A REVELATION PRINCIPLE FOR COMPETING MECHANISM GAMES

MICHAEL PETERS

ABSTRACT. This paper shows how to characterize the set of outcomes functions that can be supported as equilibrium outcome functions in competing mechanism games. We describe a set of mechanisms we refer to as *reciprocal mechanisms.* These play the same role as direct mechanisms in single mechanism designer problems in that they provide a 'canonical', though abstract, way of representing equilibrium outcomes. We use them to show that the set of outcome functions supportable as equilibrium in competing mechanism games is equivalent to the set of outcome functions supportable by a centralized mechanism designer provided the set of contracts is large enough.

1. Introduction

It has been known for some time that the 'revelation principle' doesn't hold in competing mechanism games. What this means is that modeling competing mechanism designers as if they offered the usual kind of direct mechanisms in which agent reports about their payoff types are converted into actions, will only capture some of the things that can be supported as equilibrium in competing mechanism games. This argument stems from a remark in (McAfee 1993) - since agents observe the mechanisms that were offered by the other mechanism designers, their types should be defined broadly enough to allow them to convey this payoff information. McAfee didn't make anything of this, but the subsequent literature offers many simple examples to illustrate why this actually makes a difference. Examples of equilibrium outcome functions in competing mechanism games that can't be supported as equilibria when designers offer naive direct mechanisms have been given by (Peck 1995), (Epstein and Peters 1999), (Martimort and Stole 2002), and (Peters 2001) among others.

This problem makes it difficult to understand what can be supported with competing mechanisms since there are potentially so many different ways to model a competing mechanism game and the contracts that are available to players in that game. The point of this paper is to show that competing mechanism games, like standard incentive problems, have a canonical structure that can be understood in very much the same way that standard incentive problems can be understood using the revelation principle. To this end, this paper provides a set of mechanisms for competing players can use. We refer to these as *reciprocal mechanisms.* These play the same role as direct mechanisms. They provide a way to characterize all the outcome functions that are supportable as equilibrium of competing mechanism games. They also make it possible to characterize these outcomes with a simple and very familiar set of inequalities.

June 28, 2011.

Reciprocal mechanisms differ from direct mechanisms in two ways. At the most basic level, they augment players ability to communicate by providing them with a message we refer to as a *correlating message*. This message is used to support randomization and correlation among the actions of the players. This basic difference between competing mechanisms and standard games was pointed out by (Peck 1995) who provided examples of competing mechanism games in which players were able to correlate their play by communicating through agents.

However, a more important difference is that reciprocal mechanisms recognize that some of the messages that players use in a competing mechanism game are *commitment messages*. As the name suggests these describe commitments. One way that players can use these commitment messages is by asking agents to describe the commitment messages of other players. Here we use a simpler method in which players condition directly on these commitment messages.

The paper consists of two theorems. After describing competing mechanism games in a somewhat compact way, our main theorem shows that any outcome function that can be supported as a Bayesian equilibrium associated with *some* competing mechanism game[1], can also be supported as a Bayesian equilibrium in a very simple game in which players compete in reciprocal mechanisms. This simple game involves two rounds of communication. In the first round, players simultaneously and public announce commitments. In the second round, players simultaneously and privately send type reports of the usual kind, and correlating messages from the interval $[0,1]$ to every other player. In this equilibrium, all commitment messages are identical, every player reports his type truthfully to every other player, and each player chooses his correlating message using a uniform randomization over the interval $[0,1]$.

Our second theorem shows that the set of supportable outcomes is the same as the set of outcome functions that can be supported by a centralized mechanism designer who has the ability to control the actions of the players.

The paper proceeds with a heuristic description of the method. The main theorems are then presented. This is followed by a discussion of the literature and implications of the result.

2. Program Equilibrium

The basic approach used in this paper closely resembles (but it considerably simpler than) a technique from computer science called 'program equilibrium'. This idea is based on an idea suggested by [2]. A program has a kind of duality - it is a set of instructions that converts some set of inputs into an output - but at the same time it is data that can serve as input into other programs. Programs can then be written to identify themselves, and this functionality can be used to support nice outcomes when the original games are replaced by games in which players submit programs to play for them. The logic of this method is explained in (Peters and Szentes 2008) who encode programs as integers, then illustrate how to construct self-referential programs using this logic.

We emphasize that though we can represent the equilibria of games played by programs by using reciprocal contracts, reciprocal contracts themselves involve no self-referential statements or complex logic. Instead, we simply restrict the number

[1]The game has to satisfy an inscrutability requirement that we describe below.
[2]Von Neumann

of commitments, then match every commitment with a message that expresses it. The objective is to find the smallest set of messages that can be used to support all equilibrium. In this sense, our objective is the opposite of (Epstein and Peters 1999) or (Peters and Szentes 2008) who are looking for very rich sets of commitment devices. The method more closely resembles the method in (Pavan and Calzolari 2008) who a player's *extended type*. The player's extended type includes a description of the outcome the player expects to prevail in the game. We use this approach here.

To see the rough idea, consider a simple symmetric prisoner's dilemma game with actions C and D, and restrict attention for the moment to pure outcomes. Begin by defining a message θ^*, which either player can announce publicly at the first stage of our competing mechanism game. Announcing this message commits the player to the following action which depend on the public message of the other player:

$$\theta^* \equiv \begin{cases} C \text{ if player 2 announces } \theta^* \\ D \text{ otherwise.} \end{cases}$$

Add to this two additional messages, θ^c. Announcing θ^c at the beginning of the game commits the player to use the cooperative action no matter what public message the other player uses, and θ^d which commits to the non-cooperative action no matter what message the other player announces. We then simply define the set of feasible reciprocal mechanisms as $\left\{\theta^*, \theta^c, \theta^d\right\}$. Each of the symbols θ^*, θ^c and θ^d plays a dual role as it represents a specific mechanism, but is also used as a possible message which is used by the mechanisms or contract of the other player.

The normal form of the game in reciprocal mechanisms is simply the following (replacing payoffs with outcomes):

	θ^*	θ^c	θ^d
θ^*	(CC)	(DC)	(DD)
θ^c	(CD)	(CC)	(CD)
θ^d	(DD)	(DC)	(DD)

It is immediate that in this new normal form in reciprocal mechanisms, there is an equilibrium in which both players publicly announce the mechanism θ^*. In this equilibrium cooperation occurs. Evidently (DD) can also be supported as an equilibrium.

We can now use the argument in (Pavan and Calzolari 2008, Pavan and Calzolari 2009) to extend this argument to arbitrary (finite) complete information games. Suppose there are n players, each of whom has a finite set of feasible actions A_i. Define $A = \prod_{i=1,\ldots n} A_i$ and $A_{-i} = \prod_{j\neq i} A_j$. Again, we will stick for the moment to pure strategies, and define a collection of constant mechanisms for each player. To make things a little simpler, suppose each player has k pure actions.[3] Then create k messages $\theta^i_{a_1}$ to $\theta^i_{a_k}$ for player i. Announcing any of these messages publicly commits i to play the corresponding action no matter what messages other players announce.

