof: From Theorem 1, an instance $\Gamma = \langle G, \rhd_1, \ldots, \rhd_n \rangle$ of NTSS is accepted if and only if some coalition is successful, i.e., NTSS is equivalent to deciding if G is accepted as instance of the problem NON-EMPTY GAME (the complement of the problem EMPTY GAME which accepts instances G of QCGs for which no coalition succeeds). The decision problem EMPTY GAME was shown to be co-NP-complete in [10, Theorem 35], and thus its complement decision problem – NON-EMPTY GAME – is NP-complete. This suffices to deduce that NTSS is also NP-complete. ∎

We note, in passing, that from [10, Corollary 36] it is immediate that NTSS remains NP–complete even if we restrict instances to those which are coalition monotonic, i.e. for which $\Psi(Ag, G)$ has the property that if $\Psi[C, G'] = \top$ and $C \subseteq D$ then $\Psi[D, G'] = \top$ also.

While the NP-hardness of NTSS is perhaps unsurprising, it is less obvious that the problem belongs to NP. Theorem 1, however, provides a decision method in NP that obviates any requirement to consider sets of coalitions $\mathcal{R} \subseteq 2^{Ag}$ and sets of goal sets $\mathcal{Y} \subseteq 2^G$ whose size is superpolynomial in $n+m$, by characterising the existence of non-empty stable sets in terms of the existence of successful coalitions.

3.2. Optimal Goal Sets

The maximal preferred sets – $\mu^{\rhd}(C)$ – provide one approach to defining what is meant by a goal set being optimal for a coalition C: if $G' \in \mu^{\rhd}(C)$ then there is no choice $G'' \in \mathcal{X}(C)$ that will result in *every* member of C being able to *strictly improve* upon the goals is realise within G'. Another widely studied concept of optimality is of course that of *Pareto optimality* [5, p.305].

Definition 9 *A goal set $G' \in \mathcal{X}(C)$ is* Pareto optimal *(with respect to $\sqsupset_C$) if for all other $G'' \in \mathcal{X}(C)$, should it be the case that for some i in C there is a goal $r_i \in G'' \cap G_i$ which is strictly preferred to every $s_i \in G' \cap G_i$, then there is some $s_j \in G' \cap G_j$ which is strictly preferred to every $r_j \in G'' \cap G_j$. A goal set $G' \in \mathcal{X}(C)$ is* Pareto optimal *(with respect to $\succ_C$) if for all other $G'' \in \mathcal{X}(C)$, should it be the case that for some $i \in C$ for every $s_i \in G' \cap G_i$ there is some $r_i \in G'' \cap G_i$ for which $r_i \rhd_i s_i$, then there is some $j \in C$ in which: for every $r_j \in G'' \cap G_j$ there exists $s_j \in G' \cap G_i$ with $s_j \rhd_j r_j$.*

Thus G' is a Pareto optimal goal set for a coalition C if it is feasible for and satisfies each member of C, but for any other goal set G'' that is feasible and satisfies C, if some agent can realise a more preferred goal with G'' this will be at the expense of another agent having to accept a goal that it prefers less to its optimal goals within G'. We note the distinction between G' being Pareto optimal and a maximal preferred goal set (with either $\succ$ or $\sqsupset$ as the underlying preference relations). As counterparts to $\mu^{\rhd}$, let

$$\pi^{\rhd}(C) \ = \ \{G' \in \mathcal{X}(C) \mid G' \text{ is Pareto optimal w.r.t } \rhd_C\}$$

It is *not* necessarily the case that $\pi^{\rhd}(C) = \mu^{\rhd}(C)$, although it is easily shown via the respective definitions that for every QCGP and coalition $C \subseteq Ag$, we have $\pi^{\rhd}(C) \subseteq \mu^{\rhd}(C)$. Consider, however, the following.

