Group Dominant Strategies

Ola Rozenfeld and Moshe Tennenholtz
Technion–Israel Institute of Technology
Haifa 32000, Israel

Abstract

We introduce a new solution concept for complete information games, which we call equilibrium in group dominant strategies. This concept is the strongest of all known solution concepts so far, since it encompasses both the ideas behind the concepts of dominant strategies and strong equilibrium. Because of its strength, a solution in group dominant strategies does not exist in any interesting game; however, as we show, such solutions can be achieved in various rich settings with the use of mediators.

Introduction

A *finite game in strategic form* is a tuple $\Gamma = \langle N, \{A_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ where:

- $N = \{1, \ldots, n\}$ is a finite set of players.
- For each player $i \in N$, A_i is a finite non-empty set of actions (or strategies, we use the terms interchangeably) available to player i.
- For $S \subseteq N$, A_S denotes $\prod_{i \in S} A_i$, and A_{-S} denotes $\prod_{i \in N \setminus S} A_i$. A_N is denoted by A.
- For each player $i \in N$, $u_i : A \to \Re$ is a utility function, which represents the "contentment" of the player with each specific strategy profile.
- Let $a \in A$. We will sometimes write a as (a_i, a_{-i}) for $i \in N$ and as (a_S, a_{-S}) for $S \subseteq N$.

One of the most basic questions of game theory is: given a game in strategic form, what is the solution of the game? Basically, by a "solution" we mean a *stable* strategy profile which can be proposed to all agents, in a sense that no rational agent would want to deviate from it. Many solution concepts for games have been studied, differing mainly by the assumptions that a rational agent would have to make about the rationality of other agents. For example, probably the most well known solution concept for games is the Nash equilibrium:

A profile of actions $a \in A$ is a *Nash equilibrium* (NE) if

$$\forall i \in N \; a_i \in br_i(a_{-i})$$

A shortened version of this paper appeared in the proceedings of WINE 2008, Patras, Greece.

Here, $br_i(a_{-i})$ for $i \in N$, $a_{-i} \in A_{-i}$ denotes $\arg\max_{a_i \in A_i} \{u_i(a_i, a_{-i})\}$ (the set of *best responses* of i to a_{-i}).

There are two basic problems with the Nash equilibrium as a solution concept for games:

Problem 1: A NE guarantees absence of profitable deviations to a player only in the case that all the other players play according to the suggested profile; in the case where even one of the other players deviates, we have no such guarantees. So, the assumption that this concept requires about the rationality of other players is: all the other players will stick to their prescribed strategies. But why should a rational player make that assumption?

The following stability concept takes this problem into account: A profile of actions $a \in A$ is an *equilibrium in weakly dominant strategies* if

$$\forall i \in N, b_{-i} \in A_{-i} \quad a_i \in br_i(b_{-i})$$

The above definition strengthens the concept of NE by taking care of the aforementioned problem: no unilateral deviation *can ever be* beneficial, no matter what other players do; in other words, it requires no assumptions on the rationality of other players.

Problem 2: A NE does not take into account joint deviations by coalitions of players. We usually assume that an individual will deviate from a profile if she has an available strategy that strictly increases her income. In some settings it would be natural to assume also that a group of individuals will deviate if they have an available joint strategy that strictly increases the income of each group member. For example, consider the famous Prisoner's Dilemma game:

	C	D
C	4,4	0,6
D	6,0	1,1

The strategy profile (D, D) is a NE and even an equilibrium in weakly dominant strategies; however, it is not stable in the sense that if both players deviate to (C, C), the income of each one of them will increase. The following stability concept by (Aumann, 1959) deals with this problem:

A profile of actions $a \in A$ is a *strong equilibrium* (SE) if

$$\forall S \subseteq N \quad a_S \in br_S(a_{-S})$$

Here, the concept of best response strategy is extended to multiple players as follows: for $S \subseteq N$ and $a_{-S} \in A_{-S}$,

$br_S(a_{-S})$ denotes the set of best responses of S to a_{-S}:

$$br_S(a_{-S}) = \left\{ \begin{array}{l} a_S \in A_S | \forall b_S \in A_S\ \exists i \in S : \\ u_i(b_S, a_{-S}) \leq u_i(a_S, a_{-S}) \end{array} \right\}$$

The concept of strong equilibrium indeed takes care of Problem 2; however, it again does not take Problem 1 into account. What we would ideally like to have is a solution concept that has neither of these problems: we would like to assume that players are able to cooperate for mutual benefit, and on the other hand we would also like to assume nothing about the actions of the other players. These requirements may seem conflicting. Note that simply saying that we are interested in a profile $a \in A$ that is both a SE and an equilibrium in weakly dominant strategies is not enough: for games with more than 2 players, we would have no guarantees about the absence of joint deviations for players 1 and 2, in the case that player 3 deviated.

This brings us to the stability concept that we wish to present: a profile of actions $a \in A$ is an *equilibrium in group (weakly) dominant strategies* (GDS) if

$$\forall S \subseteq N, b_{-S} \in A_{-S} \quad a_S \in br_S(b_{-S})$$

Existence of a GDS implies, for each player, that no matter what the other players choose, and no matter with whom can she unite in making her decision, they will not find a joint strategy that will be better to all of them than the proposed one. And thus, if a GDS exists in a given game, we can safely declare it to be the solution of the game. However, a GDS does not exist in any game that has ever been a subject of interest. This is not surprising, since the concept is so strong that its mere existence renders any game not interesting. For this reason, the concept was never a subject of exploration in complete information games. In incomplete information games the concept is known under the name of *group strategy proofness* and is widely studied, because in some cases such solutions can be indeed implemented by mechanism design. However, the whole approach of mechanism design is not applicable to complete information games – although we would indeed want to assume the existence of an interested party, we don't want to give it the power to design the game.

An interested party who wishes to influence the behavior of agents in a (complete information) game, which is not under his control, will be called a mediator. This concept is highly natural; in many systems there is some form of reliable party or administrator who is interested in a "good" behavior of the system. Many kinds of mediators have been studied in the literature, differing by their power in influencing the game (see e.g.(Mas-Colell *et al.*, 1995; Jackson, 2001; Aumann, 1974; Myerson, 1986)). The less power we assume on the mediator, the more applicable the positive results will be to the real world. For example, if we assume that a mediator is able to observe the chosen strategies of the players and issue arbitrarily large fines for deviating from a proposed strategy profile, then, on one hand, such mediator will trivially be able to implement any profile as a very stable solution (e.g. GDS); on the other hand, though, this model will not be applicable to almost any real life multi-agent encounter. For this reason, as the focus of this paper is to study the power of mediators in establishing equilibrium in group dominant strategies, we make some restricting assumptions: the mediator cannot design a new game, cannot enforce agents' behavior, cannot enforce payments by the agents, and cannot prohibit strategies available to the agents.

In the rich literature about mediators, two different kinds of mediators exist that adhere to our restricting assumptions: *routing mediators* and *k-implementation*. K-implementation was introduced by (Monderer and Tennenholtz, 2004). There, a mediator is a reliable authority who can observe the strategies selected by the players and commit to non-negative monetary payments based on the selected profile. Obviously, by making sufficiently big payments one can implement any desirable outcome. The question is: what is the cost of implementation? A major point in k-implementation is that monetary offers need not necessarily materialize when following desired behaviors; the promise itself might suffice. In particular, (Monderer and Tennenholtz, 2004) show that any NE of a game can be implemented as an equilibrium in dominant strategies with 0-cost.