Now we can create a set of reciprocal mechanisms which will be denoted by messages like $\theta^i_{\left(a_1,a'_i\right),\ldots,\left(a_n,a'_n\right)}$. The objects a_i are elements of the feasible set of

[3]This is without loss of generality since we can modify payoff functions to assign low payoffs when i takes actions that aren't feasible in the original game.

actions for player i. One way to think of these is that they are the actions that players expect everyone to use in the game. The object a'_i is a vector of $n-1$ actions from i's action set. Informally, these represent punishments that i will implement against each of the other players when they don't do what is expected. So the index (a_i, a'_i) has n components. There are k^{n^2} reciprocal mechanisms like this. The set of mechanisms available to each player i will then be the set of constant mechanisms plus this set of reciprocal mechanisms.

The public message $\theta_{(a_1,a'_1),\dots(a_i,a'_i),\dots,(a_n,a'_n)}$ commits player i as follows:

$$\theta_{(a_1,a'_1),\dots(a_i,a'_i),\dots,(a_n,a'_n)}\left(\theta'_{-i}\right) =$$

$$\begin{cases} a'_{ij} & \exists ! j : \theta'_j \neq \theta_{(a_1,a'_{-1}),\dots,(a_i,a'_{-i}),\dots,(a_n,a'_n)} \\ a_i & \text{otherwise.} \end{cases} \tag{2.1}$$

In the notation above $\exists!$ means "there exists a unique".

Each mechanism like this has a cooperative action a_i and an array of punishments $a'_i = \{a'_{i1}, \dots, a'_{im}\}$ for each of the other players. If all the other players offer the same reciprocal mechanism as player i, then i will respond with his cooperative action a_i. If all but one of the players offer the same reciprocal mechanism as player i, then i will respond with the punishment action a'_{ij} for that player. It is immediate that this set of reciprocal mechanisms supports every action profile in which each player receives at least his min max payoff as a Nash Equilibrium in mechanisms.

This illustrates how reciprocal mechanisms work. We now proceed to show how this method can be used to characterize all supportable allocation rules.

3. Incomplete Information Games

In a game of incomplete information, there are n players. Each player has a finite action set A_i and a finite type set T. In standard notation A, A_{-i} represent cross product spaces representing all players actions and the actions of all the players other than i, respectively. Types are jointly distributed on T^m according to some common prior and preferences of player i are given by $u_i : A \times T \to \mathbb{R}$. Players have expected utility preferences over lotteries. Note that because the distribution of types is arbitrary, it is without loss of generality to assume that all players types lie in the same space.

Let Q_i, Q and Q_{-i} represent mixtures over A_i, A and A_{-i} respectively. Since players have expected utility preferences we write $u(Q,t)$ to be the expected utility associated with the mixture Q when types are i. An *outcome function* is a mapping $\omega : T^n \to Q$. An outcome function is *incentive compatible if*

$$\mathbb{E}\left\{u\left(\omega\left(t\right),t\right)|t_i\right\} \geq \mathbb{E}\left\{u\left(\omega\left(t'_i, t_{-i}\right),t\right)|t_i\right\}. \tag{3.1}$$

A punishment $\rho_i : T_{-i} \to Q_{-i}$ is used when player i chooses not to participate in the mechanism that implements ω. The outcome function is *individually rational* if there is a punishment ρ_i for each player i such that for every i

$$\mathbb{E}\left\{u\left(\omega\left(t\right),t\right)|t_i\right\} \geq \max_{a_i \in A_i} \mathbb{E}\left\{u\left(a_i, \rho_i\left(t_{-i}\right)\right), t|t_i\right\}. \tag{3.2}$$

3.1. **Competing Mechanism Games.** Competing mechanism games come in every shape and form. This inevitably leads to an indigestible description of what a competing mechanism game is. The same is true here. The key point in this section is simply to illustrate the breadth of the revelation principle proved below. As in much of the previous literature, a competing mechanism game consists of a set of possible messages for each player, along with a commitment for every player that converts messages into actions. The best known variants of this come from the competing auction literature (for example, (Epstein and Peters 1999),(Yamashita 2010) or (Peters and Troncoso-Valverde 2009)) or the literature on common agency ((Pavan and Calzolari 2001) or (Martimort and Stole 2002) or (Bernheim and Whinston 1986)) in which mechanism designers simultaneously offer mechanisms which make commitments based on a specific group of players called agents. The theorem below covers this framework. However it also works when mechanisms are offered sequentially, as in (Pavan and Calzolari 2009) or privately as in (Segal and Whinston 2003).

No matter the details of the timing or whether or not contracting is private, players announce announce commitments which influence the actions of other players. In the references above, influence is indirect. For example, in the competing auction approach, an agent observes a message by some principal that represents a commitment, then conveys that message to another principal, who responds to it. The literature on program equilibrium (e.g., (Tennenholtz 2004)) and contractible contracts ((Peters and Szentes 2008)) assume that principals can make commitments that condition directly on the other principals' commitment messages. At a general level, papers like (Yamashita 2010) or (Peters and Troncoso-Valverde 2009) show that there is no difference between these two approaches - direct contracting acts as a kind of reduced form for the more familiar approach of committing through agents. So we want a formulation that encompasses all these different approaches, yet at the same time retains the idea that mechanisms are mappings from messages to actions.

To capture all this we assume the basic game involves a finite number of rounds. At each round, each player can transmit at most one signal from some feasible set. This feasible set can vary from round to round it needn't be the same for each player in any particular round. Sequential play is captured in this approach by by giving players degenerate message sets in stages where they cannot communicate. A message for a player is a sequence of signals. The set of feasible messages for player i is given by $\mathcal{M}_i$. The space $\mathcal{M}_i$ is partitioned into a collection E of measurable subsets (equivalence classes) $\{\mathcal{M}_i^e\}_{e\in E}$. Each element of the same equivalence class corresponds to the same commitment by player i. In this formalism, messages are already associated with commitments. So we don't need a formalism for the set of feasible mechanisms. Instead, we have for each player a single mapping$\lambda_i : \mathcal{M}_i \times \prod \mathcal{M}_j \to A_i$. If two messages m and m' are in the same equivalence class for player i, then $\lambda_i(m, \cdot) = \lambda_i(m', \cdot)$.

The picture that follows illustrates with the prisoner's dilemma example discussed above in which players can commit unconditionally to C or D (contracts θ_c or θ_d), or to cooperate if the other player offers θ^*. To make a simple point, we add an additional contract in which the player can defer choosing an action until the second stage.