Example 3 *Let Γ be a* QCGP *with $Ag = \{a_1, a_2\}$, $G = \{g_1, g_2, g_3\}$, $G_1 = \{g_1, g_2\}$, $G_2 = \{g_3\}$, $\rhd_1 = \{g_1 \rhd_1 g_2\}$, $\rhd_2 = \emptyset$, and $\Psi(Ag, G) = a_1 a_2 g_3 (g_1 \vee g_2)$. For this* QCGP *it is easily checked that*

$$\mathcal{X}(C) = \begin{cases} \{\{g_1, g_3\},\ \{g_2, g_3\}\} & \text{if} \quad C = \{a_1, a_2\} \\ \emptyset & \text{if} \quad C \neq \{a_1, a_2\} \end{cases}$$

Furthermore,

$$\begin{aligned} \mu^{\sqsupset}(\{a_1, a_2\}) &= \mu^{\succ}(\{a_1, a_2\}) \\ &= \mathcal{X}(\{a_1, a_2\}) \\ &= \{\{g_1, g_3\},\ \{g_2, g_3\}\} \end{aligned}$$

On the other hand, the Pareto optimal sets are

$$\pi^{\sqsupset}(\{a_1, a_2\}) = \pi^{\succ}(\{a_1, a_2\}) = \{\{g_1, g_3\}\}$$

In both cases strict subsets of $\mu^{\rhd}(\{a_1, a_2\})$ which also contains $\{g_2, g_3\}$. The goal set $\{g_1, g_3\}$ strictly improves the goal that can be realised by a_1 ($g_1 \rhd_1 g_2$) but does not leave a_2 less *satisfied than before. Thus for this example $\pi(\{a_1, a_2\}) \subset \mu(\{a_1, a_2\})$.*

4. Decision Problems for QCGPs

We now consider four decision problems associated with QCGPs. Although in principle one could define distinct variants of these in terms of the two different preference relationships – $\sqsupset$ and $\succ$ – this turns out to be unnecessary. For the complexity classifications that we prove, all the lower bound results, (i.e., hardness proofs), construct instances for which $G' \sqsupset_C G''$ if and only if $G' \succ_C G''$ for every coalition C. For upper bounds arguments, (i.e., membership of a given class), it is easily verified given a QCGP, coalition C, and goal sets G', G'' that the tests $G' \succ_C G''$ and $G' \sqsupset_C G''$ can both be accomplished easily. The one exception arises in the result proved in Theorem 4, a $\mathcal{C}$-completeness result for which establishing membership in $\mathcal{C}$ is a non-trivial argument involving differing constructions dependent on the exact preference relation employed.

The first problem we consider is that of whether a set of goals is in the core of a coalition.

> CORE MEMBERSHIP: (CM)
> *Instance*: QCGP $\langle G, Ag, G_1, \ldots, G_n, \Psi, \rhd_1, \ldots, \rhd_n \rangle$, coalition $C \subseteq Ag$, goal set $G' \subseteq G$.
> *Question*: Is $G' \in \kappa^{\rhd}(C)$?

Theorem 2 CM *is co-*NP*-complete.*

Proof: CM is in co-NP since $G' \in \kappa^{\rhd}(C)$ if and only if

$$\forall G'',\ D \subseteq C \ (G' \in \mathcal{X}(C) \text{ and } \neg(G'' \rhd_D G')$$

That is, if G' satisfies and is feasible for C, and is not dominated by any other goal set G'' with respect to any subset of C. We note that if $G' \notin \mu^{\rhd}(C)$ (and hence cannot belong to the core) then it will be dominated by some $G'' \in \mu^{\rhd}(C)$ with respect to C. By quantifying over all subsets of C, i.e., not simply strict subsets, this case is detected. To complete the proof, we use a reduction from UNSAT. Let $\Phi(x_1, \ldots, x_n)$ be an instance of UNSAT. We form an instance $\langle \Gamma_\Phi, C, G' \rangle$ of CM for which $G' \in \kappa^{\rhd}(C)$ if and only if $\Phi(x_1, \ldots, x_n)$ is unsatisfiable.