Routing mediators were introduced by (Rozenfeld and Tennenholtz, 2007), continuing the work of (Monderer and Tennenholtz, 2006). A routing mediator is a reliable authority which can play the game on behalf of players who give it such right. Such mediator devises a conditional contract that he proposes to all players to sign: in this contract, the mediator specifies exactly which actions he will take on behalf of the players who sign the contract, given every possible combination of actions by players who do not sign it. If a player signs the contract, the mediator is then committed to playing the game on behalf of that player by the contract specifications. So, in essence, the mediator adds a new strategy that is available to each player – to sign the contract; the payoffs of this new game are specified exactly by the contract he offers. Note that no matter which players sign the contract, in the end a strategy profile from the original game is played, and the payoffs are not altered. [1] For example, consider such a mediator in the Prisoner's Dilemma game. The mediator offers the agents the following protocol: if both agents agree to use his services then he will play C on behalf of both agents. However, if only one agent agrees to use his services then he will play D on behalf of that agent. The mediator's protocol generates a new game, where a new strategy M is available for using the mediator services:

	C	D	M
C	4,4	0,6	0,6
D	6,0	1,1	1,1
M	6,0	1,1	4,4

Note that the mediated game has a most desirable property: in this game jointly delegating the right of play to the mediator is an equilibrium in group dominant strategies! We can also note that in this example the mediator did not, in

[1] Similar ideas are explored in the extensive literature on commitments and conditional commitments. In particular, (Kalai *et al.*, 2007) shows a folk theorem result for two-player games, using a completely different model of interaction.

fact, require to be informed of the player's chosen strategy – it sufficed for him to know which agents agreed to delegate him their right of play. However, as we will show, in general such mediators will be too weak for implementing GDS; the Prisoner's Dilemma is, in a way, the only example. Therefore, in this paper we will concentrate on fully informed mediators, i.e. ones who can observe the entire action profile selected by the agents and condition their action on it.

Our results: In this paper we explore how different mediators can implement GDS. Section deals with routing mediators. In subsection we present a general sufficient condition for the existence of GDS. We show two natural classes of games that satisfy this condition; both of them are subclasses of ID-congestion games, defined in (Monderer, 2006). We show that simple monotone increasing identity-dependent [MIID] congestion games satisfy our positive criterion, and hence have a solution in GDS using a routing mediator; we show that this also holds for quasi-symmetric MIID-congestion games. Our results also imply that such implementation can be efficiently computed for these classes of games, even when the input representation is succinct.

In subsection we aim to characterize the games which have a solution in GDS using an informed routing mediator. Our goal is a polynomial algorithm that gets a game in strategic form as input, and outputs a routing mediator which implements a solution in GDS, if such exists. We present a polynomial algorithm for this problem for games with 2 and 3 players.

In subsection we present two negative results. The first is a general sufficient condition of *non*-existence of GDS; using this condition we can show that in many known classes of games we can not hope, in general, to attain a solution in GDS with an informed routing mediator (examples include job scheduling, network design, zero-sum games, monotone decreasing congestion games, and more). The second result concerns uninformed mediators – here we justify the claim that the Prisoner's Dilemma game is, in a sense, the only example where a GDS can be achieved with an uninformed mediator.

Section deals with k-implementation. Extending the results of (Monderer and Tennenholtz, 2004), we show that a profile can be implemented as GDS with 0 cost if and only if it is a strong equilibrium. In particular, this result implies that we can implement GDS with 0 cost in all settings where SE is known to always exist, such as job scheduling, network design and certain forms of monotone congestion games (see e.g. (Andelman *et al.*, 2007; Holzman and Law-Yone, 1997)). We also observe that the minimal-cost implementation of a given strategy profile can be computed in polynomial time, given an explicit representation of the game, if we assume that either the number of players or the number of distinct payoffs for each player are constant. Note that an explicit representation of a game takes exponential space in the amount of players, therefore these simplifying assumptions can be justified.

In section we investigate what happens when our mediator has the power of both routing mediators and of k-implementation; i.e. he can both play on behalf of players who give him such right and commit to non-negative payments. There, we derive our main result: the max-min fair outcome of any minimally fair game can be implemented as GDS with 0 cost. Minimally fair games are a generalization of symmetric games: a game is minimally fair if the agents have the same strategy space and, in addition, in every strategy profile agents who chose the same strategy receive the same payoff. This setting applies to many situations where the users are not identical, for example job scheduling (where users may have tasks of different sizes) or certain forms of ID-congestion games.

(Rozenfeld and Tennenholtz, 2007) showed that the max-min fair outcome of any minimally fair game can be implemented as a strong equilibrium with the aid of an informed routing mediator; therefore, our current result can be simply derived from the combination of the result of (Rozenfeld and Tennenholtz, 2007) and our result in section . Nevertheless, we consider it to be the main positive result of the paper, because of its importance: we show that a socially optimal profile of a very large class of games can be implemented as an equilibrium in group dominant strategies with 0 cost.

Routing Mediators

Preliminaries

Recall our intuition on routing mediators: a mediator is a party who can offer agents to play the game on their behalf, and whose behavior on behalf of the agents who agreed to use his services is specified by a contract. This contract can be conditioned on the choices of all other agents. Hence, in this setting, we assume that the original game can be played, in a sense, only through the mediator – for example, the mediator sits on a router that receives all messages about the actions chosen by the players. The mediator cannot alter these messages, but he can observe them; this observability can serve him as a critical tool in establishing his chosen actions on behalf of the players who delegated him their right of play.

First, we formally define routing mediators. We simplify the definitions given in (Rozenfeld and Tennenholtz, 2007), for the following two reasons: first, in this work we consider only pure strategies, and secondly, we restrict ourselves to fully informed mediators (in the notation of (Rozenfeld and Tennenholtz, 2007), we fix $\Omega = \Omega_{full}$).

Let $\Gamma = \langle N, \{A_i\}_{i\in N}, \{u_i\}_{i\in N}\rangle$ be a game in strategic form. A *(fully informed) routing mediator for* Γ is a tuple $\mathcal{M} = \langle m, (\mathbf{c}_z)_{z\in Z}\rangle$, where the following holds:

- $m \notin A_i$ for all $i \in N$. m denotes the new strategy that is now available to each player: to send a message to the mediator, indicating that the player agrees to give the mediator the right of play for him.
- $Z_i = A_i \cup \{m\}$, and $Z = \times_{i\in N} Z_i$. Given $z \in Z$, let $T(z)$ denote $\{j \in N | z_j = m\}$. That is, $T(z)$ denotes the players who agree to give the mediator the right of play for them in z.
- For every $z \in Z$, $\mathbf{c}_z \in A_{T(z)}$. That is, $\mathbf{c}$ is the conditional contract that is offered by the mediator: it specifies exactly which actions the mediator will perform on behalf

of the players who agree to use his services, as a function of the strategy profile chosen by all agents.

Every mediator $\mathcal{M}$ for Γ induces a new game $\Gamma(\mathcal{M})$ in strategic form in which the strategy set of player i is Z_i. The payoff function of i is defined for every $z \in Z$ as follows: $u_i^{\mathcal{M}}(z) = u_i(\mathbf{c}_z, z_{-T(z)})$. For $S \subseteq N$ we denote by m^S the strategy profile $(m, \ldots, m) \in Z_S$. We say that a mediator $\mathcal{M}$ implements a profile a in GDS (resp., SE), if $c_{m^N} = a$ and m^N is a GDS (resp., SE) in $\Gamma(\mathcal{M})$.

Note that when informed mediators are considered, the requirement on the implemented profile to be a GDS (and not some weaker solution concept, such as SE) makes even more sense: the mediator is able, indeed, to observe all the players' actions, so a group of players will want to sign the contract only if the mediator commits to always play in their best interests, according to how the other players play.

