RECIPROCAL MECHANISMS

MICHAEL PETERS

ABSTRACT. 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 Perfect Bayesian equilibrium in regular competing mechanism games is equivalent to the set of outcome functions supportable as Perfect Bayesian equilibria in a reciprocal contracting game. We provide a full characterization of this set of outcome functions. This characterization makes it possible to by pass game theoretic complexities in order to understand the impact of competition in competing mechanism games using a set of inequalities.

1. INTRODUCTION

In many interesting environments, competition between firms involves much more than simple price competition. Sellers at eBay attract buyers by allowing them to bid in auctions, or to purchase at a fixed (buy it now) price, or to choose between some combination of the two. Principals in a common agency vie for an agents' attention with non-linear pricing contracts, the benefits of which depend on exactly what kinds of non-linearities exist in the contracts offered by other principals. Groups of bidders in procurement auctions collude in an effort to compete both with other bidders and with the auctioneer.

One complication in dealing with these kind of competitive environments comes from the fact the main tool normally used to think about contracts - the revelation principle - doesn't work. What it means 'not to work' is that there are many outcomes that can be supported as equilibrium in competing mechanism games that cannot be supported if contract designers are restricted to direct mechanisms in which agents report their 'payoff types'. The reason is that agents have a lot of market information when they communicate with principals. Contracts that exploit this information can be used to support outcomes that look 'collusive'.

This makes for many equilibrium outcomes. This has been know at least since (Bernheim and Whinston 1986) who pointed out that common agency models have a large number of equilibrium outcomes. To find all equilibrium outcomes in common agency, it is necessary to allow competing principals to offer agents menus rather than direct mechanisms.[1] The logic is that the agent's preferences over a menu of alternatives will depend on the menus offered by other principals. Under those

April 15, 2012This paper is an extension of the paper "A Revelation Principle for Competing Mechanism Games" and should supersede it. The essential difference is that this paper deals with Perfect Bayesian Equilibrium instead of Bayesian Equilibrium. The conceptual and technical details associated this extension are considerable, so it is offered here as a new paper.

[1](Peters 2001) or (Martimort and Stole 2002).

conditions, a deviation by one principle can sometimes be met by a punishment that is meted out by the common agent.

This kind of logic is compounded in a 'multiple agency' (for example (Yamashita 2010)) in which a principal can ask his agents directly whether or not one of his competitors has deviated from some putative equilibrium, then commit himself to punish such a deviation. If there are enough agents, each of the agents will report a deviation truthfully simply because he expects the other agents to.

The upshot is that competing mechanism environments should normally be expected to support multiple equilibrium outcomes unless there are severe constraints on the contracts that players are allowed to offer or unless the environment is very specialized.

This leads to a second complication. The rules by which players interact in a competing contract environment are often obscured because the collusive outcomes they support are illegal. For example, a bidding ring in an auction supports its collusive behavior by agreeing in secret to some kind of punishment. Intellectual property lobbyists negotiate laws enforcing restrictive trade practices with politicians privately for much the same reason. If no one can see exactly what they are doing, it is hard to know what should be made illegal. This makes it even harder to understand what equilibrium outcomes look like, since we do not even know the rules of the game by which these contracts and outcomes are determined.

The objective of this paper is to provide a way around these two obstacles. We describe a game in which players make proposals about how outcomes should be determined. These proposals determine the contracts that convert the various messages that are sent during the game into commitments. We provide a full characterization of the set of outcome functions that can be supported as weak perfect Bayesian equilibria in this (relatively) simple competing mechanism game. We then show that any outcome function that can supported as a weak perfect Bayesian equilibrium in some *regular*[2] competing mechanism game can also be supported as a weak perfect Bayesian equilibrium in the reciprocal contracting game.

The advantage of this theorem is three fold. First, if we are interested in questions like whether competition will support efficiency, then we can address this by studying the impact of competition on the inequalities that characterize equilibrium rather than reasoning through complex game theoretic issues. If nothing else, conclusions about the impact of competition will then be robust to the way that competition is modeled. Furthermore, the inequalities that characterize equilibrium are just the inequalities that one might expect from a centralized mechanism designer. In this sense, the theorems in this paper bring back the revelation principle for competing mechanism games.

Second, for problems like collusion and corruption, in which the extensive form of the contracting game is unlikely to be understood, we can still use the inequalities to characterize the equilibrium outcomes that would be most (least) desirable from social point of view, then study the behavioral properties of these outcomes.

Finally, because our approach can be used to characterize equilibria in all competing mechanism games, it will also characterize equilibrium outcomes when players' ability to commit is restricted. For example, in a common agency, the agent

[2]A regular competing mechanism game is a game in which the introduction of contracts does not exogenously eliminate actions that are available as primitives to players.

can commit (to trade with a particular seller for example) but can't make this commitment contingent on any messages at all. From our main theorem, the set of outcome functions supportable under there restrictions must be contained in the set characterized here. For some of these restrictions, the approach here makes it easy to see what additional constraints are needed to understand the new set of equilibrium.[3] We briefly discuss these additional restrictions for problems where some players can't commit at all. Indeed, if no players can commit, it is easy enough to show that the set of outcomes supported by our characterization is simply the set of communications equilibrium ((Forges 1986)).