The QCGP Γ_Φ has $Ag = \{a_1, \ldots, a_n, a_{n+1}\}$, $G_i = \{g_i^{\top}, g_i^{\perp}, g_i^{\min}\}$ so that $G = \cup_{i=1}^{n+1} G_i$. The characteristic function $\Psi(Ag, G)$ is given as

$$\begin{aligned} &\left(\bigwedge_{i=1}^{n+1} a_i g_i^{\min} (\neg g_i^{\top})(\neg g_i^{\perp})\right) \ \vee \\ &\qquad (\neg a_{n+1}) \Phi(a_1 (g_1^{\top} \vee \neg g_1^{\perp}), \ldots, a_n (g_n^{\top} \vee \neg g_n^{\perp})) \end{aligned}$$

The preference relation $\rhd_i$ contains exactly two elements: $\{g_i^{\top} \rhd_i g_i^{\min},\ g_i^{\perp} \rhd_i g_i^{\min}\}$. Finally we set $C = Ag$ and $G' = \cup_{i=1}^{n+1} \{g_i^{\min}\}$ to form the instance $\langle \Gamma_\Phi, C, G' \rangle$ of CM.

We note that $G' \sqsupset_C G''$ if and only if $G' \succ_C G''$ for all $C \subseteq Ag$, and hence $\kappa^{\succ}(C) = \kappa^{\sqsupset}(C)$ for every coalition C. Now, $G' = \bigcup_{i=1}^{n+1}\{g_i^{\min}\} \in \kappa(Ag)$ if and only if $\Phi(x_1, \ldots, x_n)$ is unsatisfiable. ∎

The next problem we consider is whether or not a goal set is maximally preferred by a coalition, i.e., whether this goal set both satisfies every member of the coalition, and there is no other goal set that satisfies the coalition that is strictly preferred by it.

MAXIMAL GOAL SET: (MGS)
Instance: QCGP $\langle G, Ag, G_1, \ldots, G_n, \Psi, \rhd_1, \ldots, \rhd_n\rangle$, coalition $C \subseteq Ag$, goal set $G' \subseteq G$.
Question: Is $G' \in \mu^{\rhd}(C)$?

In addition, we consider the problem of determining whether a goal set is Pareto optimal.

PARETO OPTIMAL GOAL SET: (PO)
Instance: QCGP $\langle G, Ag, G_1, \ldots, G_n, \Psi, \rhd_1, \ldots, \rhd_n\rangle$, coalition $C \subseteq Ag$, goal set $G' \subseteq G$.
Question: Is $G' \in \pi^{\rhd}(C)$?

Corollary 2

a) MGS *is co*-NP-*complete.*

b) PO *is co*-NP-*complete.*

Proof: For (a), membership is immediate from the fact that $\langle \Gamma, C, G'\rangle$ is accepted as an instance of MGS if and only if: $G' \in \mathcal{X}(C)$ and for each G'', if $G'' \in \mathcal{X}(C)$ then it is not the case that $G'' \rhd_C G'$, a test easily accomplished by a co-NP algorithm. For (b), membership is established from the relation $\langle \Gamma, C, G'\rangle$ is accepted as an instance of PO if and only if: $G' \in \mathcal{X}(C)$ and for all $G'' \subseteq G$ for which $G'' \in \mathcal{X}(C)$:

$$(\exists i \in C \text{ and } s_i \in G'' \cap G_i \text{ with } s_i \rhd_i r_i,\ \forall r_i \in G' \cap G_i) \Rightarrow (\exists j \in C \text{ and } r_j \in G' \cap G_j \text{ with } r_j \rhd_j s_j,\ \forall s_j \in G'' \cap G_j)$$

An identical reduction from UNSAT serves to prove co-NP-hardness in both cases. We use a similar reduction from instances $\Phi(x_1, \ldots, x_n)$ of UNSAT as that of Theorem 2 but with $Ag = \{a_1, \ldots, a_n\}$, G_i, $\rhd_i$ as before and $\Psi(Ag, G)$ in Γ_Φ given by

$$\left(\bigwedge_{i=1}^{n} a_i g_i^{\min}(\neg g_i^{\top})(\neg g_i^{\bot})\right) \vee \Phi(a_1(g_1^{\top} \vee \neg g_1^{\bot}), \ldots, a_n(g_n^{\top} \vee \neg g_n^{\bot}))$$

We set $C = Ag$ and $G' = \cup_{i=1}^{n}\{g_i^{\min}\}$. It is clearly the case that $G' \in \mathcal{X}(Ag)$. Furthermore by a similar argument to that of Theorem 2 $\mathcal{X}(Ag) = \{G'\}$ if and only if $\Phi(x_1, \ldots, x_n)$ is unsatisfiable. As we noted earlier $\mathcal{X}(C)$ containing a single goal set indicates that this set is maximal, i.e., in $\mu(C)$. Equally G' is Pareto optimal if and only if $\Phi(x_1, \ldots, x_n)$ is unsatisfiable. ∎