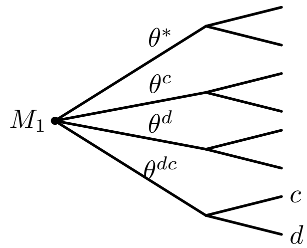

This picture depicts a situation in which player 1 can send one of four signals in the first round of the game, θ^*, θ_c, θ_d and θ_{cd} . These messages are all associated with different commitments for player 1. The first three are the same commitments we described in the simple example at the beginning of the paper. The final signal is interpreted as a commitment to defer a choice of action until the second round of the game. In the second round, the player has two signals c and d. The first two messages $\{\theta^*, c\}$ and $\{\theta^*, d\}$ represented by the two topmost branches in this tree are both in the same equivalence class since they both represent a commitment to use action c if the other player announces the commitment θ^* in the first round, and to use action d otherwise. Similarly along the branches associated with $\theta_c$and θ_d, the second round message are irrelevant because the first round messages commit the player to use action c or d respectively, independent of any other messages sent during the game. The final pair of branches $\{\theta_{cd}, c\}$ and $\{\theta_{cd}, d\}$ are different equivalence classes representing commitments chosen in the second stage of the game.

The mappings $\{\lambda_i\}$ embed all the modeler's choices concerning how the competing mechanism game works. The choices all involve measurability restrictions. For example, a principle agent model (or its extension in common agency) assumes that the principal can condition on the messages of the agent, but not conversely. A detailed understanding of the mappings λ_i isn't critical in what follows. However, it is important to note that this framework incorporates not just asymmetries, and sequential play. It will also accommodate private contracting. This just to emphasize that the theorem here covers a much broader set of models than just common agency, or competing auctions.

A competing mechanism game is defined using the collection $\{\mathcal{M}_i, \lambda_i\}_{i=1,\ldots,n}$. For each array of messages $\{m_1, \ldots, m_n\}$, type contingent payoffs are

$$\left\{u_i\left(\{\lambda_j(m_1, \ldots, m_n)\}_{j=1,\ldots n}, t\right)\right\}_i.$$

This defines a sequential Bayesian game. We focus on Bayesian equilibrium in what follows so that the theorem can be applied under any refinement. Fix an array $\{\sigma_1, \ldots, \sigma_n\}$ of type contingent strategies for each of the players.[4] The payoff to player i is

$$\mathbb{E}\left\{u_i\left(\{\lambda_j(\tilde{m}_1, \ldots, \tilde{m}_n)\}_{j=1,\ldots n}, (t_i, t_{-i})\right)|\sigma, t_i\right\}.$$

In the theorem below, we want to characterize implementable outcomes with reference to a 'default' game which might be played in the absence of contracts. This default game we take as a primitive. It could be simultaneous or sequential. Whatever the case we deal with its normal form. Some modeling assumptions

[4]These strategies are mappings from information sets into mixtures over signals at different rounds. However, these details do not affect the following argument.

about contracts will effectively change this normal form game. To imagine a simple example, suppose the default game is matching pennies with minmax value 0. If we add on top of this game an asymmetric contract structure, we can change it into a sequential game. For example suppose that player two can only announce a commitment to play either heads or tails, while player 1 can announce any commitment that is conditional on player 2's first round message. Evidently player 2 will commit to match (or mismatch depending on the payoffs) player 1's commitment. Then the value for 1 is 1 and the value for player 2 is -1 as it would be in a sequential version of the game.

To deal with this, we need a modeling convention. We adopt the approach in (Peters and Szentes 2008).

Definition 1. An array of strategies σ is *regular* for player i if there is a type dependent joint distribution $Q_{-i}(t_{-i})$ such that for each $a_i \in A_i$, i can send a message m_{a_i} that commits i to action a_i with probability 1 given σ_{-i}, and such that the joint distribution of actions of the other players conditional on m_{a_i}, σ_{-i} and t_{-i} is $Q_{-i}(t_{-i})$.

In an obvious terminology, we say an array of strategies σ is regular if it is regular for each player i. Furthermore, a competing mechanism game $\{\mathcal{M}_i, \lambda_i\}$ is regular if every array of strategies is regular.

What this restriction does is to ensure that player i has the option of 'opting out' of the contracting process completely. Here, opting out means not only that the player can simply choose an action unconditionally, but also that he can prevent others making their actions conditional on the action he chooses. The asymmetric matching pennies game described above is an example of a game that isn't regular. The prisoner's dilemma game described in the figures above is regular since the messages $\{\theta_{Cd}, c\}$ or $\{\theta_{cd}, d\}$induce commitments c and d respectively while eliciting the same response from the opponent.

A couple of remarks are in order. The formulation here simply assumes there is a set of players, all of whom could potentially offer contracts. The special models that have been examined in the literature are special cases. For example the principle agent formulation ((Epstein and Peters 1999), (Yamashita 2010), (Han 2006)) assumes that players can be divided into two groups consisting of uninformed principals who possess the ability to write contracts, and informed agents who can't write contracts at all. Formally, we simply endow 'agents' with the ability to condition only on their own messages. Common agency (Peters 2001) or (Martimort and Stole 1998) imposes the additional restriction that there is only a single agent, and prevents the 'principals' from contracting with each other.

It isn't at all obvious what the right way to model indirect mechanisms is in the competing mechanism game. Menu mechanisms in common agency, or competition in reserve prices among auctioneers ((Peters and Severinov 1997)) simply impose restrictions on what mechanism designers can do. Complex issues associated with contractual robustness, and the infinite regress associated with allowing contracts to depend on other contracts are addressed in (Epstein and Peters 1999) and (Peters and Szentes 2008). Here we are simply going to assume that the modeler has chosen methods to resolve these issues and included them in the specification of the set of messages.

Given this specification, the equilibrium in the competing mechanism game is defined in the usual way. A strategy rule specifies for each player and for each of his possible types, a randomization over the set of mechanisms and rules that specify randomizations over the set of messages sent to other players conditional on whatever information is available at the time the message is sent. We focus on Bayesian equilibrium (as does the revelation principle) because we are interested in characterizing all possible outcomes, and all the refined equilibrium are also Bayesian equilibrium.

3.2. **Reciprocal Mechanisms.** The objective here is to show that there is a relatively simple competing mechanism game that can be used to understand equilibria in all competing mechanism games. This game takes place in two stages. At the first stage each player simultaneous announces a public message that describes his commitments in the game. In the second round each player privately sends a message in $[T \times [0,1]]$ to each of the other players. The first element of this message is a type report, the second is a correlating message used to support randomization. In the formalism of the previous section the message space consists of a sequence of two signals, the first from Θ_i are all public, the second from $(T \times [0,1])^n$ are all private.

