Routing Mediators

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

Abstract

We introduce a general study of routing mediators. A routing mediator can act in a given multi-agent encounter on behalf of the agents that give it the right of play. Routing mediators differ from one another according to the information they may have. Our study concentrates on the use of routing mediators in order to reach correlated strong equilibrium, a multi-agent behavior which is stable against deviations by coalitions. We study the relationships between the power of different routing mediators in establishing correlated strong equilibrium. Surprisingly, our main result shows a natural class of routing mediators that allow to implement fair and efficient outcomes as a correlated super-strong equilibrium in a very wide class of games.

1 Introduction

In many multi-agent systems it makes little sense to assume that agents will stick to their parts when provided with suggested strategies, if a deviation can increase their payoffs. While for any multi-agent encounter, represented as a strategic form game, there always exists a strategy profile for which unilateral deviations are not beneficial, more demanding requirements are rarely satisfied. In order to tackle this issue we consider in this paper the use of *routing mediators*, extending upon previous work in game theory and AI.

A mediator is a reliable entity that can interact with the players and perform on their behalf actions in a given game. However, a mediator can not enforce behavior. Indeed, an agent is free to participate in the game without the help of the mediator. This notion is highly natural; in many systems there is some form of reliable party or administrator that can be used as a mediator. Notice that we assume that the multi-agent interaction (formalized as a game) is given, and all the mediator can do is to perform actions on behalf of the agents that explicitly allow it to do so. The mediator's behavior on behalf of the agents that give it the right of play is pre-specified, and is conditioned on the information available to the mediator on all agents' actions.

This natural setting is different from the one discussed in the theory of mechanism design (see [Jackson, 2001] for an introduction) where a designer designs a new game from scratch in order to yield some desired behavior. The simplest form of mediator discussed in the game theory literature is captured by the notion of correlated equilibrium [Aumann, 1974]. This notion was generalized to communication equilibrium by [Forges, 1986; Myerson, 1986]. Another type of mediators is discussed in [Monderer and Tennenholtz, 2004]. However, in all these settings the mediator can not perform actions on behalf of the agents that allow it to do so.

A powerful type of mediators, which can perform actions on behalf of the agents who give the mediator the right of play is considered in [Monderer and Tennenholtz, 2006]. This type of mediators turns out to be a specific case of the mediators discussed in this paper. Namely, in this restricted setting the agents need to decide whether to give the mediator the right of play, while the mediator does not possess any further information on the agents' actions.

In order to illustrate the power of a reliable mediator consider the famous prisoners dilemma game:

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

In this classical example we get that in the only equilibrium both agents will defect, yielding both of them a payoff of 1. However, this equilibrium, which is also a dominant strategy equilibrium, is inefficient; indeed, if both agents deviate from defection to cooperation then both of them will improve their payoffs. Formally, mutual defection is not a strong equilibrium; it is not stable against deviations by coalitions. Indeed, there is no strong equilibrium in the prisoners dilemma game.

Consider now a reliable mediator who offers the agents the following protocol: if both agents agree to use the mediator services then he will perform cooperate on behalf of both agents. However, if only one agent agrees to use his services then he will perform defect on behalf of that agent. Notice that when accepting the mediator's offer the agent is committed to actual behavior as determined by the above protocol. However, there is no way to enforce the agents to accept the suggested protocol, and each agent is free to cooperate or defect without using the mediator's services. Hence, the mediator's protocol generates a new game, which we call a *mediated game*:

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

The mediated game has a most desirable property: in this game there is a strong equilibrium; that is, equilibrium which is stable against deviations by coalitions. In this equilibrium both agents will use the mediator services, which will lead them to a payoff of 4 each!