We begin by discussing an example that illustrates the approach. We then provided the main theorems in the paper before we return to discuss the relationship with other papers in the literature.

2. Example

Since modeling collusion is one problem for which the approach is likely to be useful, we consider a simple example in which buyers and sellers bid in a double auction. At this point we are simply interested in whether sellers can use reciprocal contracts to enforce what, for them, is a collusive outcome.

In this story, there are two sellers and two buyers (i.e. four players in all). Each seller has a single unit of output to which he or she assigns a value of 0. Each buyer has a private valuation, either v_l or v_h ranked in the obvious way with $0 < v_l < v_h$. Payoffs to the seller are equal to the money he receives while payoffs to each buyer are equal to their private valuation when they succeed in trading, less the money they pay. 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.

In the double auction that guides the interaction between them, players submit bids. The two available goods are awarded to the two highest bidders at a price equal to the third highest bid with the proviso that if there are more than two highest bidders, then the good is awarded to buyers whenever possible and randomly otherwise. There is a continuum of (ex post efficient) Bayesian equilibrium outcomes for this game in which all bidders bid $q \in (0, v_l)$ independent of type. The best that sellers can do in any of these equilibrium outcomes is a payoff of v_l, in which case, high value buyers earn $v_h - v_l$ and low valuation buyers earn nothing. In all of these equilibrium outcomes trade occurs for sure.

We are interested in whether there is some kind of collusive outcome in which sellers do better than they do in any of these efficient equilibria. What makes this a conceptually challenging problem is the fact that it is hard to know how sellers would negotiate such an agreement and how they would enforce it.

Whatever this collusive mechanism is, it must ultimately be inefficient. The sellers must restrict their supply somehow in order to extract some of the high value buyer's information rent. Whatever agreement the sellers reach will ultimately determine a trading price in the auction for the three different informational outcomes - both high value, both low value, different values. The argument we are trying to make is that in order to understand what sellers could do, we don't need to model the collusion directly. Instead, the best they can do can be understood by maximizing payoff subject to a set of inequalities.

[3]An example of this is (Celik and Peters 2011).

To illustrate, suppose the players come to some sort of agreement that results in three prices are p_{hh}, p_{ll} and p_{hl} that prevail when both buyers have high values, both buyers have low values, or buyers have different values. The sellers have to find a way to share the trading responsibility and to make the trading outcomes incentive compatible. Such a scheme would have to be incentive compatible and have the property that all the players would want to participate. It is reasonably straightforward in this simple environment to find the scheme that maximizes the expected profit of the sellers.

For example, assuming that one of the sellers submits a very high bid (recall that they have to reduce output to get higher prices), the incentive condition for a high value buyer is

$$q\left(v_h - p_{hh}\right) + (1-q)\left(v_h - p_{hl}\right) \geq$$
$$(1-q)\frac{1}{2}\left(v_h - p_{ll}\right). \tag{2.1}$$

He trades for sure if his value is high at a price that might depend on the value of the other buyer. Since one of the sellers submits a high bid, he will fail to trade if he pretends to be low value and the other buyer has a high value. If the other other buyer's value is low, he will have the same chance to trade as the other buyer - $\frac{1}{2}$.

Similarly, the incentive condition for the low value buyer is

$$\frac{q}{2}\left(v_l - p_{ll}\right) \geq q\left(v_l - p_{hl}\right) + (1-q)\left(v_l - p_{hh}\right).$$

As for participation, it is quite straightforward here. If a buyer or seller refuses to participate in the mechanism, the others can simply commit themselves to bids in the double auction that prevent the defector from earning a surplus. So if a buyer refuses to participate, the others all bid v_h independent of type, if a seller refuses to participate, the others all bid zero.

Sellers' expected surplus is

$$q\left(\frac{1}{2}p_{hh} + \frac{1}{4}p_{ll}\right) + (1-q)\frac{1}{2}p_{hl}. \tag{2.2}$$

We want to maximize this subject to the incentive constraints given above, and subject to each player earning non-negative expected surplus. Once we characterize the solution to this problem, we can show how to implement it as a perfect Bayesian equilibrium in the reciprocal contracting game.

Since a low value bidder only trades at price p_{ll}, he will not participate if $p_{ll} > v_l$. It is then immediate that $p_{ll} = v_l$ at the solution to this problem. The solution to the problem is then apparent from Figure 2.1:

The steeper of the two curves in figure describes the set of (p_{hh}, p_{hl}) pairs that make the high type buyer indifferent between revealing his type and pretending to be a low value buyer. This presumes that the price when both buyers claim to have low values is v_l. The high value buyer pays prices p_{hh} and p_{hl} when he trades. So if these are too high, he will be better off pretending to be low value and getting nothing. The set of pairs that are incentive compatible for the high value buyer are those below the curve for this reason.

The flatter of the two curves[4] represents the set of price pairs that make the low value buyer indifferent between revealing his type truthfully and pretending to be

[4]The curves have different slopes because the low and high value buyer have different beliefs about whether or not the other buyer has a high value.