Before we proceed with our results, we show an alternative definition of implementing a profile in GDS with the use of a routing mediator; this version is easier to work with and it will serve us in our proofs.

Let $S \subseteq N$, $a_S \in A_S$. We define a game $\Gamma' = (\Gamma \upharpoonright a_S)$ (the *subgame of* Γ *induced by* a_S) as follows: $\Gamma' = \langle N', \{A_i\}_{i \in N'}, \{u'_i\}_{i \in N'} \rangle$ where $N' = N \setminus S$ and $u'_i : A_{N'} \to \Re$ is defined as follows: for any $a_{N \setminus S} \in A_{N'}$, $u'_i(a_{N \setminus S}) = u_i(a_S, a_{N \setminus S})$.

We say that Γ' is a *subgame of* Γ if there exist $S \subseteq N$, $a_S \in A_S$ so that $\Gamma' = (\Gamma \upharpoonright a_S)$. In particular, Γ is a subgame of itself (we call it the *full subgame*).

Let $a, b \in A$ be two strategy profiles. We say that a *strictly dominates* b (or b is *strictly dominated by* a) if $\forall i \in N \ u_i(a) > u_i(b)$. We say that $b \in A$ is *strictly dominated* if there exists $a \in A$ that strictly dominates b.

Note that for any $S \subseteq N, a_S \in A_S, b_{-S} \in A_{-S}$, a_S is not strictly dominated in the subgame $(\Gamma \upharpoonright b_{-S})$ if and only if $a_S \in br(b_{-S})$.

Let Γ be a game and $\mathcal{M} = \langle m, (\mathbf{c}_z)_{z \in Z} \rangle$ a mediator for Γ. We say that $\mathcal{M}$ implements a profile a in GDS if and only if for any $S \subseteq N$, $a_{-S} \in A_{-S}$, m^S is a SE in $(\Gamma(\mathcal{M}) \upharpoonright a_{-S})$. Note the equivalence to the original definition; note also that when checking the requirements for SE it will be enough to check that no profile $b_T \in A_T$ for $T \subseteq S$ strictly dominates m^T in $\big(\Gamma(\mathcal{M}) \upharpoonright (a_{-S}, m^{S \setminus T})\big)$.

Positive Results

Now we begin our exploration of the power of informed routing mediators in establishing GDS. The following theorem presents a sufficient condition for existence of GDS with the aid of a fully informed mediator:

Theorem 1 *Let Γ be a game which satisfies the following property: in any subgame Γ' of Γ there exists a non-empty $S \subseteq N'$ and a profile $a_S \in A_S$, such that for each $i \in S$ and every $b_{-S} \in A_{N' \setminus S}, c \in A_{N'}$ $u'_i(a_S, b_{-S}) \geq u'_i(c)$. In words, the profile a_S guarantees each member of S the highest possible payoff for her in the subgame Γ', no matter what the remaining players in $N' \setminus S$ do. Then:*

1. *There exists a profile $a^* \in A$ that is a SE.*
2. *An informed routing mediator can implement a^* as a GDS.*

Proof:

1. Suppose that the game Γ satisfies the above property. Then, we iteratively define the profile a^* as follows: Γ is in particular a subgame of Γ, therefore there exists a non-empty $S^0 \subseteq N$ and a profile $a_{S^0} \in A_{S^0}$ that satisfies the requirements of the theorem: a_{S^0} guarantees all players in S^0 the highest payoff in Γ. We take $a^*_{S^0} = a_{S^0}$, and consider the subgame $\Gamma' = (\Gamma \upharpoonright a^*_{S^0})$. By the conditions of the theorem, there exists a non-empty $S^1 \subseteq N \setminus S^0$ and a profile $a_{S^1} \in A_{S^1}$ that satisfies the requirements. We take $a^*_{S^1} = a_{S^1}$. We continue in the same manner until the profile a^* is fully defined. Since in every step k the subset S^k is non empty, we need at most $|N|$ steps to define the profile.

   Now we must show that a^* is a strong equilibrium. We show by induction on k that no $i \in S^k$ can be a member of a deviating coalition. It is clear that no member of S^0 will want to deviate, since by playing a^* they guarantee themselves the highest possible payoff in the game. From the definition of a^* we see that the same logic can be used for the induction step: no player in S^{k+1} will want to deviate, since a^* was chosen so that all players in S^{k+1} guarantee themselves the best payoff in the subgame where players in $S^0 \cup \ldots \cup S^k$ play according to a^*.

2. We have to fully define the conditional contract that the mediator offers; in other words, for every $z \in Z$ we have to define the profile $\mathbf{c}_z \in A_{T(z)}$ that the mediator commits to playing on behalf of $T(z)$. We define this profile iteratively, in the similar manner that we defined a^*: we start with the subgame of Γ induced by $z_{-T(z)}$, and fix the action of the set S^0 of players who can guarantee the highest payoff in the subgame; then we fix the action of the set S^1 of players who can guarantee themselves the highest possible payoff in the resulting new subgame; etc. Now the game $\Gamma(\mathcal{M})$ is defined, it remains to verify that in every subgame $\Gamma(\mathcal{M}) \upharpoonright a_{-S}$, playing m^S is a SE. This can be proved in the same manner as (1): by induction on k we can show that no member of S^k will participate in a deviating coalition. We show the induction step: suppose in a profile z players in $S^0 \cup \ldots \cup S^k$ choose m; we must prove that no member of S^{k+1} will want to join a deviating coalition T and play according to some $w_T \in A_T$. Let us denote $z' = (w_T, z_{-T})$ (the profile after the deviation of T). The important thing to notice here is that $(\mathbf{c}_z)_i = (\mathbf{c}_{z'})_i$ for all $i \in S^0 \cup \ldots \cup S^k$ – this follows from our definition of $\mathbf{c}_z$ and the induction hypothesis. Then we can use the same logic as in the proof of (1) to derive the result.

Note the computational implications of the above proof: suppose we have a game which satisfies the conditions of Thm. 1, and we want to implement a solution in GDS efficiently. We can treat the mediator as a kind of oracle: given a profile z, we want to be able to compute $\mathbf{c}_z$ efficiently. It follows from the proof of Thm. 1 that all that we need in order to achieve this goal is the ability to efficiently compute, for any given subgame Γ', the $S \subseteq N'$ and $a_S \in A_S$ whose existence is guaranteed by the theorem. As we will see, in some natural classes of games such computation can be done efficiently, even when the game is given in a succinct representation.

Now we will show two classes of games which satisfy the condition of Thm. 1.

A *monotone increasing identity-dependent [MIID]-congestion game* is defined as follows:

- A finite set of players, $N = \{1, \dots, n\}$.
- A finite non-empty set of facilities, M.
- For each player $i \in N$ a non-empty set $A_i \subseteq 2^M$, which is the set of actions available to player i (an action is a subset of the facilities).
- With every facility $m \in M$ and set of players $S \subseteq N$ a real number $v_m(S)$ is associated, having the following interpretation: $v_m(S)$ is the payoff to each user of m when the set of users of m equals S.
- For each $m \in M, S \subseteq N, T \subseteq S: \ v_m(T) \leq v_m(S)$, meaning that the payoff from a resource is non-decreasing with its the users.

The utility function of player i, $u_i : A \to \Re$, is then defined as follows:

$$u_i(a) = \sum_{m \in a_i} v_m(\{i | m \in a_i\})$$

MIID-congestion games are not congestion games in the original sense of (Rosenthal, 1973), since we allow the payoff from a resource to depend on the identity of its users. It is a particular subclass of ID-congestion games, defined in (Monderer, 2006), with the restrictions to non-player-specific version (users occupying the same resource get the same payoff) and monotone-increasing payoffs. MIID-congestion games can be used to model situations such as buyers clubs, where players choose providers and get discounts based on the group of people they buy with; also they can be used in various situations of non-symmetric sharing of the cost of a resource by the occupying players.