Although the above example illustrates the potential power of a mediator who can act on behalf of agents who give him the right of play, the above mediator was very restricted. In general, a *routing mediator* possesses information about the actions taken by agents. Consider a router in a typical communication network. Messages submitted to the system must pass through this router. In addition, the router can suggest a protocol to the agents who may wish to use its services also in selecting appropriate routes. Hence, in this situation it is most natural to assume that the router can observe selected actions of all agents, and not only of those who give it the right of play. In order to illustrate the power of such routing mediator consider the following example:

	A	B	C
A	4,4	6,0	1,1
B	0,6	4,4	0,6
C	1,1	6,0	1,1

Given the above game, assume we are interested in obtaining a socially optimal outcome, in which each agent gets a payoff of 6. It is easy to see that the column player can guarantee himself an expected payoff of 6.5 by playing B or C with probability 0.5 each. As a result, a routing mediator who is not informed about the pure realizations of the column player's strategies cannot implement any outcome in which that player gets a payoff less than 6.5. However, a routing mediator who can see the pure realizations of the deviating player and choose the punishment appropriately can enforce the outcome $(6, 6)$ as a strong equilibrium! This can be achieved, for example, by the following simple protocol: If only agent i selects the mediator's services, the mediator will copy the action of the other player, j, on behalf of i; if both agents select to use the mediator's services then the mediator will perform C on behalf of both agents.

This example illustrates the important role of information in the definition of a routing mediator. Indeed, routing mediators can be classified based on the information they can access about the agents' actions. In this paper, we formally define routing mediators as a function of their available information. We distinguish between two types of available information:

1. Information about the actual action instantiations.
2. Information about the agents' programs (i.e. correlated strategies).

As we will show, having full access to the agents' actions may sometimes be more powerful than having full access to their programs, and vice versa.

Given the general setting of routing mediators, we prove that they can indeed significantly increase the set of desired outcomes that can be obtained by multi-agent behaviors which are robust against deviations by coalitions. Namely, we show that any *minimally fair* game possesses a fair and efficient correlated super-strong mediated equilibrium. In a correlated super-strong equilibrium, correlated deviations by any subset of agents can not benefit one of them without hurting another. A minimally fair game is a game in which all actions are available to all agents, and all agents who select identical actions get identical payoffs. A particular example of minimally fair games are the symmetric games, but the family of minimally fair games is much wider. In particular, our results are applicable to symmetric congestion games and job-shop scheduling, two major topics of study in the interface between computer science and game theory.

2 Preliminaries

2.1 Games in strategic form

A game in strategic form is a tuple $\Gamma = \langle N, (X_i)_{i\in N}, (u_i)_{i\in N}\rangle$, where N is a finite set of players, X_i is the strategy set of player i, and $u_i : X \rightarrow \Re$ is the payoff function for player i, where $X = \times_{i\in N} X_i$. If $|N| = n$, whenever convenient we assume $N = \{1, \cdots, n\}$. Γ is *finite* if the strategy sets are finite.

For every $S \subseteq N$ we denote $X_S = \times_{i\in S} X_i$. When the set, N is clear, $X_{N\setminus S}$ will be also denoted by X_{-S}, and moreover, $X_{-\{i\}}$ will be also denoted by X_{-i}.

When we discuss subsets of a given set we implicitly assume non-emptiness unless we specify otherwise. Let Y be a set. The set of probability distributions with finite support over Y is denoted by $\Delta(Y)$. That is, every $c \in \Delta(Y)$ is a function $c : Y \rightarrow [0, 1]$ such that $\{y \in Y | c(y) > 0\}$ is finite, and $\sum_{y\in Y} c(y) = 1$. For every $y \in Y$ we denote by δ_y the probability distribution that assigns probability 1 to y. Let I be a finite set, and let $c_i \in \Delta(Y_i)$, $i \in I$, where I is a finite set of indexes, and Y_i is a set for every $i \in I$. We denote by $\times_{i\in I} c_i$ the product probability distribution on $Y = \times_{i\in I} Y_i$ that assigns to every $y \in Y$ the probability $\prod_{i\in I} c_i(y_i)$.

Let $\Gamma = \langle N, (X_i)_{i\in N}, (u_i)_{i\in N}\rangle$ be a game in strategic form, and let $S \subseteq N$. Every $c \in \Delta(X_S)$ is called a *correlated strategy for* S . A correlated strategy for the set of all players N is also called a *correlated strategy*, and for every i, a correlated strategy for $\{i\}$ is also called a *mixed strategy* for i. Let $c \in \Delta(X)$ and let $S \subseteq N$, the marginal probability induced by c on X_S is denoted by $c_{[S]}$. That is,

$$c_{[s]}(x_S) = \sum_{x_{-S} \in X_{-S}} c(x_S, x_{-S})$$

The expected payoff of i with respect to a correlated strategy c is denoted by $U_i(c)$. That is,

$$U_i(c) = \sum_{x \in X} u_i(x) c(x)$$

For every i, and for every $x_i \in X_i$, the mixed strategy δ_{x_i} is called a *pure strategy*.