The result of Corollary 2(a), indicates one potential difficulty for C considering whether or not to realise G'. Although given $\Psi(Ag, G)$ it may be efficiently checked (in terms of the formula size) that $\Psi[C, G'] = \top$ and that G' satisfies each member of C, unless significant computational effort is expended, it may not be clear as to whether G' is a "best" outcome achievable by C. In the same way, even if this is guaranteed, Theorem 2 presents further difficulties in that C's realisation of G' is subject to the threat of $C' \subset C$ forming to realise a set of outcomes that it prefers to those offfered in G' if $G' \notin \kappa^{\rhd}(C)$. In this way a coalition C may face a non-trivial strategic choice in planning whether to bring about a given choice G': namely, if C should ignore the possibility that G' may not be "optimal" thereby avoiding the significant computational effort that might be required to validate $G' \in \mu^{\rhd}(C)$ but, in doing so engendering both the risk that G' could be improved upon and the possibility that some strict subset C' may secede on the grounds that $G' \notin \kappa^{\rhd}(C)$. Of course such considerations would be redundant if one could guarantee both $\mathcal{X}(C) \neq \emptyset$ – C *can* succeed – and $\mathcal{X}(C) = \kappa^{\rhd}(C)$ – every feasibly satisfying set of outcomes is optimal and not subject to attack by any strict subset of C. Such questions form the basis of the problem we introduce as *Core Completeness*. As might be expected, this turns out to be complete for a complexity class – D^p the class of languages expressible as the intersection of a language in NP with a language in co-NP – considered to be "harder" than either NP or co-NP.

We note that via near identical constructions to those of Theorem 2 and Corollary 2 it is a trivial matter to show that given $\langle \Gamma, C\rangle$ the problem of deciding $\mathcal{X}(C) = \mu(C)$ is co-NP-complete: use the construction of Corollary 2 together with the observation that for this construction $|\mathcal{X}(Ag)| = 1$, (thence giving $\mathcal{X}(C) = \mu(C)$), if and only if Φ is unsatisfiable.

Our observations earlier that the sets of Pareto optimal goal sets for a given coalition C may be a *strict subset* of the set of maximally preferred coalitions for C, motivates the following decision problem

MAXIMAL ONLY GOAL SET: (MOGS)
Instance: QCGP $\langle G, Ag, G_1, \ldots, G_n, \Psi, \rhd_1, \ldots, \rhd_n\rangle$, $C \subseteq Ag$, $G' \subseteq G$.
Question: Is $G' \in \mu^{\rhd}(C) \setminus \pi^{\rhd}(C)$?

Theorem 3 MOGS *is* D^p*–complete.*

Proof: Omitted. ∎

To conclude, consider the following problem.

CORE COMPLETENESS: (CC)
Instance: QCGP $\langle G, Ag, G_1, \ldots, G_n, \Psi, \rhd_1, \ldots, \rhd_n\rangle$, $C \subseteq Ag$.
Question: Is $\kappa^{\rhd}(C) = \mathcal{X}(C)$ and $\mathcal{X}(C) \neq \emptyset$?

Thus, CC is concerned with the question of whether C is successful *and* every feasible goal set that satisfies each member of C is in the core. The proof of this result is omitted: the D^p-hardness proof is straightforward, but establishing membership of D^p involves a rather elaborate construction, for which we do not have space here.

Theorem 4 *CC is D^p-complete.*

5. Discussion & Further Work

The model of Qualitative Coalitional Games with Preferences that we have introduced has largely been considered with respect to "classical" concepts from Game Theory, e.g. we have formulated analogues of "core", "stability", and "minimality". While the fact that it is possible sensibly to define such concepts for QCGPs provides some indication that the formalism is sufficiently powerful, the investigation of, say, solution concepts in QCGPs that originate independently of game-theoretic ideas would be of some interest.