We say that a MIID-congestion game is *simple* if $\forall i \in N, a \in A_i : |a| = 1$.

We say that a MIID-congestion game is *quasi symmetric* if $\forall i, j \in N, A_i = A_j = A$.

Proposition 1 *Let Γ be a MIID-congestion game. Then, Γ satisfies the conditions of Thm. 1 if either one of the following holds:*

1. *Γ is quasi symmetric.*
2. *Γ is simple.*

Proof: Note that if Γ is a simple MIID-congestion game, then any subgame Γ' of Γ is also a simple MIID-congestion game. Same can be said about quasi-symmetric MIID-congestion games. Therefore, to prove both cases it is enough to show the existence of $S \subseteq N, a_S \in A_S$ in the full subgame Γ such that all players is S receive the highest payoff in the game when playing a_S, regardless of other players' actions.

1. Suppose that Γ is quasi symmetric. Let $S = N$, and $a \in \arg\max_{a \in A} \left(\sum_{m \in a} v_m(\{i | m \in a\})\right)$. We define a_S by taking $a_i = a$ for all $i \in S$. It is easy to see that a_S grants each player the highest possible payoff in the game.
2. Suppose that Γ is simple. For any $m \in M$, let $S_{max}(m)$ denote $\{i | \{m\} \in A_i\}$. Let $m \in \arg\max_{m \in M} (v_m(S_{max}(m)))$, and let $S = S_{max}(m)$, and a_S defined as $a_i = \{m\}$ for all $i \in S$. Since the game is monotone increasing, each player in S gets the highest possible payoff in the game.

Note that in both simple and quasi symmetric MIID-congestion games the $S \subseteq N, a_S \in A_S$ of Thm. 1 can be efficiently computed, even if the games are given in a succinct representation; this implies, as we showed, that a solution in GDS can be implemented efficiently.

We have to remark that quasi-symmetric MIID-congestion games would usually be considered trivial – the symmetric socially optimal outcome where each player gets the highest possible payoff in the game is a SE, so where is the problem? The problem is that SE is not GDS. The simplest example of such apparently "trivial" game is the coordination game:

	A	B
A	1,1	0,0
B	0,0	1,1

Obviously, (A, A) and (B, B) are both SE, but what would be a good advice to play? A routing mediator will be able to solve this game by the following contract: if both players cooperate, the mediator plays (A, A), and if one deviates, the mediator will copy her action on behalf of the cooperating player. This solution is more than SE – it is GDS, and it is non-achievable here without a mediator.

Note also that the condition of Thm. 1 is sufficient for being able to implement GDS with an informed mediator, but it is not necessary: the Prisoner's Dilemma game does not satisfy this condition, however the profile (C, C) is implementable as GDS by an informed mediator, as we saw in the introduction. In the next section we will attempt to derive a necessary and sufficient condition for existence of GDS with an informed mediator.

Characterization for $n = 2$ and $n = 3$

Our goal is to characterize all the games in which a GDS is implementable using an informed mediator. We begin with the simple case $n = 2$:

Proposition 2 *Let Γ be a 2 player game, and let a be a strategy profile. a is implementable as GDS using an informed mediator if and only if a is not strictly dominated and*

$$\forall i \in N \ u_i(a) \geq \max_{b_i \in A_i} \{ \min_{b_{3-i} \in br_{3-i}(b_i)} \{u_i(b)\}\} \qquad (*)$$

Proof: $\Rightarrow$ Suppose that a is implementable as GDS using an informed mediator. If a was strictly dominated by b, then (m, m) wouldn't be a SE in the full subgame $\Gamma(\mathcal{M})$, since both players could jointly deviate to b. Therefore, a is not strictly dominated. Next, suppose for contradiction that (*) does not hold. That means, w.l.o.g. there exists $b_1 \in A_1$

such that $u_1(a) < u_1(b_1, b_2)$ for every $b_2 \in br_2(b_1)$. Consider the subgame $\Gamma' = (\Gamma(\mathcal{M}) \upharpoonright b_1)$, and let $b_2 = \mathbf{c}_{(b_1,m)}$ (the response of the mediator on behalf of player 2 in Γ'). If $b_2 \notin br_2(b_1)$, then m_2 is not a SE in Γ'. Therefore, it must hold that $b_2 \in br_2(b_1)$. But in this case, (m, m) is not an SE in $\Gamma(\mathcal{M})$, since player 1 has a profitable deviation to b_1. Contradiction.

$\Leftarrow$ Suppose that both conditions hold. We define an informed mediator $\mathcal{M} = \langle m, (\mathbf{c}_z)_{z \in Z} \rangle$ as follows: if both players cooperate, the mediator plays a; if only player i cooperates, and the player $3-i$ plays b_{3-i}, the mediator plays $b_i \in \arg\min_{b_i \in br_i(b_{3-i})}\{u_{3-i}(b)\}$. Formally, $\mathbf{c}_{(m,m)} = a$, $\mathbf{c}_{(m_i, b_{3-i})} \in$ $\arg\min_{b_i \in br_i(b_{3-i})}\{u_{3-i}(b)\}$. Note that the mediator always plays the best response strategy: $m \in br_i(b_{3-i})$. The two players cannot jointly deviate, since a is not strictly dominated. Finally, (*) implies that $m \in br_i(m_{3-i})$, which completes the proof that (m, m) is a GDS in $\Gamma(\mathcal{M})$.

Proposition 3 *There exists a polynomial algorithm that accepts as input a 3 player game Γ in explicit form and a strategy profile a, and if a is implementable in GDS by an informed mediator, outputs such a mediator.*

Proof: To specify an informed mediator, we must determine which actions the mediator will commit to playing on behalf of the cooperating players, depending on the actions of others; formally, we must specify $\mathbf{c}_z$ for each $z \in (Z \setminus A)$. We know that $\mathbf{c}(m, m, m) = a$. We first determine the sets of "feasible" outcomes that the mediator can choose: consider the profile (m, m, x_3) for some $x_3 \in A_3$. We say that a profile $x = (x_1, x_2, x_3)$ is feasible for a mediator to choose (as $\mathbf{c}(m, m, x_3)$) if the following conditions hold:

1. $u_3(x) \le u_3(a)$.
2. The profile (x_1, x_2) is not strictly dominated in the subgame $(\Gamma \upharpoonright x_3)$.
3. There is no profile (y_1, y_2, x_3) that satisfies condition 1 and for which $\forall i \in \{1, 2\}$ $u_i(y_1, y_2, x_3) \ge u_i(x)$ and for at least one $i \in \{1, 2\}$ the inequality is strict.

We denote by $Pos(m, m, x)$ the set of all feasible profiles for (m, m, x). The set satisfies the following property:
Let $z, z' \in Pos(m, m, x)$. Then $u_1(z) > u_1(z') \Leftrightarrow u_2(z) < u_2(z')$.
More importantly, we can show that if a mediator exists that implements a as GDS, then there also exists such mediator that commits to a profile from $Pos(m, m, x)$ when (m, m, x) is played. In other words, we don't lose generality in restricting our attention to $Pos(m, m, x)$. It is easy to see that the first two conditions are necessary: the first guarantees that player 3 cannot deviate from (m, m, m), and the second guarantees that players 1 and 2 cannot jointly deviate from (m, m, x). Condition 3 states that once condition 1 is satisfied, we can never "lose" by increasing the payoff of the cooperating players – no new deviations are made possible. Therefore, we can safely choose a profile which maximizes the reward of the cooperating players.