For every S, the set of mixed-strategy profiles is denoted by Q_S. That is, $Q_S = \times_{i \in S} \Delta(X_i)$. We will use Q for Q_N, and Q_i for $Q_{\{i\}}$.

The *mixed extension* of the game Γ is the game $(N, (Q_i)_{i \in N}, (w_i)_{i \in N})$, where for every $q \in Q$, $w_i(q) = U_i(q_1 \times \cdots \times q_n)$.

2.2 Strong equilibrium concepts

Let $\Gamma = \langle N, (X_i)_{i \in N}, (u_i)_{i \in N} \rangle$ be a game in strategic form, and let $x \in X$. We say that x is a *strong equilibrium of type I* [Aumann, 1959] in Γ if the following holds:

For every subset, S of players and for every $y_S \in X_S$ there exists $i \in S$ such that $u_i(y_S, x_{-s}) \leq u_i(x)$.

Let $q = (q_1, \cdots, q_n)$ be a profile of mixed strategies.

We say that q is a *strong equilibrium of type II in* Γ if q is a strong equilibrium of type I in the mixed extension of Γ. That is, q is a strong equilibrium of type II in Γ if for every subset of players S, and for every profile of mixed strategies $(p_i)_{i \in S}$ there exists $i \in S$ such that $U_i(\times_{i \in S} p_i \times_{i \in N \setminus S} q_i) \leq U_i(\times_{i \in N} q_i)$.

Obviously, if $x \in X$, and $(\delta_{x_1}, \cdots, \delta_{x_n})$ is a strong equilibrium of type II in Γ then x is a strong equilibrium of type I, but the converse does not hold.

Let $c \in \Delta(X)$. We say that c is a *strong equilibrium of type III*, or *correlated-strong equilibrium*, if for every subset of players S , and for every correlated strategy for S, $\xi_S \in \Delta(X_S)$ there exists $i \in S$ such that $U_i(\xi_S \times c_{[-S]}) \leq U_i(c)$.

Let $q \in Q$. Obviously q is a strong equilibrium of type II if $q_1 \times q_2 \times \cdots \times q_n$ is a strong equilibrium of type III, but not vice versa.

Let $x \in X$. The requirement that x is a strong equilibrium of type II seems to be acceptable in an environment in which the players believe that they and others could not possibly correlate their behavior (e.g., when every player is sitting in a separate room, and there is no communication between the players). However, every player can perform a private randomization. In an environment in which the players do not correlate their strategies in X, but they may fear/hope that such a correlation is possible, we expect x to be a strong equilibrium of type III in order to be believed/played by the players. In the scope of this article, we will consider such correlation possible, and therefore when we use the term strong equilibrium, we refer to correlated strong equilibrium.

The above definitions assumed that a player would agree to join a deviating coalition only if the deviation would result in a strict improvement of his payoff. A much stronger stability concept is needed if we fear that players may be persuaded to join a deviating coalition also in cases where such deviation would leave their payoff unchanged, but would increase the payoff of at least one of the other coalition members.

We say that $c \in \Delta(X)$ is a *correlated super-strong equilibrium* (or *super-strong equilibrium*) if there is no $S \subseteq N$ and $\xi_S \in \Delta(X_S)$, such that:

1. $\forall i \in S \quad U_i\left(\xi_S, c_{[-S]}\right) \geq U_i\left(c\right)$
2. $\exists i \in S \quad U_i\left(\xi_S, c_{[-S]}\right) > U_i\left(c\right)$

A super-strong equilibrium is, in particular, a strong equilibrium, but not vice versa.

2.3 Information sets

Let Y be a set. We say that Ω is an *information set on* Y if Ω is a partition of Y.

Let $y \in Y$. By $\omega(y)$ we refer to the equivalence class of y in Ω, i.e. to $P \in \Omega$ s.t. $y \in P$.