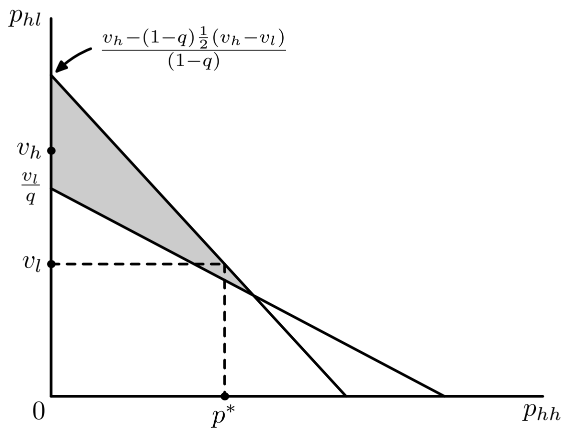

FIGURE 2.1. The Solution

high value. Reversing the reasoning above, if the prices (p_{hh}, p_{hl}) are too low, the low value buyer will want to pretend to have a high value so he can buy at these low prices. As a result, the prices that are incentive compatible for the low value buyer are those above this curve. The set of prices that are incentive compatible for both are those in the shaded area,

The sellers' iso-profit function has the same slope as the steeper of the two curves, as is readily seen by comparing (2.1) and (2.2) above. As a result, the sellers' best scheme implements any pair of prices on the upper right edge of the shaded triangle. For example, the solution we implement below has price equal to v_l when one of the two buyers has a low value, and p^* when both buyers have high values. It is straightforward that expected profits for sellers exceed their profits in every Bayesian equilibrium of the double auction provided v_h is high enough.

Our objective in this example is to show how this optimal collusive outcome for sellers can be implemented with reciprocal contracts. In words, what will happens is that sellers will coordinate such that one of them (chosen randomly) will bid p^* while the other will bid v_l. Buyers will bid p^* when they have high values and v_l when they have a low value. When both bidders have high values there are three bids at p_h^* so both buyers trade and pay p^*. When one buyer has a low value, one seller (the one who bid p^*) retains the good, the other pair trades at the third highest bid v_l. When both bidders are low, one of them trades with the low bid seller. Notice that this is an agreement, not an equilibrium in the bidding in the double auction. The high value bidders would like to reduce their bids since that would lower their trading price. So buyers are going to be part of this collusive agreement.

In this game, players simultaneously make proposals about how the game should be played and declare a type and 'correlating' message. The proposals are always public, however the information players have about the types and correlating messages depends on the proposals. If the proposals agree, the types and correlating messages are revealed to the players and the proposals are implemented. If the proposals disagree, a cheap talk game is played during a second stage. If all but one of the proposals agree, then the types and correlating messages of the players are revealed to everyone except the unilateral dissenter before the cheap talk game occurs. Otherwise, players do not get to see the type and correlating messages.

A 'proposal' is basically a list of direct mechanisms. When there is an agreement, these direct mechanisms use the type and correlating declarations to specify an outcome for each player. To make this clear, we need to describe these mechanisms.

The actions A_i for each player are bids q_i in some set. To keep a bound on notation, we assume this set is an interval, though in the main body of the paper it is assumed that actions are taken from a finite set.

Write T as the set of pairs of types of the two buyers, and $X = [0,1]^4$ as a set of four tuples of correlating messages. Let D_i be the set of measurable mappings from $T \times X$ into the set of bids. We refer to these mappings as direct mechanisms, though they are enhanced a bit by the addition of the correlating messages. Let $\delta_i = \left\{d_i, \{p_j\}_{j\neq i}\right\}$ be a list of 4 direct mechanisms for player i, with the interpretation that d_i is the mechanism i will use when there is an agreement, while p_j is the mechanism that i will use when player j refuses to participate in the agreement. Let $\delta = \{\delta_i\}_{i=1,\ldots,4}$ with Δ the set of all δ. The set Δ represents the set of proposals about how the double auction should be played.

Each player's first period choice is an element of$\Delta \times T_i \times [0,1]$. The game form λ that defines the outcomes associated with each profile of choices from this set is given by:

$$\lambda^i\left(\delta^1,\delta^2,\delta^3,\delta^4,t,x\right) =$$

$$\begin{cases} d_i\left(t,x\right) & \delta^1=\delta^2=\delta^3=\delta^4 \\ p_j\left(t,x\right) & \delta^j\neq\delta^i;\delta^i=\delta^{i'}\forall i'\neq j \\ a_i\in A_i & \text{otherwise.}\end{cases}$$

We can now explain the mechanisms that support the collusive outcome described above. Let $x = \{x_1, x_2, x_3, x_4\}$ be the vector of correlating messages. Define $\gamma(x) = \lfloor \sum_i x_1 \rfloor$, i.e., the fractional part of the sum of the correlating messages. For seller 1, define the direct mechanism

$$d_1^*\left(v,x\right) =$$

$$\begin{cases} p^* & \gamma\left(x\right) > \frac{1}{2} \\ v_l & \gamma\left(x\right) \leq \frac{1}{2}\end{cases}$$

while for seller 2 the mechanism

$$d_2^*\left(v,x\right) =$$

$$\begin{cases} p^* & \gamma\left(x\right) \leq \frac{1}{2} \\ v_l & \gamma\left(x\right) > \frac{1}{2}.\end{cases}$$

For buyers, define direct mechanisms

$$d_i^*\left(v,x\right) =$$

$$\begin{cases} p^* & v_i = v_h \\ v_l & v_i = v_l\end{cases}$$

For punishments, define the mechanism

$$\rho_b^*\left(v,x\right) = v_h$$

when the deviating player j is a buyer, and $\rho_s^*(v,x) = 0$ when the deviating player is a seller.