Similarly, we define the sets $Pos(m, x, m)$ and $Pos(x, m, m)$ for all $x \in A$. In a similar manner, we want to limit the mediator to feasible actions on behalf of a single cooperating player: we say that a profile $x = (x_1, x_2, x_3)$ is feasible for a mediator to choose for (m, x_2, x_3) if the following conditions hold:

1. $\exists i \in \{2, 3\}$ $u_i(x) \le u_i(a)$.
2. $x_1 \in br_1(x_2, x_3)$.
3. There is no profile (y_1, x_2, x_3) that satisfies conditions 1 and 2 and for which $\forall i \in \{2, 3\}$ $u_i(y_1, x_2, x_3) \le u_i(x)$ and for at least one $i \in \{2, 3\}$ the inequality is strict.

Similarly, we denote by $Pos(m, x, y)$ the set of all feasible profiles for (m, x, y). The set satisfies the following property:
Let $z, z' \in Pos(m, x, y)$. Then $u_2(z) > u_2(z') \Leftrightarrow u_3(z) < u_3(z')$.
Similarly, we can show that if a mediator exists that implements a as GDS, then there also exists such mediator that commits to a profile from $Pos(m, x, y)$ when (m, x, y) is played. As before, it is easy to see that the first two conditions are necessary: the first guarantees that players 2 and 3 cannot jointly deviate from (m, m, m), and the second guarantees that player 1 cannot deviate from (m, x, y). Condition 3 states that once conditions 1 and 2 are satisfied, we can never "lose" by punishing the deviating players – no new deviations are made possible. Therefore, we can safely choose a profile which minimizes the reward of the deviating players. So, in similar manner we define the sets $Pos(x, m, z)$ and $Pos(x, y, m)$ for all $x \in A_1, y \in A_2, z \in A_3$. Note that finding the sets $Pos(\overrightarrow{m})$ takes polynomial time (we assume that the game was given explicitly); obviously, if any one of the sets is empty, the algorithm should stop and return false.

The main idea of the algorithm is: now that all that is left is to find the assignment from each profile $\overrightarrow{m} \in \prod_{i \in N} A_i \cup \{m\} \setminus A \setminus \{(m, m, m)\}$ to $Pos(\overrightarrow{m})$, we are going to reduce the problem to 2-SAT, and solve the resulting 2-SAT instance. Our (binary) variables will be assertions of the form $w_i(\overrightarrow{m}) \ge C$ or $w_i(\overrightarrow{m}) \le C$, meaning that player i gets at least (at most) C in the profile $\overrightarrow{m}$. For any $x_3 \in A_3$, let $\overrightarrow{m} = (m, m, x_3)$. We define

$$\begin{aligned} Var(\overrightarrow{m}) = \quad & \{w_1(\overrightarrow{m}) \le C | \exists x \in Pos(\overrightarrow{m}) :\ u_1(x) = C\} \cup \\ & \{w_1(\overrightarrow{m}) \ge C | \exists x \in Pos(\overrightarrow{m}) :\ u_1(x) = C\} \cup \\ & \{w_2(\overrightarrow{m}) \le C | \exists x \in Pos(\overrightarrow{m}) :\ u_2(x) = C\} \cup \\ & \{w_2(\overrightarrow{m}) \ge C | \exists x \in Pos(\overrightarrow{m}) :\ u_2(x) = C\} \end{aligned}$$

Similarly, for any $x_2 \in A_2$ and $\overrightarrow{m} = (m, x_2, m)$, we define

$$\begin{aligned} Var(\overrightarrow{m}) = \quad & \{w_1(\overrightarrow{m}) \le C | \exists x \in Pos(\overrightarrow{m}) :\ u_1(x) = C\} \cup \\ & \{w_1(\overrightarrow{m}) \ge C | \exists x \in Pos(\overrightarrow{m}) :\ u_1(x) = C\} \cup \\ & \{w_3(\overrightarrow{m}) \le C | \exists x \in Pos(\overrightarrow{m}) :\ u_3(x) = C\} \cup \\ & \{w_3(\overrightarrow{m}) \ge C | \exists x \in Pos(\overrightarrow{m}) :\ u_3(x) = C\} \end{aligned}$$

and for any $x_1 \in A_1$ and $\overrightarrow{m} = (x_1, m, m)$, we define

$$\begin{aligned} Var(\overrightarrow{m}) = \quad & \{w_2(\overrightarrow{m}) \le C | \exists x \in Pos(\overrightarrow{m}) :\ u_2(x) = C\} \cup \\ & \{w_2(\overrightarrow{m}) \ge C | \exists x \in Pos(\overrightarrow{m}) :\ u_2(x) = C\} \cup \\ & \{w_3(\overrightarrow{m}) \le C | \exists x \in Pos(\overrightarrow{m}) :\ u_3(x) = C\} \cup \\ & \{w_3(\overrightarrow{m}) \ge C | \exists x \in Pos(\overrightarrow{m}) :\ u_3(x) = C\} \end{aligned}$$

In the same manner we define $Var(\overrightarrow{m})$ for all other $\overrightarrow{m} \in \prod_{i \in N} A_i \cup \{m\} \setminus A \setminus \{(m, m, m)\}$, and we let

$Var = \bigcup_{\overrightarrow{m}} Var(\overrightarrow{m})$. Now we will translate the problem of finding our mediator into binary clauses on variables in Var.

Transitivity: For all $i \in N$, whenever $C < D$ we add the following clauses:

$$w_i(\overrightarrow{m}) \le C \to w_i(\overrightarrow{m}) \le D$$
$$w_i(\overrightarrow{m}) \ge D \to w_i(\overrightarrow{m}) \ge C$$

Relation between $\ge$ and $\le$: For all $i \in N$, whenever D is a successor of C, we add:

$$\neg w_i(\overrightarrow{m}) \le C \to w_i(\overrightarrow{m}) \ge D$$
$$\neg w_i(\overrightarrow{m}) \ge D \to w_i(\overrightarrow{m}) \le C$$

Anti-symmetry: For all $i \in N$, whenever $C < D$, we add:

$$w_i(\overrightarrow{m}) \le C \to \neg w_i(\overrightarrow{m}) \ge D$$
$$w_i(\overrightarrow{m}) \ge D \to \neg w_i(\overrightarrow{m}) \le C$$

Solution borders: For any maximal C such that $w_i(\overrightarrow{m}) \le C$ is defined, we add:

$$w_i(\overrightarrow{m}) \le C$$

For any minimal D such that $w_i(\overrightarrow{m}) \ge D$ is defined, we add:

$$w_i(\overrightarrow{m}) \ge D$$

Solution domain: Now we use the fact that $Pos(\overrightarrow{m})$ satisfies the aforementioned properties in order to add clauses of the form:

$$w_i(\overrightarrow{m}) \le C \to w_j(\overrightarrow{m}) \ge D$$
$$w_j(\overrightarrow{m}) \ge D \to w_i(\overrightarrow{m}) \le C$$

GDS requirements: Finally, we encode the requirements for GDS: that playing m is a SE in any subgame. We already took care of some of the requirements in our definitions of the sets $Pos(\overrightarrow{m})$; namely, that (m, m, m) is a SE; that there is no beneficial deviation from m for the single cooperating player when the others deviate; and, that two cooperating players don't have a beneficial joint deviation. All that is left to encode is the requirements that one of the two cooperating players must never have a beneficial deviation. They are encoded as follows: let $x \in A_3, y \in A_1$. We add he clause

$$w_1(m, m, x) \le C \to w_1(y, m, x) \le D$$

Here, D is the highest such that $D \le C$, and $w_1(y, m, x) \le D$ is defined. Similarly, we add the clauses:

$$w_2(m, m, x) \le C \to w_2(m, y, x) \le D$$
$$w_1(m, x, m) \le C \to w_1(y, x, m) \le D$$
$$w_3(m, x, m) \le C \to w_3(m, x, y) \le D$$
$$w_2(x, m, m) \le C \to w_2(x, y, m) \le D$$
$$w_3(x, m, m) \le C \to w_3(x, m, y) \le D$$

In each of the above, x, y are chosen from the appropriate domains, and C, D are set according to the definition of the appropriate variable.