We use these definitions in order to model situations in which an element $y \in Y$ is chosen, and an agent has only partial information about the chosen element. It is assumed that, if Ω is the information set that is available to such agent, the agent is informed of $\omega(y)$.

Let Ω_1, Ω_2 be information sets on Y. We say that Ω_1 *refines* Ω_2 if for all $y_1, y_2 \in Y$ if $\omega_1(y_1) = \omega_1(y_2)$ then $\omega_2(y_1) = \omega_2(y_2)$. That is, an agent with available information Ω_1 is informed at least as well as an agent with Ω_2.

The above relation is a partial ordering of the information sets, which forms a lattice: the maximal possible information on Y is $\Omega_{full} = \{\{y\} | y \in Y\}$; having such information means that for any choice of $y \in Y$, , the agent knows exactly which element was chosen. The minimal possible information on Y is $\Omega_\emptyset = \{Y\}$; that means that the agent has no information at all about any chosen element $y \in Y$.

3 Routing mediators

We now introduce routing mediators, a general tool for coordinating and influencing agents' behavior in games. A routing mediator is always assumed to be reliable, and he is endowed with a single ability: to play for the players who explicitly give him the right to play for them. The mediator cannot enforce the players to use his services, neither can he affect the players' utility values. However, routing mediators differ in the information they have on the strategy profile which is being chosen by the players.

Let Γ be a game in strategic form. A *routing mediator for* Γ is a tuple $\langle m, \Omega, \Psi, (\mathbf{c}_{\omega,\psi})_{\omega \in \Omega, \psi \in \Psi} \rangle$, where the following holds:

- $m \notin X_i$ 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.[1]
- Ω is an information set on Z, where $Z = \times_{i \in N} Z_i$, and $Z_i = X_i \cup \{m\}$. Ω represents the information available to the mediator about each pure realization of the players' correlated strategy.

[1] [Monderer and Tennenholtz, 2006] show that no power is added if the communication between the mediator and the players includes more than one message type.

Ω refines Ω_{basic}, which is defined as follows:
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. Then, Ω_{basic} is the partition induced by the following equivalence relation on Z: $z_1 \equiv z_2$ iff $T(z_1) = T(z_2)$ (For $\omega \in \Omega$ we will use $T(\omega)$ to denote $T(z)$ for some $z \in \omega$).
That is, we demand that in each pure realization of the agents' strategy the mediator knows exactly which players give him the right of play for them.

- Ψ is an information set on $\Delta(Z)$, which represents the information available to the mediator about the correlated strategy of the players. We place no restrictions on Ψ, in particular Ψ may equal $\Psi_{basic} = \{\Delta(Z)\}$, meaning that the mediator has no such information.
- for every $\omega \in \Omega, \psi \in \Psi$, $\mathbf{c}_{\omega,\psi} \in \Delta\left(X_{T(\omega)}\right)$. 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 information available to it about the strategy profile chosen by the 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 $c \in \Delta(Z)$ as follows:

$$U_i^{\mathcal{M}}(c) = \sum_{z \in Z} c(z) U_i(\mathbf{c}_{\omega(z),\psi(c)} \times (\times_{j \in N \setminus T(z)} \delta_{z_j}))$$

The above definitions express the following interaction between the mediator and the players:

1. The players choose a profile $c \in \Delta(Z)$
2. A joint coin flip is issued, determining a pure realization $z \in Z$ of c
3. The mediator is informed of $\omega(z)$ and $\psi(c)$
4. The mediator issues a coin flip to determine a pure realization $\delta_{T(z)}$ of $\mathbf{c}_{\omega(z),\psi(c)}$
5. The payoff of each player i is given according to $u_i\left(\delta_{T(z)}, z_{-T(z)}\right)$

The payoffs of the players from the profile c are assumed to be their expected payoffs from both coin flips in steps 2,4.

The mediator is guaranteed to behave in a pre-specified way, however this behavior is conditional on the actions of the other agents. For this reason, although $U_i(c)$ is well defined for any profile $c \in \Delta(Z)$, in particular for pure profiles $z \in Z$, $U_i(c)$ does not necessarily equal $\sum_{z \in Z} U_i(z) c(z)$, unless $\Psi = \Psi_{basic}$. Nevertheless, the concepts of strong equilibrium can still be applied to $\Gamma(\mathcal{M})$.