The public messages in Θ_i are all tied to commitments which may be of two possible types. We describe the messages then associate a mechanism with each of them. First, we focus on a subset of the set of measurable mappings $c_i : (T \times [0,1])^n \times (T \times [0,1])^{n-1} \to A_i$. The arguments of these functions are messages the player sends to and receives from other players. Notice that there are n messages that the player sends privately to other players instead of $n-1$. The informed player can commit based on his type, but defer revealing his type information until the second round. This is just inscrutability as in (Myerson 1983). Also note that we allow the principal to make commitments based on all his messages. Thus he can verify to any player the set of messages that he sent to any other player. It is nonetheless important that these messages are sent privately so that the player has the option of hiding these messages from some players if he wants to.

A mechanism d_i is referred to as a *direct mechanism* if it is a measurable function from $(T \times [0,1])^n \times (T \times [0,1])^{n-1} \to A_i$, if d_i implements some arbitrary fixed action unless the first n messages are all the same, and if it depends only on the fractional part of the sum of the n correlating messages. We explain how these correlating messages are used below. A *punishment mechanism* p_i^j for player i to use against player j is a direct mechanism involving all the players except player j. Let D_i be the set of direct mechanisms available to i, and P_{ij} the set of punishment mechanisms available to player i when he wants to punish player j..

The second kind of commitment we are interested in is referred to as a *reciprocal mechanism*. To construct these, we begin with the public messages that represent them. Let $\delta_i = \left\{ d_i, \left\{ p_i^j \right\}_{j \neq i} \right\}$ consist of a direct mechanism and a list of punishment mechanisms to be used against each of the other players. The notation $\delta = \{\delta_1, \ldots, \delta_m\}$ represents a list of such lists, while Δ is the set of all δ. We are going to index the set of feasible mechanism for player i by $\Theta_i = \Delta \cup D_i$. If $\theta_i \in D_i$, then θ_i commits player i to the corresponding direct mechanism.

We need to associate a unique commitment with each element of Δ. Use the notation θ_i^δ to be the mechanism associated with the message δ, and let d_i^δ to refer the direct mechanism associated with the i^{th} element of the list δ. The *reciprocal* *mechanism* θ_i^δ is given by the mapping

$$\theta_i^\delta(\theta') = \begin{cases} p_i^j & \exists! j : \theta'_j \neq \theta_j^\delta \\ d_i^\delta & \text{otherwise.} \end{cases} \tag{3.3}$$

The notation ∃! means "there exists a unique". This commits player i to the direct mechanism d_i^δ (from the list δ) unless there is a unique player j who fails to use the mechanism θ_j^δ. In that case, the mechanism commits i to use the punishment mechanism p_i^j designed for that player.

Now a game in reciprocal mechanisms provides each player with a message space consisting of pairs in $\Theta_i \times [T \times [0,1]^n$ sent over two rounds. The mapping $\lambda_i$is given by (3.3) for reciprocal mechanisms and is defined in the obvious way for direct mechanisms. Since a mechanism that commits to the same fixed action for all messages is a direct mechanism, and since each player is free to send the same message for every possible value of his type, the game in reciprocal mechanisms is *regular* as defined in (1).

4. Theorem

The set of Θ_i of mechanisms available to each player is small. One can imagine mechanisms more complicated than direct mechanisms. Reciprocal mechanism especially are restrictive. They allow players to respond to deviations, but only in a manner that is independent of what the deviation is. Nonetheless, we can show that when players are restricted to choose their mechanism in Θ_i, any incentive compatible and individually rational outcome function can be supported as an equilibrium in the corresponding competing mechanism game.

Let ω be an outcome function that is incentive compatible and individually rational in the sense of (3.1) and (3.2). Recall that an outcome function has the property that for every array of types t, $\omega(t)$is some joint distribution on A.

Theorem 2. *There is a Bayesian equilibrium in the contracting game with reciprocal mechanisms Θ that supports the outcome function ω if and only if ω is implementable by a mechanism designer in the sense that is satisfies (3.1) and (3.2). In this equilibrium, all players use a common mechanism θ^δ, players report their types truthfully at the second round to each player, and each player chooses a correlating message uniformly from the interval* $[0,1]$.

Proof. Index the action profiles in A in some arbitrary way. Let $\omega^k(t)$ be the probability assigned to action profile a^k by the outcome function ω when player types are given by the vector t. The notation a_i^k means the action taken by player i in action profile a^k. As mentioned above, we restrict attention to mechanisms where i constrains himself to send the same message to all the other players. Let $x = (x_i, x_{-i})$ be the vector of signals sent by all the players. Now define

$$d_i^\omega(t,x) = \left\{ a_i^k : k = \min_{k'} : \sum_{\tau=1}^{k'} \omega^\tau(t) \geq \lfloor \sum_j x_j \rfloor \right\}. \tag{4.1}$$

The notation $\lfloor y \rfloor$ means the fractional part of the real number y. This functions aggregates the signals x into a number between 0 and 1, uses this to choose the index of the action profile in A that depends on the type reports of all the players, then directs player i to take his part in this action profile.

To see how this is going to implement the desired outcome, suppose first that each of the players chooses x_i using a uniform distribution. The random variable $\lfloor \sum_j x_j \rfloor$ is uniformly distributed on $[0,1]$. In that case, the function $\{d_1^\omega(t,x), \dots, d_n^\omega(t,x)\}$ implements the action profile a^k with probability $\omega^k(t)$ as the outcome function requires. The more interesting property of this construction is that for each value of x_i, $\lfloor x_i + \sum_{j \neq i} x_j \rfloor$ is also uniformly distributed on $[0,1]$.[5] This means that whatever else is happening, it is a best reply for each player i to select a number x_i using a uniform distribution provided he believes the others are doing the same thing.

We use exactly the same construction for the punishment

$$p_{ij}^{\rho_j}(t,x) = \left\{ a_i^k : k = \min_{k'} : \sum_{\tau=1}^{k'} \rho_j^\tau(t) \geq \lfloor \sum_{\tau \neq j} x_j \rfloor \right\}. \tag{4.2}$$

Let $\delta_i^\omega = \left\{ d_i^\omega, \{p_{ij}^{\rho_j}\}_{j \neq i} \right\}$ where $d_i^\omega$and $p_{ij}^{\rho_j}$ are defined by (4.1) and (4.2), and $\delta^\omega = \{d_i^\omega\}_{i=1,\dots,n}$. We now claim that if the outcome function ω is implementable, then there is a Bayesian equilibrium in the competing mechanism game in which each player i offers reciprocal mechanism $\theta_i^{\delta^\omega}$as defined by (3.3) using the list δ_i^ω.