Our claim is that there is a perfect Bayesian equilibrium for the reciprocal contracting game in which each player makes the proposal

$$\{(d_1^*, \rho_s^*, \rho_b^*, \rho_b^*), (d_2^*, \rho_s^*, \rho_b^*, \rho_b^*), (d_3^*, \rho_s^*, \rho_s^*, \rho_b^*), (d_4^*, \rho_s^*, \rho_s^*, \rho_b^*)\} \tag{2.3}$$

each buyer declares his type truthfully, and each player chooses a correlating message using a uniform distribution on $[0,1]$. The outcome function supported by this equilibrium is the collusive outcome described above.

To see why, notice that if each player makes this announcement at the first stage, then each of them is committed to use the mechanism d_i^* to determine their final bid. Suppose that each of the players is expected to choose his or her correlating message uniformly from the interval $[0,1]$, and that buyers are expected to declare their types truthfully. If the correlating messages are all uniform, $\gamma(x)$ will be uniform. This device will ensure that half the time seller 1 sets a low price v_l and trades no matter what the buyer valuations, while seller 2 sets a high price p_h^* and trades only if both buyers values are high.

The transformation γ has the property that if each of the players chooses his correlating message uniformly, then $\gamma\left(x_i + \sum_{j\neq i} x_j\right)$ is uniformly distributed independent on x_i.[5] So it is sequentially rational to choose a correlating message using a uniform distribution.

Buyers need to carry out their part of the bargain, which commits them to bid p_h^* when they have high types and v_l when their types are low. The consequence is that there are three bids at p_h^* when both buyers have high values and both sellers trade at that price. If one of the buyers has a low value, then the seller who bids v_l trades with the high value bidders at price v_l. Finally when both buyers have low values, the seller who bid v_l trades at that price.

Notice that this is not part of an equilibrium in the bidding game - buyers are committed to bid p^* despite the fact that they realize they could lower the trading price by bidding less. Their contracts compel them to make this bid. To to check sequential rationality, it is only necessary to check that a high type bidder would rather bid p_h^* than to bid v_l. By definition, p_h^* is the highest price that has this property, so incentive compatibility is built in by design.

The consequence of deviating and announcing some other proposal in the first stage is to commit the others to a punishment that makes it impossible to earn surplus in the double auction. If a buyer deviates, the others all bid v_h. If a seller deviates the others all bid 0.

The upshot is that each player (including buyers) is better off proposing (2.3) than they are making some other proposal because any other proposal results in the other players punishing them much in the manner of a repeated game. The appeal of this extensive form is that it makes the competing mechanism logic trivial. This ought to make it much easier to understand how various contracting restrictions work.

This game describes a competing mechanism problem as a bargaining model between mechanism designers. It isn't an implausible descriptive story. Yet, there are many other stories. The reason this one is interesting is two fold. First, we

[5]See (Peters and Troncoso-Valverde 2009), who develop the idea from (A.T Kalai and Samet 2010).

have found a way to implement the most collusive outcome as a perfect Bayesian equilibrium. We found this most collusive outcome using standard mechanism design logic - i.e., constrained maximization. What we are going to show below is that very generally, the set of outcomes supportable as perfect Bayesian equilibrium in reciprocal contracting games coincides with the set of outcome functions that are implementable in the mechanism design sense. So despite the fact that the revelation principle applied to competitive problems doesn't 'work', there is nevertheless a way to use standard revelation principle arguments to understand outcomes supportable as equilibria in competing mechanism games.

Second, even if one has strong reason to believe that the commitment ability players possess in the reciprocal contracting game doesn't exist in some environments, there is a way that the game can be used to understand these environments. Our second main theorem below shows that any outcome that is supportable as a perfect Bayesian equilibrium in a *regular* competing mechanism game can also be supported as an equilibrium in the reciprocal contracting game.

In this sense, the reciprocal contracting game plays the same role for competing mechanism environments as the revelation principle does for single principle environments. Precisely, in a single principle environment, any incentive compatible and individually rational outcome function can be implemented as a Bayesian equilibrium using a direct mechanism. However indirect mechanisms may restrict players ability to communicate and commit. So not every incentive compatible and individually rational outcome can be supported with specific indirect mechanisms. Nonetheless, an outcome function supportable by any indirect mechanism can be supported with a direct mechanism.