We note also that our analysis of decision problems arising from QCGPs has concentrated on issues of computational complexity, and less so on "positive" algorithmic aspects. Thus, to choose just two examples from the many questions that may merit further study, we have: whether there are classes of propositional function for which those decision questions that are intractable in general, admit efficient algorithms; and, the examination of feasible negotiation mechanisms by which coalitions with particular properties can be encouraged to form. We note several points about the first of these. Although we have not provided details, the reductions used to obtain our hardness results, typically procede by constructing the feasibility function, Ψ, in a QCGP via a propositional formula, Φ presented as part of a hard satisfiability related decision problem. It is, of course, well-known that there are a number of restricted classes of formula whose associated satisfiability problems range from bordering on the trivial, e.g. monotone and Horn clause formulae, to allowing polynomial-time algorithms, e.g. 2-CNF formulae. Hence we can raise two related questions of some interest: are there restricted classes of formula which admit efficient algorithmic solutions for any of the decision problems addressed in this article; and to what extent do such classes allow descriptions of "realistic" QCGP contexts.

If we consider the (non-preference) model of QCGs as introduced in [10] the answer to the first of these questions seems far from clear: we noted earlier that for QCGs imposing a monotonicity condition, in general, has no effect on a problem's computational complexity. Similarly there are some indications to suggest that restricting the propositional formulae for QCGs to be Horn clause forms may fail to result in tractable instances. The principal reason why classes of formulae with efficient satisfiability methods, do not always yield similar efficiencies for QCGs, is that the concept of a coalition, C, "succeeding" depends not only on there being a set of goals, G', for which $\Psi[C, G'] = \top$ but also on the requirement for such a set to satisfy C. For the decision problems we have examined above, such as MGS, it open whether one can exploit these requirements to construct proofs of intractability for special cases.

6. Conclusions & Related Work

We have introduced Qualitative Coalitional Games with Preferences, a variation of Qualitative Coalitional Games in which agents are assumed to have preferences over goals. We defined a number of solution concepts for such games, established some properties of these solution concepts, and investigated their computational complexity.

Probably closest to our work is that of Conitzer and Sandholm, who investigated the complexity of determining non-emptiness of the core in a subclass of conventional coalitional games [2]. By assuming superadditivity, they were able to derive a succinct representation of characteristic functions, and proved that determining non-emptiness of the core assuming this representation was NP-complete, irrespective of whether or not utility was transferable. Bilbao and colleagues derived a number of complexity results for other subclasses of cooperative games by interpreting these games over combinatorial structures of various kinds (e.g., minimal spanning trees) [1].

A number of approaches to coalition structure generation and related problems have been described in the literature [8, 9, 7]. For example, Shehory and Kraus developed algorithms for coalition structure formation in which agents were modelled as having different capabilities, and were assumed to benevolently desire some overall task to be accomplished, where this task had some complex structure [8, 9]. Sandholm and colleagues developed algorithms to find optimal coalition structures within some given ratio bound k of optimal [7].

Acknowledgements

The work reported in this paper was carried out with the support of EPSRC Grant GR/R60836/01.

References

[1] J. Bilbao, J. Fernández, and J. López. Complexity in cooperative game theory. Manuscript.

[2] V. Conitzer and T. Sandholm. Complexity of determining nonemptiness of the core. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-03)*, pages 613–618, Acapulco, Mexico, 2003.

[3] P. E. Dunne and M. Wooldridge Qualitative coalitional reasoning with preferences. Technical Report, ULCS-04-008, Dept. of Computer Science, Univ. of Liverpool, March 2004.

[4] J. Y. Halpern Defining relative likelihood in partially-ordered preferential structures. *Jnl. of Artificial Intelligence Research*, 7, 1–24, 1997

[5] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press: Cambridge, MA, 1994.

[6] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley: Reading, MA, 1994.

[7] T. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohmé. Coalition structure generation with worst case guarantees. *Artificial Intelligence*, 111(1–2):209–238, 1999.

[8] O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1-2):165–200, 1998.

[9] O. Shehory and S. Kraus. Feasible formation of stable coalitions among autonomous agents in non-super-additive environments. *Computational Intelligence*, 15(33):218–251, 1999.

[10] M. Wooldridge and P. E. Dunne. The computational complexity of qualitative coalitional games. Technical Report ULCS-04-007, Dept. of Computer Science, Univ. of Liverpool, 2003. (to appear, *Artificial Intelligence*, 2004).