To see why, imagine first that all players offer this reciprocal mechanism. Then all players are constrained to use the function d_i^ω to translate messages into actions. We have already explained that whatever type report a player sends to other players, he is completely indifferent about the message that he sends as long as he believes the others messages are chosen uniformly. As a consequence, there is a continuation equilibrium in which each player chooses his message uniformly. If all other players are revealing their true types to other players, (recall we are assuming here that every player is using a direct mechanism which constrains each of them to send the same message to every other player), then the payoff to player i if he also reveals his true type is

$$\begin{aligned} \mathbb{E}\left\{ u\left(d_1^\omega(t,x), \dots, d_m^\omega(t,x), t\right) | t_i \right\} &= \mathbb{E}\left\{ u\left(\omega(t_i, t_{-i}), t\right) | t_i \right\} \\ &\geq \mathbb{E}\left\{ u\left(\omega(t_i', t_{-i}), t\right) | t_i \right\} \end{aligned}$$

where this last expression is the payoff he gets if he lies about his type. The inequality follows from (3.3) and (3.1).

Now suppose that i deviates to some alternative contract c'. This must be unprofitable. To ensure this, we need to construct continuation play that makes i worse, conditional on the fact that, as a unilateral deviator, the other players are committed by the reciprocal contract, to carry out the punishments $p_{-i,i}^{\rho_i}$ against i. We assume that each player j continues to report type truthfully when implementing punishments, and that whenever i's contract asks the others to make reports, they report some arbitrary messages $(\bar{t}_{-i}, \bar{x}_{-i})$. Then since the action that player i takes cannot depend on the types of the other players, the payoff when player i of type t_i deviates is

$$\max_{\{t_{ij}, x_{ij}\}_{j \neq i}} \mathbb{E}\left\{ u\left(c'\left(\{t_{ij}, x_{ij}\}_{j \neq i}, (\bar{t}_{-i}, \bar{x}_{-i}) \right), p_{-i,i}^{\rho_i}(t_{-i}, x_{-i}), t \right) | t_i \right\} \leq$$

[5] This device is from the paper (A.T. Kalai and Samet 2010). A proof of this last property is given in (Peters and Troncoso-Valverde 2009). This proof also shows why $\lfloor \sum_j x_j \rfloor$ must be uniformly distributed.

$$\max_{a_i} \mathbb{E}\left\{u\left(a_i, p^{\rho_i}_{-i,i}\left(t_{-i}, x_{-i}\right), t\right) | t_i\right\} \leq$$

$$\mathbb{E}\left\{u\left(\omega\left(t_i, t_{-i}\right), t\right) | t_i\right\}$$

where the last line follows from (3.2).

To prove the other direction, begin with an outcome function ω that is supportable as a Bayesian equilibrium in reciprocal mechanisms. Player i's equilibrium payoff is

$$\mathbb{E}\left\{u\left(\omega\left(t\right), t\right) | t_i\right\}. \tag{4.3}$$

Let c' be any contract for which i's action is an onto function only his own correlating message x_i. In any equilibrium relative to reciprocal mechanism, this deviation must be unprofitable. In response to such a deviation, the other players will respond with some randomization over punishment mechanisms $\left\{p^{\rho_j}_{ij}\right\}_{j \neq i}$. The randomization is possible because players may be mixing when choosing among mechanisms. The messages and punishments that the others use can't depend on i's choice of action. Furthermore, choosing a payoff maximizing message at the second stage is equivalent to choosing a payoff maximizing action at the second stage since the contract is onto. It then follows

$$\mathbb{E}\left\{\max_{t'_i} \mathbb{E}_{\left(t'_{-i}, t_{-i}, x_{-i}\right)} u\left(a^{t'_i}_i, \tilde{p}_{-i}\left(t'_{-i}, x_{-i}\right), \left(t_i, t_{-i}\right)\right) | t_i\right\} \leq$$

$$\mathbb{E}\left\{u\left(\omega\left(t\right), t\right) | t_i\right\}.$$

This is almost the expression we want, except that when taking expectations on the left hand side. To complete the argument, let $\rho'\left(t_{-i}\right)$ be the probability distribution over the actions A_{-i} conditional on t_{-i} that is induced by the equilibrium strategies of the players other than i in the continuation equilibrium. Then from the inequality above

$$\max_{a_i} \mathbb{E}\left\{u\left(a_i, \rho'\left(t_{-i}\right), \left(t_i, t_{-i}\right)\right) | t_i\right\} \leq \mathbb{E}\left\{u\left(\omega\left(t\right), t\right) | t_i\right\},$$

which completes the proof. □

5. Equilibrium in regular Competing Mechanism Games

Theorem 2 provides a characterization of outcome functions supportable in reciprocal mechanism games. Reciprocal mechanism are convenient because their logic involves a very simple version of the logic that is used in repeated games. A outcome can be supported in equilibrium if every player prefers it to being punished. Outcomes are evaluated using direct mechanisms which are by n ow familiar. On the other hand, reciprocal mechanisms seem abstract. It is difficult to imagine a real world environment which would literally work this way. Part of the reason for this is simply that direct mechanisms themselves are abstract. As direct mechanisms provide a reduced form logic to understand single mechanism designer problems, reciprocal mechanisms provide a reduced form logic for regular competing mechanism games in the sense that every outcome function that can be supported as equilibrium of some regular competing mechanism game can be supported as an equilibrium outcome function in a reciprocal contracting game. In view of Theorem 2, all that is required to show this is to show that every outcome function supportable as an equilibrium in a regular competing mechanism game satisfies (3.2).

Theorem 3. *If an outcome function ω is supportable as an equilibrium in in a regular competing mechanism game, then it is supportable as an equilibrium in the reciprocal contracting game.*

Proof. By Theorem 2, we only need to show that there is a type contingent punishment $\rho_i\left(t_{-i}\right)$ such that (3.2) holds. Let $\omega\left(t\right)$ be the outcome function supported by some equilibrium of a regular competing mechanism game. Player i's payoff is given by (4.3). Since the game is regular, there are a collection of messages $\left\{m_{a_i}\right\}_{a_i \in A_i}$ for i each of which commits i to a different action, and each of which elicits some common type contingent response $Q_{-i}\left(t_{-i}\right)$ from the other players. Since each of these deviations is unprofitable in equilibrium

$$\max_{a_i} \mathbb{E}\left\{u\left(a_i, Q_{-i}\left(t_{-i}\right), \left(t_i, t_{-i}\right)\right)|t_i\right\} \leq \mathbb{E}\left\{u\left(\omega\left(t\right), t\right)|t_i\right\}.$$

Now simply set $\rho_i\left(t_{-i}\right) = Q_{-i}\left(t_{-i}\right)$ to prove the theorem. □

Theorem 3 differs slightly from theorem 2 since it only goes in one direction. Not all the outcome functions that can be supported with reciprocal mechanisms can be supported as equilibrium in an arbitrary regular competing mechanism game. This is as it should be. It is precisely restrictions on the players' ability to contract with one another that will narrow down the set of supportable outcomes. This is one of the advantages of Theorem 2. It benchmarks the restrictions imposed by the equilibrium logic of contracts. The makes it easier to assess the impact of more specific modeling restrictions.