In order to finish the reduction, we must show how we define a mediator given a satisfying assignment to the variables: first, we claim that any satisfying assignment to the variables has a certain form: Consider all the variables $w_i(\overrightarrow{m}) \le C_1, \ldots, w_i(\overrightarrow{m}) \le C_k$ and $w_i(\overrightarrow{m}) \ge C_1, \ldots, w_i(\overrightarrow{m}) \ge C_k$ for a given $i, \overrightarrow{m}$ where $C_1, \ldots, C_k$ are sorted in decreasing order. Then, for any satisfying assignment there exists C_l such that $w_i(\overrightarrow{m}) \le C_j$ is satisfied if and only if $j \le l$ and $w_i(\overrightarrow{m}) \ge C_j$ is satisfied if and only if $j \ge l$. This property follows from our requirements on transitivity, anti-symmetry and relations between $\ge$ and $\le$. What this means is that any satisfying assignment determines exactly the payoffs of the two players of interest (for profiles where two player cooperate, these are the cooperating players, and for profiles where two players defect, these are the defecting players); therefore, a satisfying assignment enables us to choose a single profile from $Pos(\overrightarrow{m})$ (the profile is guaranteed to exist, due to the "solution domain" clauses that we included and the way $Pos(\overrightarrow{m})$ was defined).

From all the above we should be convinced that:

1. Any assignment to variables in Var that satisfies all the clauses results in a feasible mediator that implements a in GDS;
2. If a mediator exists that can implement a in GDS, then there exists a satisfying assignment for our clauses.

Of course, the number of clauses of the 2-SAT problem is polynomial in the input size (the game is given explicitly); the time required for the preliminary work, as well as the definition of the mediator, is also polynomial. Therefore the above algorithm provides a complete characterization of achieving GDS in 3-player games by an informed mediator.

For non-constant number of players, the explicit representation of a game is infeasible; so, in a sense, it would not help us much to find an algorithm for general n whose running time is polynomial in the size of the input. We conjecture that even for $n = 4$, the decision problem of whether a given profile can be implemented as GDS by an informed mediator is NP-hard.

Negative Results on Routing Mediators

Now we present a simple necessary condition for existence of GDS, which can be used to derive many negative results. We will need the definition of a symmetric game: A *permutation* of the set of players is a one-to-one function from N onto N. For every permutation π, and for every action profile $a \in A$ we denote by πa the permutation of a by π. That is, $(\pi a)_{\pi i} = a_i$ for every player $i \in S$. Γ is a *symmetric game* if $A_i = A_j$ for all $i, j \in N$, and $u_i(a) = u_{\pi(i)}(\pi a)$ for every player i, for every action profile $a \in A$, and for every permutation π.

Proposition 4 *Let Γ be a symmetric game, and let V denote the optimal social surplus. If there exists a set $S \subseteq N$ and a profile $a_S \in A_S$ such that $\forall a_{-S} \in br_{-S}(a_S), i \in S$ $u_i(a) > V/n$ then there is no profile which can be implemented as a GDS by an informed mediator.*

Proof: Suppose, for contradiction, that there exists a set $S \subseteq N$ and a profile $a_S \in A_S$ such that $\forall a_{-S} \in br_{-S}(a_S), i \in S$ $u_i(a) > V/n$, but there exists a mediator $\mathcal{M} = \langle m, (\mathbf{c}_z)_{z \in Z} \rangle$ that implements a solution in GDS.

Consider the subgame $(\Gamma(\mathcal{M}) \upharpoonright a_S)$, and suppose all the players in $N \setminus S$ cooperate with the mediator. The mediator must choose a profile $\mathbf{c}_{(a_S, m_{-S})} \in br_{-S}(a_S)$, otherwise the players in $N \setminus S$ will have a beneficial deviation. However, that means that all the players in S must receive more than V/n each in the profile (a_S, m_{-S}). Since m^N is a SE in the full subgame $\Gamma(\mathcal{M})$, it follows that each of the players in S must receive more than V/n in the profile m^N as well. However, same reasoning applies to any $S' \subseteq N$ such that $|S'| = |S|$, since the game is symmetric; therefore, every player has to receive more than V/n in m^N – contradiction.

A good example to explain the intuition of the above proposition is the Chicken-Dare game:

	C	D
C	0,0	-1,1
D	1,-1	-10,-10

In this game, a solution implemented by an informed mediator can never be a GDS: any individual player can guarantee herself a payoff of 1 by playing D, because she knows that if she plays D, the mediator will have to play C (the best response) on behalf of her opponent. Since no profile exists where both players receive a payoff of 1, one of them will have an incentive to deviate.

Proposition 4 can be used to derive many negative results on the existence of GDS. For example, in settings such as the network design game, monotone decreasing congestion games, zero-sum games, job scheduling, and more, there is no hope to achieve GDS with an informed mediator.

Our next negative result concerns uninformed mediators. One can remark that the solution in GDS that we presented in the introduction for the Prisoner's Dilemma game did not require very strong assumption that the mediator has the ability to see the actions that players choose in the underlying game. Formally, we say the a mediator $\mathcal{M} = \langle m, (\mathbf{c}_z)_{z \in Z} \rangle$ for a game Γ is uninformed if $\mathbf{c}_z = \mathbf{c}_{z'}$ for all $z, z' \in Z$ such that $T(z) = T(z')$. Such mediators require no access to the players' chosen actions; in fact, any reliable party can serve as an uninformed mediator. It turns out, though, that the class of games where GDS is implementable by such a weak mediator is extremely limited:

Proposition 5 *Let Γ be a game, $n = 2$, and let a be a strategy profile. Then, a can be implemented as GDS by an uninformed mediator if and only if:*

(a) a is not strictly dominated

(b) Γ has an equilibrium b in weakly dominant strategies, for which $u_i(a) \geq u_i(b)$ for $i = 1, 2$

In particular, a solution in GDS is implementable in Γ by an uninformed mediator if and only if Γ possesses an equilibrium in weakly dominant strategies.

Proof: $\Leftarrow$ Suppose a satisfies both conditions. We define the mediator $\mathcal{M} = \langle m, (\mathbf{c}_z)_{z \in Z} \rangle$ as follows: $\mathbf{c}_{(m,m)} = a$, $\forall x \in A \quad \mathbf{c}_{(m,x_2)} = b_1$, $\mathbf{c}_{(x_1,m)} = b_2$. We need to show that (m, m) is GDS in $\Gamma(\mathcal{M})$. Since a is not strictly dominated, the two players do not have a joint beneficial deviation from (m, m). Now w.l.o.g. we need to prove $m \in br_1(x_2)$ for all $x_2 \in Z_2$. For all $x_1 \in A_1$ it holds that $u_1^{\mathcal{M}}(x_1, m) = u_1^{\mathcal{M}}(x_1, b_2) \leq u_1^{\mathcal{M}}(b) \leq u_1^{\mathcal{M}}(a) = u_1^{\mathcal{M}}(m, m)$, therefore $m \in br_1(m)$. Let $x_2 \in A_2$. For all $x_1 \in A_1$ it holds that $u_1^{\mathcal{M}}(x_1, x_2) \leq u_1^{\mathcal{M}}(b_1, x_2) = u_1^{\mathcal{M}}(m, x_2)$, $m \in br_1(x_2)$. So (m, m) is GDS in $\Gamma(\mathcal{M})$.