The same logic applies here. Every appropriately defined incentive compatible and individually rational outcome function can be supported by perfect Bayesian equilibrium in the reciprocal contracting game. The same thing won't be true for arbitrary contracting games. However, if something can be supported as an equilibrium in some contracting game, then it can also be supported in a reciprocal contracting game.

The rest of the paper proves these things for the general case.

3. Incomplete Information Games and Mechanism Design

The basic approach in what follows is to add the contracting game on top of a basic game of incomplete information. We refer to this basic game as the *default game*. In default game, there are n players. Each player has a finite action set A_i and a finite type set T_i. 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. Similarly, define $T = \prod_i T_i$, and $T_{-i} = \prod_{j \neq i} T_j$. Types are jointly distributed on T according to some common prior.

Let q be a mixture over the set of action profiles A. The notation Q is used to represent the set of all such mixtures. For any action profile a, we write q_a to be the probability of a under q, and $q_{a_i} = \sum_{a_{-i}} q_{a_i, a_{-i}}$. We use notation q_{A_i} to represent the marginal distribution over A_i and $q_{A_{-i}}$ to be the marginal distribution over A_{-i}. We assume that players have expected utility preferences over lotteries. Then players preferences are given by $u_i : Q \times T \to \mathbb{R}$ where u_i is linear in q. An *outcome function* is a mapping $\omega : T \to Q$. So player i's payoff from this outcome function is $\mathbb{E}\left\{u_i\left(\omega\left(t\right), t\right) | t_i\right\}$.

One way to resolve this game is to have a mechanism designer collect information from the players, then tell each player what action to take. We think of the mechanism designer as an enforcer here, not just a coordinator - players have to carry out the action the mechanism designer tells them to whether they want to or not.

Obviously, the mechanism designer can only implement an outcome function ω if it is *incentive compatible.* Specifically, for every i, t_i and t_i',

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

Incentive compatibility as defined above is completely standard so there is no need to discuss it further. On the other hand, a player cannot be coerced into participating in this mechanism - he has to agree at the interim stage to be bound by the mechanism designer's ex post instruction. What makes this problem somewhat complex is the fact that the payoff to a player who chooses not to participate is not exogenous since he can still choose whatever action he wants in the default game.

In that event, we allow the mechanism designer to implement a punishment that includes a recommendation to the non-participating player about which action he should take. This recommendation may depend on the types of the participating players and should obey the usual obedience constraint which requires that the non-participating player should want to carry out any recommendation given his beliefs conditional on receiving that recommendation.

We can characterize the set of outcome functions that are implementable by a mechanism designer. Let $\rho_i : T \to Q$ be an outcome function that is implemented when player i chooses not to participate in the mechanism that implements ω. We refer to this outcome function as a *punishment.* The outcome function ω is *individually rational* if there is a collection of punishments $\{\rho_i\}_{i=1,n}$ such that for every player i,

$$\mathbb{E}\left\{u_j\left(\omega\left(t\right),t\right)|t_i\right\} \geq$$

$$\mathbb{E}\left\{u_i\left(\rho_i\left(t\right),t\right)|t_i\right\} \geq$$

$$\max_{t_i'}\sum_{\tilde{a}_i}\left\{\max_{a_i}\mathbb{E}\left\{u_i\left(a_i,\rho_{A_{-i}}\left(t_i',t_{i-i}\right),t\right)|t_i,\tilde{a}_i,t'\right\}\right\}\mathbb{E}\left\{\rho_{\tilde{a}_i}\left(t_i',t_{-i}\right)|t_i\right\}. \tag{3.2}$$

What follows shows that an outcome function ω is supportable as a weak perfect Bayesian equilibrium in *some* competing mechanism game if and only if there is a collection of punishments such that (3.1) and (3.2) hold. There are two parts to this. First, beginning with an outcome function that satisfies these constraints, we need to construct a competing mechanism game which can be used to support the outcome function as an equilibrium. We have already described this game in the introduction, it is the reciprocal contracting game.

Most of the work in the rest of the paper is devoted to explaining exactly how the reciprocal contracting game accomplishes this. This borrows a number of methods that are probably unfamiliar, so we break them up a bit in the discussion that follows to explain heuristically how they work. They are combined in the proof of the main theorem which appears in the appendix.

The other part of the proof is to show why 'regular' competing mechanism games have enough structure to ensure that the constraints given above are all satisfied in equilibrium. We defer this discussion until later in the paper.

3.1. **Reciprocal Contracting.** The overall approach in this paper is to model a competing mechanism game as an extensive form game of incomplete information. This extensive form is defined by a sequence of public and private messages along with a mapping that fixes the actions of the players as functions of the history of messages in the game. The reciprocal contracting game is a special case. It takes place in two stages. In the first stage, which we refer to as the commitment stage, players make public proposals about how the game should be played, and 'privately' state a type and a number in the interval $[0,1]$ which we refer to as a correlating message. These statements determine a set of mechanisms that each player will use to choose and action in the default game. These mechanisms condition directly on the mechanisms of the other players, somewhat in the manner of (Peters and Szentes 2012), though in a restricted and simple way.