Let Γ be a game in strategic form, $m \notin \bigcup_{i \in N} X_i$, Ω an information set on Z, Ψ an information set on $\Delta(Z)$. A correlated strategy $c \in \Delta(X)$ is a *strong mediated equilibrium relatively to* Ω, Ψ if there exists a mediator for Γ, $\mathcal{M} = \langle m, \Omega, \Psi, (\mathbf{c}_{\omega,\psi})_{\omega \in \Omega, \psi \in \Psi} \rangle$, with $\mathbf{c}_{\omega(m^N),\psi(m^N)} = c$, for which m^N is a correlated strong equilibrium in $\Gamma(\mathcal{M})$. The notion of *super-strong mediated equilibrium* can be defined similarly.

Note that the above definition coincides with the one of [Monderer and Tennenholtz, 2006] in the case where $\Omega = \Omega_{basic}$, $\Psi = \Psi_{basic}$; i.e. they considered the case in which the mediator is only informed of the players that chose to delegate him the right to play for them in step 2; he has no information about the actions of the other players, neither about the correlated strategy that was chosen in step 1.

We think that there may be quite a few situations in real life where it would be relevant to consider $\Omega \neq \Omega_{basic}$, or $\Psi \neq \Psi_{basic}$. For example, a situation which can be modeled by $\Psi = \Psi_{full} = \{\{c\} | c \in \Delta(Z)\}$ is when the mediator doesn't decide/reveal its action before the game takes place, and players agree to deviate only if the mediator can't *possibly* harm anyone of them (i.e. they consider the *worst case* of the unknown action). Another such situation might be when the deviating players fear that an information leak might be possible, i.e. the mediator will find out about their plan (discover the strategy ξ_S) and respond accordingly. Lastly, it could be that the correlated deviation is implemented using a program which is submitted and executed on a server, and the mediator is located on this server; in this case, the mediator indeed is fully informed about the correlated strategy of the players.

$\Omega \neq \Omega_{basic}$ can be used to model all the situations in which the mediator does indeed, as the name implies, "stand between" the players and the actual game, e.g. situations in which the play is only available through the mediator. For example, if the players' strategies are messages in a network and the mediator sits on the router, it might be reasonable to assume that $\Omega = \Omega_{full}$.

Lastly, we can imagine many situations in which the mediator is only partially informed, both in terms of Ω and of Ψ.

Observations: Suppose $c \in \Delta(X)$ is a strong (super-strong) mediated equilibrium relatively to Ω, Ψ.

1. Let $w(c) = (w_1, \ldots, w_n)$ be the payoff vector induced by c. Then, for any $c' \in \Delta(X)$ such that $w(c') = w(c)$, c' is a strong (resp. super-strong) mediated equilibrium relatively to Ω, Ψ. In light of this, we will sometimes use the following notation: a payoff vector w can be implemented as a strong (super-strong) mediated equilibrium relatively to Ω, Ψ if there exists $c \in \Delta(X)$ such that $w(c) = w$ and c is a strong (resp. super-strong) mediated equilibrium relatively to Ω, Ψ.
2. If $\Psi = \Psi_{full}$, c can be implemented as a strong (resp. super-strong) mediated equilibrium relatively to Ω, Ψ even if the mediator is restricted to use only pure actions (i.e. satisfies $\mathbf{c}_{\omega,\psi} \in X_{T(\omega)}$).
3. If Ω' refines Ω, then c is also a strong mediated equilibrium relatively to Ω', Ψ
4. If Ψ' refines Ψ, then c is also a strong mediated equilibrium relatively to Ω, Ψ'

Observations 3, 4 show that the possible values of Ω, Ψ induce a four-end lattice of strong mediated equilibrium concepts, where the strongest (most stable) concept is obtained from $\Omega_{basic}, \Psi_{basic}$ (since it assumes nothing on the information available to the mediator), and the weakest concept is obtained from $\Omega_{full}, \Psi_{full}$. Part of this lattice is depicted below (arrows mean increase in stability):