$\Rightarrow$ Suppose that an uninformed mediator $\mathcal{M} = \langle m, (\mathbf{c}_z)_{z \in Z} \rangle$ implements a in GDS. Obviously, a is not strictly dominated, otherwise both players would jointly deviate from (m, m). Let $b = (\mathbf{c}_{(m,x_2)}, \mathbf{c}_{(x_1,m)})$. The profile b is well defined, because the mediator is uninformed. Since (m, m) is a GDS in $\Gamma(\mathcal{M})$, $m \in br_1(x_2)$ for all $x_2 \in A_2$; therefore, for any $x_2 \in A_2, x_1 \in A_1$ we have $u_1(b_1, x_2) = u_1^{\mathcal{M}}(b_1, x_2) = u_1^{\mathcal{M}}(m, x_2) \geq u_1^{\mathcal{M}}(x_1, x_2) = u_1(x_1, x_2)$, therefore b_1 is a weakly dominant strategy for player 1; similarly, b_2 is a weakly dominant strategy for player 2. Since $m \in br_1(m)$, it must hold that $u_1(a) = u_1^{\mathcal{M}}(a) = u_1^{\mathcal{M}}(m, m) \geq u_1^{\mathcal{M}}(b_1, m) = u_1^{\mathcal{M}}(b) = u_1(b)$; similarly, from $m \in br_2(m)$ we derive that $u_2(a) \geq u_2(b)$, which completes the proof.

Note that if an equilibrium b in weakly dominant strategies exists in Γ, we can always take a profile a to be a Pareto optimal profile s.t. $u_i(a) \geq u_i(b)$ for $i = 1, 2$, and both conditions will hold; thus, we have that a solution in GDS is implementable in Γ by an uninformed mediator if and only if Γ possesses an equilibrium in weakly dominant strategies.

When $n \geq 3$ we do not have any hope, in general, to implement GDS with an uninformed mediator – even if an equilibrium in weakly dominant strategies exists. The reason for this is that in order to achieve a GDS a mediator has to play a best response on behalf of all the cooperating players, no matter what the others chose; when a group of players cooperates, their joint best response is not necessarily playing their individual weakly dominant strategies; so, without the information about the played profile, the mediator may not know the best response. So, in a sense, the Prisoner's Dilemma is the only example of achieving GDS with an uninformed mediator.

K-Implementation

In this section we turn to a different kind of mediators, introduced by (Monderer and Tennenholtz, 2004). We assume that the mediator is an interested party who has the power to alter the game by committing to non-negative monetary transfers to the players, conditioned on the outcome of the game. Formally, given a game $\Gamma = \langle N, A, U \rangle$, such a mediator is defined by a payoff function vector $V = \{v_i\}_{i \in N}$, where each $v_i : A \rightarrow \Re$ is non-negative. Given a mediator a game Γ and a mediator V, the mediated game $\Gamma(V)$ is simply $\langle N, A, U + V \rangle$.

Note that the above definition implicitly makes two important assumptions:

- Output observability: The interested party can observe the actions chosen by the players.
- Commitment power: The interested party is reliable in the sense that the players believe that he will indeed pay the additional payoff defined by V.

Note also that unlike routing mediators discussed in the previous section, here the mediator does not play the game on behalf of the agents. Similarly to routing mediators, though, he observes players actions and offers a reliable contract conditioned on these actions; he also does not restrict the players' actions in any way, and does not enforce behavior.

Given a game Γ and a profile $a \in A$, we say that *a has a k-implementation in weakly dominant strategies* if there exists a V such that:

1. a is an equilibrium in weakly dominant strategies in $\Gamma(V)$
2. $\sum_{i \in N} v_i(a) \leq k$

Similarly, we define a *k-implementation in group dominant strategies*. It is easy to see that a k-implementation of any profile always exists; in particular, if we denote the maximal difference of payoffs in the game matrix by D, it is easy to see that an $D \cdot n$ implementation of any profile always exists. Obviously, our goal is to find *cheap* implementations; in particular, we are interested in 0-implementation.

Theorem 2 *(Monderer and Tennenholtz, 2004) Let Γ be a game and a a strategy profile. Then, a has a 0-implementation in weakly dominant strategies if and only if a is a NE.*

The above result can be extended into the following:

Theorem 3 *Let Γ be a game and a a strategy profile. Then, a has a 0-implementation in GDS if and only if a is a SE.*

Proof: $\Leftarrow$ Suppose a is a SE in Γ. Consider the following mediator V: $v_i(a) = 0$ for all $i \in N$; for $b \neq a$, $v_i(b) = D$ if $a_i = b_i$, and 0 otherwise. Obviously, $\sum_{i \in N} v_i(a) = 0$. We claim that a is a GDS in $\Gamma(V)$. First, we note that a is still a SE in $\Gamma(V)$: this holds because we set $v_i(b) = 0$ whenever $a_i \neq b_i$, so profitable deviations were not added. Let b be a strategy profile, and $S \subseteq N$. If $a_{-S} \neq b_{-S}$, then obviously $a_S \in br_S(b_{-S})$, since all players in S get the highest payoff in the game when (a_S, b_{-S}) is played. But when $a_{-S} = b_{-S}$, we have that $a_S \in br_S(a_{-S})$ because a is a SE in $\Gamma(V)$. Therefore, in both cases $a_S \in br_S(b_{-S})$, which makes a a GDS.

$\Rightarrow$ Suppose V is a 0-implementation of a in GDS, and suppose for contradiction that b_S is a profitable deviation of a coalition $S \subseteq N$ in Γ. We denote $b = (b_S, a_{-S})$. Since $\sum_{i \in N} v_i(a) = 0$, in particular we have $v_i(a) = 0$ for all $i \in S$; therefore, $\forall i \in S \quad u_i(a) + v_i(a) = u_i(a) < u_i(b) \leq u_i(b) + v_i(b)$, which means that b_S strictly dominates a_S in the subgame $\Gamma(V) \upharpoonright a_{-S}$; therefore $a_S \notin br_S(a_{-S})$, so a is not GDS in $\Gamma(V)$. Contradiction.

This result implies that we can implement GDS with 0 cost in all settings where SE is known to always exist, e.g.: job scheduling, network design (Andelman *et al.*, 2007) and certain forms of monotone congestion games (Holzman and Law-Yone, 1997; Rozenfeld and Tennenholtz, 2006).

Now we turn to the computational question of finding the optimal k-implementation. (Monderer and Tennenholtz, 2004) showed a polynomial algorithm for finding the optimal k-implementation in dominant strategies; now we would like to extend their results to implementation in GDS.