What players learn about commitments at the end of the first stage depends on what the proposals are. If the proposals all agree, then the players learn the type declarations and correlating messages of each of the other players. If there is a single dissenting proposal, then the dissenter learns nothing about the type declarations or correlating messages of the others (though the non-dissenters again see all the type declarations and correlating messages). In every other case, the type declarations and correlating messages remain private.

In the second stage of the game, players send cheap talk messages to one another privately before any player who remains uncommitted at that point chooses his action.

In the first stage, players choose from the set $\Delta \times T_i \times [0,1]$. The set Δ describes the set of proposals that each player is allowed to make. A proposal is a list of mechanisms.

To understand the elements of this list, let $\hat{T} \equiv \prod_j [T_j \times [0,1]]$ be the set containing the type declarations and correlating messages of all the players. Let D_i be the set of measurable mappings $d_i : \hat{T} \to A_i$. We refer to each of these mappings as a *direct mechanism.* A proposal is a list $\delta = \{\delta_1, \ldots, \delta_n\}$ where $\delta_i = \left\{ d_i, \left\{ p_i^j \right\}_{j \neq i} \right\}$ and each d_i and p_j^i is a direct mechanism.

Each element δ_i of a proposal is a description of what player i should do. A proposal says that player i should use d_i along with the the type declarations and correlating messages to determine his action. However, if player j unilaterally refuses to go along with this agreement, the proposal says that player i should use p_j^i to determine his action. In this sense p_j^i is a punishment that player i will use against player j. Then a proposal $\delta^i \in \Delta$ by player i is a description of what player i thinks that each of the players should do. Abusing notation slightly, the reciprocal contracting game converts first stage messages into commitments in the following way:

$$\lambda_i\left(\delta^i, \delta^{-i}, t, x\right) = \begin{cases} d_i\left(t, x\right) & \exists \delta^* : \delta^j = \delta^* \forall j \\ p_i^j\left(t, x\right) & \exists \delta^* : \exists ! j : \delta^j \neq \delta^* \wedge j \neq i \\ a_i \in A_i & \text{otherwise.} \end{cases} \tag{3.3}$$

In this notation, p^i_j and d_i represent the corresponding elements of δ^i and the notation $\exists!$ means "there exists a unique".

By (3.3), each array of proposals announced by the players in the first stage publicly commits that player. In the case of agreement, each player is committed to a direct mechanism. When there is a single dissenter, the others are committed to direct mechanism while the dissenter simply chooses an action in the default game. In all other cases, the players simply play a cheap talk game in the second stage.

This formalism now makes it possible to state the first theorem.

Theorem 1. *If there are four or more players and ω is an outcome function satisfying (3.1), and (3.2), then there is a weak Perfect Bayesian equilibrium in the reciprocal contracting game that supports ω. Furthermore, along the equilibrium path of this game, all players make a common proposal δ^*, declare their type truthfully in the first stage, and choose a correlating message uniformly from the interval* $[0, 1]$.

The full proof is contained in Section 7.1. There are two basic complications involved in proving the theorem. The first, stems from the fact that the outcome function $\omega(t)$ involves a joint randomization, possibly involving correlation, over the actions of all the players, while any commitment δ by player i only commits him to a randomization over his own actions. This is where the correlating messages are used. We adopt a method from (A.T Kalai and Samet 2010) (and a slight generalization of it in (Peters and Troncoso-Valverde 2009)) which converts the private correlating messages into something that works like a public randomizing device that the players cannot manipulate. Once we have created this device (details are in the proof), it is straightforward to construct the contracts that implement ω and its various punishments ρ_j.

The second complication stems from the fact that when the mechanism designer punishes a player who refuses to participate, he might need to send the player an informative recommendation. Since there is no centralized mechanism designer in the reciprocal contracting game, we need to find a way to induce the punishing players to send the right recommendations on their own. These recommendations have to depend on the types of all the players. At the same time, there cannot be any incentive for the players who are communicating these recommendations to manipulate them. We accomplish this by having the deviating player ignore recommendations from the others unless they agree. This is where the assumption that there are four or more players is used. The deviator will hear at least three recommendations, and will follow them provided at least two of them agree. At the point where the punishing players send their recommendations, they (think that) they know the types of all the other punishing players. The reason they don't manipulate their recommendation is that they anticipate a very specific type contingent recommendation from the other punishing players and believe their recommendation will be ignored if it doesn't match. Then continuation play supports an outcome in which the punishing players send the same recommendations because each of them expects the others to send that recommendation.

3.2. **Competing Mechanism Games.** The reciprocal contracting game is but one example of a large set of potential games. The vast array of modeling choices available when modeling competing mechanisms makes it difficult to come to specific conclusions about things like the impact of competition. This is one reason that characterizing the set of outcomes supportable as equilibrium is useful, since it is possible to sidestep complicated game theoretic details. Of course, restricting players' ability to contract will impose restrictions on what can be supported in equilibrium. We return to this issue below. At this point, what we want to do is to show is that whatever restrictions seems appropriate in a particular application, the equilibrium outcome functions supported in the application can also be supported with reciprocal contracts.