This is one reason that both Theorem 2 and Theorem 3 refer to Bayesian equilibrium of competing mechanism games rather than equilibrium satisfying some kind of refinement. As discussed in (Peters and Troncoso-Valverde 2009), standard refinements can't be applied in competing mechanism games if players have too much contracting ability.[6]

6. Literature

(Epstein and Peters 1999) provides a type space and set of mechanisms which allows agents to convey market information along with information about their payoff type. They show that every mechanism that is offered in the equilibrium of a principal-agent type competing mechanism game coincides with a mechanism in *universal set of mechanisms* in which agents report types that convey all their market information. The set of mechanisms that is feasible in a particular game maps into a small subset of the universal set of mechanism. Nonetheless, they were able to show that provided mechanism were not restricted in how they dealt with payoff types, pure strategy equilibria are typically robust to expansion of the set of feasible mechanisms. Thus pure strategy equilibrium in 'naive' direct mechanisms (for example, the equilibrium in competing direct mechanisms described by (McAfee 1993)) can be supported as equilibrium relative to the universal set of

[6]This has nothing to do with reciprocal contracts. If message spaces are continuous and players can freely select functions that map these messages into commitments, then it is possible for them to commit to mechanisms that don't admit any continuation equilibrium. All of the well known refinements involving sequential rationality must then fail. To get them to work, restrictions must be imposed on message spaces or feasible mechanism. The theorem here suggests that the way to assess the refinement is to begin by characterizing Bayesian equilibria with restricted messages and mechanisms in order to understand the implications of those restrictions before assessing the implications of the refinement.

mechanisms. The difficulty with naive direct mechanisms is that they cannot be used to characterize some of the outcomes that can be supported as equilibrium relative to the universal set of mechanisms.

The literature on common agency (many competing principals, but only a single agent) tried to remedy this by abandoning the revelation principle, and simply asking for some set of indirect mechanisms that could be used to support all outcomes that might quality as common agency equilibrium. (Martimort and Stole 2002) and (Peters 2001) show that every (robust) equilibrium relative to any set of indirect mechanisms in common agency is an equilibrium relative to the set of menus. (Pavan and Calzolari 2009) show a similar result for common agency using what they call the set of 'extended direct mechanisms'. All robust pure equilibrium in common agency are equilibrium relative to the set of extended direct mechanisms.

As useful as the common agency tools are, they have two shortcomings. First, common agency is special since there can only be one agent, and principals can't communicate. Second, though the set of mechanisms (menus) that this literature offers is considerably simpler than the universal set of mechanisms, they are not sufficiently structured to allow a characterization of supportable outcomes.[7]

(Yamashita 2010) has recently suggested a way to extend the common agency logic to problems in which each principal has many agents. As in common agency, principals simply ask agents what to do, and commit themselves to carry out the recommendation as long as the majority of the recommendations agree. A characterization theorem for this case is given by (Peters and Troncoso-Valverde 2009) for competing mechanism games with at least four players and at least one principal.

One consequence of these theorems is that the set of allocations that can be supported as equilibrium with competing mechanisms is large. This fact has been observed before. Starting with the large literature on delegation games ((Fershtman and Judd 1987, Fershtman and Kalai 1997)), a number of papers have shown large equilibrium sets for special cases ((Katz 2006, Tennenholtz 2004, Yamashita 2010, Peters and Troncoso-Valverde 2009)). Our paper differs from these in two ways. First we impose no restrictions on the environment. (Katz 2006, Tennenholtz 2004), for example, assume complete information. (Yamashita 2010) assumes that players who offer contracts have no private information, and restricts the number of players. (Peters and Troncoso-Valverde 2009) restricts the number of players. We impose none of these restrictions.

Secondly, like the papers by (A.T. Kalai and Samet 2010)[8] and (Peters and Szentes 2008) we provide a complete characterization of supportable equilibrium outcomes rather than simply illustrating that a large number of equilibrium outcomes can be supported. However we do not assume, as do (A.T. Kalai and Samet 2010) that players have complete information.

Finally, this paper in many way relies on the arguments developed in (Peters and Szentes 2008). They show that the simple "cooperate or be punished" logic that is used to construct reciprocal contracts can be used to understand equilibrium in competing mechanism games no matter how rich the space of contracts is. There are three differences between the two papers. First, the space of feasible contracts

[7]Characterizations of outcomes for special environments have been given by (Peters and Troncoso-Valverde 2009). Though it might not be apparent why yet, we would also include (Tennenholtz 2004) and (A.T. Kalai and Samet 2010) in this category.

[8]We borrowed the randomizing trick in (4.1) from this paper.

in (Peters and Szentes 2008) is as large as it could be. In other words, the set of feasible contracts in that paper is much larger than the set considered here. The critical result in that paper is that equilibrium outcomes must satisfy (3.1) and (3.2) even when the space of contracts is very large. This makes it possible in this paper to turn the result around and show that equilibrium in a reciprocal contracting game can be used to support the same outcomes as can equilibrium with definable contracts. Absent the theorem in their paper there is no way to know whether the equilibrium outcomes described here are robust to the introduction of new contracts.[9]

The second difference is that (Peters and Szentes 2008) restrict players ability to communicate once contracts are announced. Here we allow players to continue to communicate, all the while making commitments about how this communcation will be handled. This is what is used to expand the set of equilibrium outcomes supportable with reciprocal contracts to include all outcome functions satisfying (3.1) and (3.2) rather than just a subset, as was the case in (Peters and Szentes 2008).

Finally, and least important, (Peters and Szentes 2008) restrict attention to pure strategy equilibrium. As a consequence, their characterization does not capture the randomization and correlation that are possible in competing mechanism games. As a consequence, their approach cannot be immediately adapted for a revelation principle.

7. An Example

To illustrate, we consider a well known issue in mechanism design - selling to buyers with correlated valuations. According to the well known argument in (Cremer and McLean 1988), correlation makes it possible to design selling mechanisms such that all players' payoffs are the same as they would be under complete information. This is accomplished by making the transfer each buyer is required to pay to depend on the types of the other buyers. An obvious use of competing mechanism logic is to try to explain away this odd result by using the argument that it can't be an equilibrium for competing sellers to leave both buyers with zero expected surplus. In this section, we illustrate how to support this result with reciprocal mechanisms. We also show how to undo the result by imposing restrictions on players' ability to contract.