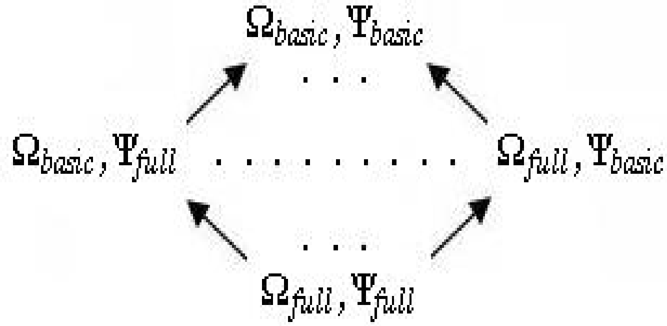

A natural question to ask is whether there exists a general relationship between the strong mediated equilibria concepts induced by $\Omega_{basic}, \Psi_{full}$ and $\Omega_{full}, \Psi_{basic}$. The question is, what empowers the mediator more – increasing the information available to him on each pure outcome, or increasing the information about the correlated strategies?

As it turns out, this question does not have a general answer:

Proposition 1 *There exists a game Γ and a strategy profile c which is a strong mediated equilibrium relatively to $\Omega_{full}, \Psi_{basic}$, but not a strong mediated equilibrium relatively to $\Omega_{basic}, \Psi_{full}$.*

An example of such game and such profile appears in the introduction. A less intuitive result is the following:

Proposition 2 *There exists a game Γ and a strategy profile c which is a strong mediated equilibrium relatively to $\Omega_{basic}, \Psi_{full}$, but not a strong mediated equilibrium relatively to $\Omega_{full}, \Psi_{basic}$.*

Proof: (sketch) Let Γ be the following 3-person game:

Player 3: A

	A	B	C
A	12,0,0	0,0,0	0,0,0
B	0,0,0	0,12,0	0,0,0
C	0,0,0	0,0,0	5,5,5

Player 3: B

	A	B	C
A	24,-12,0	0,0,0	0,0,0
B	0,0,0	0,12,0	0,0,0
C	0,0,0	0,0,0	0,0,0

Let $c = (C, C, A)$. In order to see that c is not a strong mediated equilibrium in Γ relatively to $\Omega_{full}, \Psi_{basic}$, consider the possible actions of the mediator on behalf of player 3 given the profile $(A, A, m) \in Z$. The mediator has to choose a strategy on behalf of player 3 without knowing the correlated strategy of players $\{1, 2\}$ that caused the realization (A, A). No matter what distribution between A and B will the mediator choose, the total payoff of the deviators will be 12, with player 1 getting more than player 2; note that the pure profile (B, B) guarantees the deviators the payoffs $(0, 12)$; therefore, no matter what strategy the mediator will choose, players $\{1, 2\}$ will be able to randomize between (A, A) and (B, B) so that each of them will get a payoff of 6, making the deviation beneficial for them.

However, note that the players $1, 2$ were able to guarantee themselves the payoffs $(6, 6)$ as a *response* to any contract offered by the mediator. In a sense, the information set Ψ_{basic} means that the players are more informed than the mediator – the strategy of the mediator in this setting is fixed, and cannot depend on the correlated strategy chosen by the players; while their correlated strategy can be a best response to the mediator's strategy. When the information set Ψ_{full} is considered, the mediator becomes more informed than the players: he can devise a response based on their correlated strategy. In this example, a mediator who is aware of the exact probabilities that the players $1, 2$ assign to the profiles (A, A) and (B, B) can always respond by a pure strategy A or B on behalf of the third player in a way that will cause one of the players to get a payoff less than 5, making the deviation not beneficial for them.

A formal proof of the above intuition is technically cumbersome, and therefore is left to the full version. ∎

4 The existence of strong mediated equilibrium

The following result in due to [Monderer and Tennenholtz, 2006]:

Proposition 3 *Every 2-person game has a strong mediated equilibrium relatively to $\Omega_{basic}, \Psi_{basic}$.*

That is, all two-person games can be solved using a mediator in the strongest possible sense – even when the mediator is completely uninformed.