The main complication in doing this is to try to describe in some fairly general way what a 'competing mechanism' game is. This is hard because there are so many different ways to approach competing mechanisms. 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. However a useful description should also capture models in which mechanisms are offered sequentially, as in (Pavan and Calzolari 2009) or privately as in (Segal and Whinston 2003).

Rather than trying to develop this tedious formalism, we take a slightly different approach here. We interpret a competing mechanism game as an extensive form game of incomplete information. We interpret the nodes of this game as opportunities for players to send messages. A path through the game is an ordered sequence of messages. Some of these messages convey commitments, some type information, while some are just cheap talk. In order to interpret the messages, we use an outcome function λ which assigns a profile of actions to each path through the game tree. The profile of actions indirectly determines each player's payoff.

The picture that follows shows the reciprocal contracting version of a simple prisoner's dilemma game (with the cheap talk part left out to make it simpler).

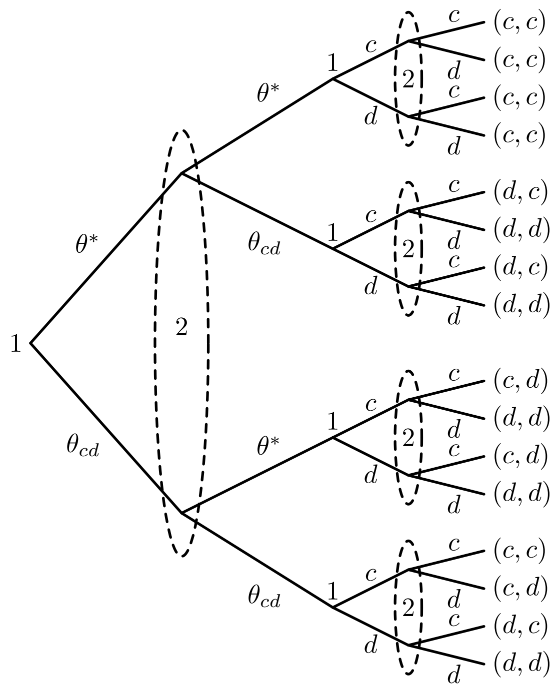

In this game, players announce public messages representing commitments over two rounds. In the first round, each player can offer a contract that conditions directly on the other player's contract. This is the contract θ^* in the picture. The alternative is a contract θ_{cd} that allows the player to defer his choice until the second round. The outcome function λ is displayed on the far right of the picture. Notice that if player 1 announces the message θ^* in his first information set, then the outcome function forces him to use action c in every history in which player 2 uses message θ^*, and to use action d in every other history following that choice. So from the outcome function λ, the interpretation of the message θ^* is that it is a reciprocal contract that commits player 1 to use action c if player to sends signal θ^*, and to use d otherwise.

Any set of behavioral strategies specify a possibly random path consisting of a sequence of messages. The outcome function λ converts every sequence of messages into a profile of actions for the players. Player i's payoff in the history in which players send the sequence of messages m is given by $u_i(\lambda(m), t)$. In the extensive form version of the reciprocal contracting game pictured above, the history of messages $\{\theta^*, \theta^*, d, d\}$ supports the profile of actions $\{c, c\}$.

Let $\{\sigma_i, b_i\}_{i=1,\ldots n}$ be behavioral strategies and beliefs for the players specifying mixtures over messages available to players in each of their information sets, and beliefs about the history of play prior to the information set. Let ι be an information set for player i. The continuation game associated with ι is the extensive form game of incomplete information in which each player's type is his payoff type from the original game along with his information about the history of play prior to ι. Beliefs for player i in this continuation game are given by $b_i(\iota)$. For every other player j, the player's type t_j in the continuation game describes (among other things) the most recent information set ι_j in which he sent a message. So player j's belief in the continuation game are his beliefs in the information set ι_j. Associated with each history $h \in \iota$, there is an outcome function that describes type contingent mixtures over action profiles when all players use the continuation strategies associated with $\{\sigma_i, \sigma_{-i}\}$ from the information set ι onward. Using i's beliefs in the information set ι, we write $\rho(t_i, t_{-i}|\sigma_i, \sigma_{-i}, \iota)$ as the outcome function conditional on attaining this information set when players are using the continuation strategies associated with (σ_i, σ_{-i}).

Given an array of behavioral strategies $\{\sigma_i, \sigma_{-i}\}$, a collection of information sets $\mathcal{I}$ is *attainable with probability* π by player i in the continuation game associated with ι if there is a continuation strategy for i at ι such that an information set in $\mathcal{I}$ is reached with probability at least π given i's beliefs $b_i(\iota)$ and the continuation strategies σ_{-i}.

An information set ι for player i has the *no-commitment* property if (i) the outcome function $\rho_{A_{-i}}(t_i, t_{-i}|\sigma_i', \sigma_{-i}, \iota)$ is independent of σ_i', and (ii) for each $a_i \in A_i$, there is a strategy σ_i' such that $\rho_{A_i}(t_i, t_{-i}|\sigma_i', \sigma_{-i})$ assigns probability 1 to the action a_i. In words, a no-commitment information set is one in which i can carry out any action he likes without changing the behavior of the other players. We say that a player i is uncommitted in information set ι if he has a continuation strategy that attains an information set having the no-commitment property with probability 1.