For the sake of this simple example, suppose there are two sellers and two buyers (i.e. four players in all). Each seller has a single unit of output to sell to which he or she assigns a value of v_m. This fact is common knowledge. Each buyer has a private valuation, either v_l or v_h ranked the obvious way with $v_l < v_m < v_h$. Payoffs to the seller are equal to the money he receives less his value if he trades. Payoffs to each buyer are equal to their private valuation when they succeed in trading, less the money they pay. Each seller offers a mechanism, each buyer chooses to participate in one and only one mechanism. We assume that valuations are correlated. To make life simple suppose that both valuations are the same with probability $q > \frac{1}{2}$ and that they are equally likely to (both be) v_h or v_l in that case.

As for feasible actions, each seller can choose to give his good to either of the two buyers, or to keep it. He can also choose to make transfers to either or both

[9]Also note that the definition of regularity of competing mechanism games is adapted from their paper. They show that games in which players can use definable contracts are regular.

buyers. We will simply ignore the fact that the set of feasible transfers isn't finite. Similarly, buyers can offer to trade with either of the sellers or not to trade at all. They can also make transfers to either or both sellers.

First of all, we can illustrate a centralized mechanism which implements the result we want. In this mechanism, each buyer submits a type report to a centralized mechanism designer. After seeing the messages, the mechanism designer assigns the goods to to their final owners, and tells the players what transfer payments they need to make.

We would like to describe a 'direct' mechanism which is ex post efficient, and which provides each buyer type an interim payoff 0. Assuming that the other buyer will report his type truthfully, let p_h be the transfer that bidder i makes to the seller when bidder j announces type v_h. The efficient outcome has the bidder trading with some seller if and only if his type is high. Then in order to implement an efficient outcome in which each bidder's interim payoff is 0, we need to find a pair of transfers (t_l, t_h) such that

$$v_h - v_m + qt_h + (1-q)\,t_l = 0$$

and

$$(1-q)\,t_h + qt_l = 0.$$

Since $q > \frac{1}{2}$ and $v_h > v_m$, this pair exists and is unique. It is straightforward that $t_l > 0$, and $t_h < 0$. Adding the two equations together gives $(v_h - v_m) + t_h + t_l = 0$

Our centralized mechanism has each buyer report his or her type to a centralized and disinterested mechanism designer. The allocation rule is that a buyer who reports v_h is given a unit of output and pays the seller's cost v_m plus an additional fee that depends on the report of the other buyer. He pays this additional fee t_h if the other buyer's reported type is v_h, and gets t_l back if the other buyer's reported type is v_l. A buyer who reports a type v_l doesn't trade but receives the same conditional payments. Conditional on participation, neither bidder type wants to lie about his type. It is immediate that a high type buyer who claims to be low will expect to pay the seller a positive amount, while a low type buyer who claims to be high will expect a positive payment that is smaller than what he loses by buying a good he doesn't want. Since the expected surplus of both buyer types is zero, neither bidder can do strictly better by refusing to participate. The surplus earned by sellers under this scheme is then $(v_h - v_m)/2$, which is the surplus sellers earn with complete information.

By Theorem 3, the outcome function associated with this centralized mechanism can be supported as an equilibrium outcome function in the reciprocal contracting game. In applications, that is all that is likely to be important. We come back to that momentarily, but for now we can illustrate exactly how this would be done. The reciprocal contracting game takes place over two rounds. The messages in the first round specify commitments, the messages in the second round convey type information. For the sake of simplicity, we leave out the correlating messages and assume that ties are resolved efficiently. Then we can assume that sellers send no messages at all in the second round. Buyers send type reports, either v_h or v_l, to each seller (and to each other).

First round messages specify commitments. Commitments are either constant mechanisms or reciprocal mechanisms. Constant mechanisms for sellers are just ordinary direct mechanisms. Constant mechanisms for buyers are just simple unconditional commitments to participate in one of the seller's mechanism, both of them, or neither of them.

Reciprocal contracts are indexed by lists of direct mechanisms. Since there are four players, every message associated with a reciprocal contract will consist of a list of 16 direct mechanisms - four direct mechanisms, one for each player, and four lists of three punishments. Here, we describe the messages and contracts that implement the Cremer McLean outcome as above. Let d_s^* be a direct mechanism that

- sells to buyer 1 at price $v_m + t_l$ if buyer 1 reports type v_h and buyer 2 reports type v_l or refuses to participate;
- sells to buyer 2 at price $v_m + t_l$ if buyer 2 reports type v_h and buyer 1 reports v_l or refuses to participate;
- sells to one of the two buyers at price $v_m + t_h$ if both buyers report v_h;
- provides the transfer t_h or t_l to a buyer who reports v_l depending on whether the other buyer reports v_h or v_l.

Implicit in this description is an assumption that ties are resolved in obvious ways. For example, if both buyers report v_h to both sellers, then both buyers will trade. If one buyer reports v_h and the other v_l, then seller 1 will trade if the buyer reporting v_h is buyer 1, etc.

For buyers d_b^* is simply a commitment to participate in both sellers' mechanisms. For punishment, suppose p^* simply means refuse to trade.

As above, the message $\{d_s^*, p^*, p^*, p^*, d_s^*, p^*, p^*, p^*, d_b^*, p^*, p^*, p^*, d_b^*, p^*, p^*, p^*\}$ is tied to a commitment to use d_s^* or d_b^* provided all players send this message in the first round, and to refuse to trade otherwise. When all players send this message, sellers get full information payoffs, buyers have expected payoff 0, and any deviation yields expected payoff 0.

As is apparent from this discussion, all four players participate in the contracting process. They are symmetric in terms of their bargaining ability, though very asymmetric in terms of the actions they control. What this means here is that buyers actively participate in punishing a seller who deviates to some other mechanism. The outcome has seller 1 trading with buyer 1 when buyer 1 has a high value. Seller 1 might want to deviate to a mechanism that allowed him to sell to buyer 2 in the event that buyer 1 has a low value. He could do this by sending some alternative message in the reciprocal contracting game. If he does, he knows that buyer 2 has committed himself to refuse to trade with him. So he won't be able to accomplish this even if the commitment he deviates to is more attractive for buyer 2.

A concern about this argument is that it seems too abstract. It is hard to imagine sellers sending messages like this. We emphasize that reciprocal contracting isn't intended as a descriptive model of contracting. Like direct mechanisms, reciprocal mechanisms are intended to characterize what is feasible. What is perhaps more unconventional about the mechanisms in this example is that sellers make commitments that are contingent on messages of other *sellers*. It is more common to restrict sellers so that their commitments can depend only on messages of buyers. The message of much of the literature on common agency is that unless one is willing to make further restrictions, reciprocal contracting occurs anyway. The simplest variant of this argument is in (Yamashita 2010). If buyers messages allow them to send recommendations about which direct mechanism a seller should use, then the Cremer McLean outcome can again be implemented. They way this is done is to have sellers commit to use a direct mechanism if all buyers recommend it.[10] Buyers enforce the cooperative agreement by recommending the mechanism d_1^* or d_2^* because they expect all other buyers to. If a seller deviates and tries to implement some other mechanism than the agreement mechanism just described, then buyers recommend to the non-deviating sellers that they 'maxmin' the deviator.