Before proceeding with other existence results, we need to define the notion of symmetry.

A *permutation* of the set of players is a one-to-one function from N to N. For every permutation π and for every action profile $x \in X$ we denote by πx the permutation of x by π. That is, $(\pi x)_{\pi i} = x_i$ for every player $i \in N$.

Let Γ be a game in strategic form. Γ is a *symmetric game* if $X_i = X_j$ for all $i, j \in N$ and $u_i(x) = u_{\pi(i)}(\pi x)$ for every player i, for every action profile $x \in X$ and for every permutation π.

Needless to say that symmetric games are most popular in computerized settings. For example, the extremely rich literature on congestion games in computer science deals with particular form of symmetric games. However, our main result requires a weaker notion of symmetry:

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

That is, a game is minimally 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.

Note that in a minimally fair game, unlike in a symmetric game, the symmetric socially optimal payoff is not necessarily feasible. In order to capture our notion of fair and efficient outcomes in such games, we need to define the notion of lexicographic ordering on vectors of payoffs:

Let $sort : \Re^n \to \Re^n$ be a function that sorts the elements of real vectors in increasing order. Given two vectors $w, w' \in \Re^n$ we say that $w \prec_l w'$ (w is lexicographically smaller than w') if there exists $1 \le i \le n$ such that $sort(w)_i < sort(w')_i$ and for all $1 \le j < i \quad sort(w)_j = sort(w')_j$.

Let Γ be a minimally fair game. We say that a vector $w \in \Re^n$ is the *fair efficient payoff* of Γ if w is the lexicographically maximal feasible payoff vector in Γ. It is easy to see that such vector always exists, and is unique. Also, such vector is, by definition, Pareto optimal.

The fairness criterion we use is in fact the well-known min-max fairness criterion, following the work of Rawls [Rawls, 1971] on fairness as justice.

We are now ready to state our main result:

Theorem 1 *Let Γ be a minimally fair game. Then, the fair efficient payoff of Γ can be implemented as a super-strong mediated equilibrium relatively to $\Omega_{full}, \Psi_{basic}$.*

Proof: Let w be the fair efficient payoff vector of Γ. Suppose w.l.o.g. that the vector w is sorted in increasing order and the players in N are renamed appropriately.

Let $m \notin \bigcup_{i \in N} X_i$. We define a routing mediator $\mathcal{M} = \langle m, \Omega_{full}, \Psi_{basic}, (\mathbf{c}_{\omega,\psi})_{\omega \in \Omega_{full}, \psi \in \Psi_{basic}} \rangle$ as follows: Ψ_{basic} contains a single element, therefore it suffices to define $\mathbf{c}_{\omega(z)}$ for all $z \in Z$. Let $z \in Z$. If $T(z) = N$, we define $\mathbf{c}_{\omega(z)}$ to be the strategy profile that implements w. Otherwise, let $T = T(z)$, let $i = \max \{N \setminus T\}$ (i.e. i is the highest indexed player among those who don't cooperate with the mediator in z), and let $a = z_i$. We then define $\mathbf{c}_{\omega(z)} = \times_{i \in T} \delta_a$.

That is, the action of the mediator for each pure realization of the correlated profile is to have all the cooperating players copy the action of the deviator with the highest payoff in w.

Let $S \subseteq N, \xi_S \in \Delta(Z_S)$. Let $\xi = \xi_S \times \left(\times_{i \in N \setminus S} \delta_m\right)$ be the profile in $\Gamma(\mathcal{M})$ that is obtained from the deviation of S to ξ_S. Suppose, for contradiction, that S weakly benefits from deviating to ξ_S in $\Gamma(\mathcal{M})$ relatively to m^N; that is, $\forall i \in S \quad U_i(\xi) \ge w_i$ and $\exists i \in S \quad U_i(\xi) > w_i$.

Let k be the highest indexed player in S. In every pure realization z of ξ, all the players in $N \setminus S$ play the same action as k: either because $k \in T(z)$, or because k is, in particular, the highest indexed player in $N \setminus T(z)$. Therefore, $U_j(\xi) = U_k(\xi)$ for every player $j \in N \setminus S$; therefore, for every $j \in \{1, ..., k\} \setminus S$, $U_j(\xi) \ge w_k \ge w_j$, which means that all the players in $S' = \{1, ..., k\}$ can deviate to the correlated profile ξ', in which the players in S play by ξ and the players in $S' \setminus S$ copy the strategy of player k, and S' would also weakly benefit from this deviation.