Proposition 6 *Let $\Gamma = \langle N, \{A_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ be a game, with $|N| = n$, $|A_i| \leq m$, and $|\{u_i(a)|a \in A\}| \leq p$ for all $i \in N$. Then, an exhaustive brute-force algorithm for finding an optimal k-implementation in GDS of a given profile $a \in A$ runs in $O(p^n \cdot n \cdot m^n)$.*

Proof: To implement a profile a in GDS, we have to find a profile of payments $v_i(a)$ to the agents such that for each subset S of agents, and each joint deviation b_S, there exists an agent in $i \in S$ such that $u_i(a) + v_i(a) \geq u_i(a_{-S}, b_S)$ (note that in all profiles $b = (a_S, b_{-S})$ where $\forall i \notin S \; b_i \neq a_i$ we w.l.o.g. can set $v_j(b) = D$ for any $j \in S$, and set $v_i(b) = 0$ for any $i \notin S$). Note also that given a k-implementation V, we can verify the validity of the implementation as follows: simply go over all the possible deviations b_S, and ensure that each deviation is covered (one of the players does not benefit). This takes at most $O(n \cdot m^n)$ steps. Note also that the amount of all possible mediators to check is at most p^n, because w.l.o.g. we can consider only possibilities in which a player receives a total payoff that equals one of his possible payoffs in Γ. Therefore, the simple brute force algorithm that checks all possible mediators takes at most $O(p^n \cdot n \cdot m^n)$ steps.

In general, p is bounded by m^n; note that if the game is given explicitly, p is at most polynomial in the size of the input, and n is at most logarithmic in the size of the input. Therefore, in the case where either n or p are constant, the brute-force algorithm that checks all the possibilities is polynomial in the size of the input[2].

Combining Routing Mediators with K-Implementation

In this section we consider mediators who combine the power of routing mediators and k-implementation. Our goal is to implement a good solution in GDS in an interesting class of games. First, we formally define combined mediators:

Let Γ be a game in strategic form. A *combined mediator for* Γ is a tuple $(\mathcal{M}, V)$, where $\mathcal{M}$ is a routing mediator for Γ and V is a payoff function vector for $\Gamma(\mathcal{M})$ (as defined in the previous section).

We say that a combined mediator $(\mathcal{M}, V)$, where $\mathcal{M} = \langle m, (\mathbf{c}_z)_{z \in Z} \rangle$, implements a profile a in GDS with cost k, if:

- $c_{m^N} = a$
- V is a k-implementation in GDS of m^N in the game $\Gamma(\mathcal{M})$

Let Γ be a game in strategic form. Γ is a *minimally fair game* (Rozenfeld and Tennenholtz, 2007) if for all $i, j \in N, X_i = X_j$ and for every action profile $x \in X$, $x_i = x_j$ implies that $u_i(x) = u_j(x)$. That is, a game is minimally

[2]It is an interesting question to consider the computational complexity of finding the optimal k-implementation for non-constant n and p, when the game is given explicitly. It is very unlikely that the problem is NP-hard, since, as we saw, the size of the witness is $O(t^{\log t})$ (where t represents the input size). In fact, several complexity classes have been defined that are good candidates for this problem (Papadimitriou and Yannakakis, 1996); we conjecture that finding the optimal k-implementation in GDS is LOGSNP-complete.

fair if players who play the same strategy get the same payoff. The exact value of the received payoff may depend on the identities of the players who chose the strategy, as well as on the rest of the profile. In particular, every symmetric game is a minimally fair game; however, minimally fair games capture a much wider class of settings. For example, typical job-shop scheduling games are minimally fair games.

In order to define what solution is considered "good", we employ the standard model of max-min fairness (Kleinberg *et al.*, 1999; Kumar and Kleinberg, 2000). We call an allocation of strategies to players max-min fair if the utility of any player cannot be increased without decreasing the utility of a player who was facing an already lower utility. In many settings max-min fairness is a natural social optimality criterion.

Now we are ready to state our main result:

Theorem 4 *Let Γ be a minimally fair game and let a be a max-min fair profile of the game. Then, a can be implemented in GDS by combined routing mediator with 0 cost.*

Proof: Applying Thm. 1 of (Rozenfeld and Tennenholtz, 2007), there exists a routing mediator $\mathcal{M} = \langle m, (\mathbf{c}_z)_{z\in Z}\rangle$ that implements a as a SE; that is, $c_{m^N} = a$ and m^N is a SE of the game $\Gamma(\mathcal{M})$. Applying Thm. 3 to the game $\Gamma(\mathcal{M})$, we have that there exists a 0-implementation V of m^N in GDS. Therefore, $(\mathcal{M}, V)$ is the desired combined mediator.

Further Work

We see several possible directions for further work on the subject of group dominant strategies. First of all, this paper did not consider the possibility of using mixed strategies. Ideally, in the extension of the definition of GDS to the mixed case, we would like to consider the possibility of correlated deviations; we would like to use the concept of correlated strong equilibrium (defined in (Rozenfeld and Tennenholtz, 2006)) and extend it to hold without any rationality assumptions on the non-cooperating players (in a manner similar to the GDS concept). This, of course, will result in a much stronger solution concept than GDS; in particular, Thm. 4 will no longer hold. However, we note that Thm. 1 still holds in this case – which implies, in particular, that the positive results of Prop. 1 regarding simple and quasi-symmetric MIIID-congestion games still apply under this concept.

Another direction is to consider the tradeoff between the cost of an implementation and the quality of the implemented profile. In section , we implicitly assumed that the cheaper implementation is always better, in particular, that our goal is 0-implementation. However, it seems natural to us that the decision making of the interested party is based, in general, on both the cost of an implementation and the quality of the implemented profile – where the quality can be, for example, the social surplus. It will be interesting to consider both the computational issues of finding the optimal implementation (where optimal now combines these two criteria) and the following variant of price of stability: to compare the optimal social surplus with the cost of the optimal implementation in GDS (where this cost, again, somehow combines the two criteria).

References

N. Andelman, M. Feldman, and Y. Mansour. Strong price of anarchy. In *SODA*, 2007.

R.J. Aumann. Acceptable points in general cooperative n-person games. In A.W. Tucker and R.D. Luce, editors, *Contribution to the Theory of Games, Vol. IV, Annals of Mathematics Studies, 40*, pages 287–324. 1959.

R.J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.

R. Holzman and N. Law-Yone. Strong equilibrium in congestion games. *Games and Economic Behavior*, 21:85–101, 1997.

Matthew O. Jackson. A crash course in implementation theory. *Social Choice and Welfare*, 18(4):655–708, 2001.

Adam Tauman Kalai, Ehud Kalai, Ehud Lehrer, and Dov Samet. Voluntary commitments lead to efficiency. Discussion Papers 1444, Northwestern University, Center for Mathematical Studies in Economics and Management Science, April 2007. available at http://ideas.repec.org/p/nwu/cmsems/1444.html.

Jon M. Kleinberg, Yuval Rabani, and Eva Tardos. Fairness in routing and load balancing. In *IEEE Symposium on Foundations of Computer Science*, pages 568–578, 1999.

A. Kumar and J. Kleinberg. Fairness measures for resource allocation. In *Proc. 41th IEEE Symp. on Foundations of Computer Science*, 2000.

A. Mas-Colell, M.D. Whinston, and J.R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

Dov Monderer and Moshe Tennenholtz. K-Implementation. *Journal of Artificial Intelligence Research*, 21:37–62, 2004.

D. Monderer and M. Tennenholtz. Strong Mediated Equilibrium. In *Proceedings of AAAI-06*, 2006.

D. Monderer. Solution-based congestion games. 8:397–409, 2006.

R. B. Myerson. Multistage games with communication. *Econometrica*, 54(2):323–358, 1986.

Christos H. Papadimitriou and Mihalis Yannakakis. On limited nondeterminism and the complexity of the v-c dimension. *Journal of Computer and System Sciences*, 53(2):161–170, 1996.

R.W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

Ola Rozenfeld and Moshe Tennenholtz. Strong and Correlated Strong Equilibria in Monotone Congestion Games. In *Proc. of 2nd international Workshop on Internet & Network Economics*, 2006.

Ola Rozenfeld and Moshe Tennenholtz. Routing Mediators. In *Proc. 20th International Joint Conference on Artificial Intelligence*, 2007.