By theorem 3, the allocation rules supported by these 'agreement' mechanisms can be implemented with reciprocal mechanisms as well. This suggests that one way to use reciprocal mechanisms is as a kind of reduced form method for understanding outcomes that are associated with contracting processes that aren't necessarily easy to observe in practice. One example like this might be the way that collusion is carried out. As collusion is often illegal, it would be unusual to have detailed information about how colluding parties come to agreements.[11] Reciprocal contracts provide a characterization of the set of outcome functions supportable under collusion without the need for an explicit description of the way the collusive process works.

To see how this works, we can imagine that both sellers simply offer direct mechanisms (not reciprocal mechanisms) d_s^* as above. and ask how buyers could collude against these mechanisms. Now we can simply use Theorem 3. In the collusive situation, bidders are trying to control their type reports to the two sellers. They will attempt to arrive at some mutual (and enforceable) agreement about how to report. Since we don't know exactly how they will go about doing this, Theorem 3 provides a useful approach. A collusive mechanism is a mapping that takes a pair of types and converts them into commitments involving the 4 messages that the buyer send to sellers. Theorem 3 say that any collusive mechanism that is incentive compatible and individually rational as in (3.2) can be implemented as an equilibrium in the reciprocal contracting game.

This simplest way for buyers to collude against the sellers mechanisms is for them to agree that both of them will send the type report v_l to both sellers. In that case each of them would receive a transfer t_l. Since this outcome is independent of any messages the buyers might exchange among themselves, it is easy to check that the buyers are both strictly better off doing this than they are refusing and receiving 0 payoff by participating in the sellers' mechanisms. One implication of this simply argument is that a common assumption in the collusion proofness literature finds some justification from reciprocal mechanisms. The assumption is that when a mechanism designer considers what is feasible for a group of colluding agents, he imagines that the agents can hire a disinterested coordinator who will implement the best collusive mechanism for them. This assumption is sometimes criticized for the fact that collusive proposals should come from individuals who are involved in the collusion. Reciprocal mechanisms illustrate how colluding players can implement allocations that appear to be guided by a disinterested planner, which are in fact, supported with wholly decentralized contracting.

---

[10] Formally, at least three buyers are required to make this work. The mechanism should commit to a direct mechanism if all buyers, or all but one of the buyers recommend it.

[11] One interesting counterexample is the paper by (Clark and Houde 2010) since the courts released transcripts of phone taps.

8. Conclusion

We have shown that all equilibria of competing mechanism games can be understood using reciprocal mechanisms. The advantage of this is that reciprocal mechanisms are conceptually no more difficult to work with than ordinary direct mechanisms. So reciprocal mechanisms provide a useful analytic approach for problems in which a broad class of mechanisms is feasible.

Like direct mechanisms, reciprocal mechanisms make it possible to understand equilibrium outcomes with competition without worrying about the intricacies of particular indirect mechanisms that are used in practice. Apart from the standard logic of incentive constraints, reciprocal mechanisms simply add the logic that if everyone else wants to do something, it is simple to write a contract that commits you to do it too.

References

A.T. Kalai, E. Kalai, E. L., and D. Samet (2010): "A Commitment Folk Theorem," *Games and Economic Behavior*, 69(1), 127–137.

Bernheim, B. D., and M. D. Whinston (1986): "Common Agency," *Econometrica*, 54(4), 923–942.

Clark, R., and J.-F. Houde (2010): "Collusion with Asymmetric Retailers: Evidence from a gasoline price fixing case," Discussion paper, Universite de Montreal.

Cremer, J., and R. McLean (1988): "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Economtrica*, 56(6), 1247–1259.

Epstein, L., and M. Peters (1999): "A Revelation Principle for Competing Mechanisms," *Journal of Economic Theory*, 88(1), 119–161.

Fershtman, C., and K. Judd (1987): "Equilibrium Incentives in Oligopoly," *American Economic Review*, 77, 927–940.

Fershtman, C., and E. Kalai (1997): "Unobserved Delegation," *International Economic Review*, 38(4), 763–74.

Han, S. (2006): "Menu theorems for bilateral contracting," *Journal of Economic Theory*, 127(1), 157–178, available at http://ideas.repec.org/a/eee/jetheo/v131y2006i1p157-178.html.

Katz, M. (2006): "Observable Contracts as Commitments: Interdependent Contracts and Moral Hazard," *Journal of Economics and Management Strategy*, 15(3), 685–706.

Martimort, D., and L. Stole (1998): "Communications Spaces, Equilibria Sets and the Revelation Principle under Common Agency," University of Chicago unublished manuscript.

Martimort, D., and L. Stole (2002): "The Revelation and Taxation Principles in Common Agency Games," *Econometrica*, 70(4), 1659–1673.

McAfee, P. (1993): "Mechanism Design by Competing Sellers," *Econometrica*, 61(6), 1281–1312.

Myerson, R. (1983): "Mechanism Design by an Informed Principal," *Econometrica*, 51, 1767–1797.

Pavan, A., and G. Calzolari (2001): "A Markovian Revelation Principle for Common Agency Games," Northwestern University Discussion Paper.

Pavan, A., and G. Calzolari (2008): "On the use of Menus in Sequential Common Agency," *Games and Economic Behavior*, 64, 329–334.

——— (2009): "Sequential Contracting with Multiple Principals," *Journal of Economic Theory*, 50, 503–531.

Peck, J. (1995): "Competing Mechanisms and the Revelation Principle," *Ohio State University Manuscript*, p. 12.

Peters, M. (2001): "Common Agency and the Revelation Principle," *Econometrica*, 69(5), 1349–1372.

Peters, M., and S. Severinov (1997): "Competition Among Sellers who offer Auctions Instead of Prices," *Journal of Economic Theory*, 75(1), 141–179.

Peters, M., and B. Szentes (2008): "Definable and Contractible Contracts," Working paper, University of British Columbia.

Peters, M., and C. Troncoso-Valverde (2009): "A Folk Theorem for Competing Mechanisms," manuscript, University of British Columbia.

Segal, I., and M. Whinston (2003): "Robust Predictions for Bilateral Contracting with Externalities," *Econometrica*, 71(3), 757 – 791.

Tennenholtz, M. (2004): "Program Equilibrium," *Games and Economic Behavior*, 49(2), 363–373.

Yamashita, T. (2010): "Mechanism Games with Multiple Principals and Three or More Agents," *Econometrica*, 78(2), 791–801.