Therefore, we can assume w.l.o.g. that $S = \{1, \ldots, k\}$ for some $1 \le k < n$. Let w' be the payoff vector that is induced by ξ. We know that for every $k \le j \le n$, $w'_j = w'_k \ge w_k$. Since w is Pareto optimal, we know that $k < n$ (otherwise, w' would Pareto dominate w); w.l.o.g. we can assume that $w_k < w_{k+1}$.

We define a payoff vector w'' as follows:

$$w'' = \begin{cases} w' \text{ if } \max_{i \in S} \{w'_i\} \le w'_k \\ \varepsilon w' + (1 - \varepsilon) w \text{ otherwise} \end{cases}$$

where $\varepsilon \in \arg\min_{i: w'_i > w'_k} \left\{ \frac{w_{k+1} - w_i}{w'_i - w_i + w_{k+1} - w'_{k+1}} \right\}$.

First, we note that $(w''_1, ..., w''_k)$ Pareto dominates $(w_1, ..., w_k)$: if $w'' = w'$ this follows from the fact that S weakly benefits from the deviation that results in w'; in the other case, the same holds since $\varepsilon > 0$.

Next, we note that the choice of ε ensures that $w''_i \le w''_j$ for $j \le k$, $i > k$; combined with our previous observation, this implies that $w \prec_l w''$, contradicting our choice of w. ∎

Below we present two additional results, that can be proven using similar logic as behind the proof of Theorem 1. The proofs of these results are therefore omitted due to lack of space. Both results attempt to relax the requirements on the information set Ω, at the price of restricting attention to the class of symmetric games (Theorem 2) or considering strong instead of super-strong equilibrium (Theorem 3).

The relaxation of the requirements on Ω is based on the idea that in order to "punish" a set of deviating agents, we can do without seeing the actions of all of them; it suffices, under some conditions, to observe only the action selected by a single deviator.

Formally, Ω must satisfy the following: there exists an ordering $<_P$ on N such that for all $z, z' \in Z$ if $z_{\max_{<_P}\{N \setminus T(z)\}} \neq z'_{\max_{<_P}\{N \setminus T(z)\}}$ then $\omega(z) \neq \omega(z')$. That is, Ω allows to distinguish between different actions of the highest ranked player (by $<_P$) among the deviators.

Theorem 2 *Let Γ be a symmetric game, and suppose than an information set Ω on Z satisfies the above property.*

Then, the symmetric optimal payoff vector $w = (V, \ldots, V)$ in Γ can be implemented as a super-strong mediated equilibrium relatively to Ω, Ψ_{basic}.

Theorem 3 *Let Γ be a minimally fair game, and suppose than an information set Ω on Z satisfies the above property.*

Then, the fair efficient payoff of Γ can be implemented as a strong mediated equilibrium relatively to Ω, Ψ_{basic}.

5 Further work

The general aim of this research is to find rich settings in which routing mediators can help to achieve stability and/or efficiency. In this paper we concentrated on achieving stability against deviations by coalitions via the concept of correlated strong equilibrium; other desirable solution concepts can be of interest, e.g. pure equilibria concepts (Nash, strong) and weakly dominant strategies.

Other models of multi-agent interaction, e.g. repeated games, or games with incomplete information, are also under consideration.

References

[Aumann, 1959] 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.

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

[Forges, 1986] F. M. Forges. An approach to communication equilibria. *Econometrica*, 54(6):1375–85, 1986.

[Jackson, 2001] M. Jackson. A Crash Course in Implementation Theory. *Social Choice and Welfare*, 18(4):655–708, 2001.

[Monderer and Tennenholtz, 2004] D. Monderer and M. Tennenholtz. K-Implementation. *Journal of Artificial Intelligence Research (JAIR)*, 21:37–62, 2004.

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

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

[Rawls, 1971] J. Rawls. *A Theory of Justice*. Belknap Press, 1971.