Definition 2. A contracting game is said to be *regular* if for every profile σ of strategies, each player i has a strategyσ_i' that attains some no-commitment information set with probability 1.[6]

This restriction is imposed because we are interested in adding contracts that enhance players' strategy sets, not in contracting games that impose arbitrary restrictions on what players can do. For example, consider the complete information game of matching pennies (with payoffs 1 and -1). This game has a unique Nash equilibrium in which both players' payoff is zero, which makes it pretty predictive by game theoretic standards. We already know we could change the outcome of this game by removing actions, or changing timing. We want to know whether contracts that both players would want to use might change the set of equilibrium outcomes for the game.

Suppose we specify the following contracting game: player 2 is allowed to choose one of two contracts. The first commits him to tails, the second to heads. We now allow player 1, still moving simultaneously with player 2, to commit himself in a manner that depends on the commitment made by player 2.[7] The only equilibrium would then have player 1 committing to match (or mismatch) the commitment of player 2. Payoffs would then be 1 for player 1 and 0 for player 2.

In this example, we would say equilibrium strategies are not regular for player 2. If player 1 is using his equilibrium strategy, then there are no strategies available to player 2 that allow him to change actions without simultaneously changing player 1's response. The contracting game we build on top of the matching pennies is simply depriving player 2 of the ability to select his action simultaneously with player 1.

It is possible to use the methods we describe below to analyze irregular games. For example, in the matching pennies example, if the asymmetric contract structure

[6]This definition is inspired by a similar assumption in (Peters and Szentes 2012).

[7]This is how the meet the competition argument works.

seems the right one for some reason, then we could analyze it by changing the original game from matching pennies to sequential matching pennies. The contracting game would then be regular with respect to this sequential game.

4. The Equivalence of Competing Mechanisms and Reciprocal Contracting

We can now state the main theorem.

Theorem 3. *Suppose the outcome function ω can be supported as a weak Perfect Bayesian equilibrium in some regular contracting game. Then there is a collection of punishments $\{\rho_i\}_{i=1,\ldots,n}$ such that (3.1), and (3.2) hold.*

Proof. Let $\omega(t)$ be the outcome function supported by some equilibrium of a regular competing mechanism game in which strategies are σ^*. It satisfies (3.1) by the usual revelation principle.

The game is regular, so i has a behavioral strategy σ_i' that attains a no-commitment information set with probability 1. Write

$$\hat{\rho}\left(t|\sigma_i'\right) \equiv \mathbb{E}_\iota\left\{\rho\left(t|\sigma^*,\iota\right)|t,\sigma_i'\right\}. \tag{4.1}$$

In words, $\hat{\rho}$ is the outcome function that prevails when player i uses the behavioral strategy σ_i' then reverts to σ_i^* once a no-commitment information set is attained.

The payoff associated with σ_i' is

$$\mathbb{E}\left\{u_i\left(\hat{\rho}\left(t|\sigma_i'\right),t\right)|t_i\right\}.$$

A player with type t_i can mimic the behavior of a player of type t_i' by adopting the same mixture over feasible messages in each of his information sets as the type t_i' player does in each of his corresponding information sets. Modifying the behavioral strategy in this way provides a new behavioral strategy that attains a no-commitment information set with probability 1. The payoff to player a player of type t who does this is

$$\mathbb{E}\left\{u_i\left(\hat{\rho}\left(\left(t_i',t_{-i}\right)|\sigma_i'\right),\left(t_i,t_{-i}\right)\right)|t_i\right\}.$$

Furthermore, once this new behavioral strategy reaches a no-commitment information set, we can modify the strategy again by having the player with type t_i adopt his original strategy σ^* from that information set on. So if i has a behavioral strategy that attains a no-commitment information set with probability 1, then there must be a strategy such that

$$\mathbb{E}\left\{u_i\left(\hat{\rho}\left(t|\sigma_i'\right),t\right)|t_i\right\} =$$

$$\sum_{A_i}\mathbb{E}\left\{u_i\left(a_i,\hat{\rho}_{A_{-i}}\left(t|\sigma_i'\right),t\right)|t_i,a_i\right\}\mathbb{E}\left\{\hat{\rho}_{a_i}\left(t|\sigma_i'\right)|t_i\right\} \geq$$

$$\max_{t_i\in T_i}\sum_{A_i}\max_{a_{i'}\in A_i}\mathbb{E}\left\{u_i\left(a_i',\hat{\rho}_{A_{-i}}\left(\left(t_i',t_{-i}\right)|\sigma_i'\right),t\right)|t_i,a_i\right\}\mathbb{E}\left\{\hat{\rho}_{a_i}\left(\left(t_i',t_{-i}\right)|\sigma_i'\right)|t_i\right\}.$$

The equality follows from the law of iterated expectations and the fact that the joint distribution of actions of the other players is independent of a_i in every no-commitment information set. The inequality follows from the fact that σ_i' attains a no-commitment information set with probability 1. This verifies that the punishment $\hat{\rho}\left(\cdot|\sigma_i'\right)$ satisfies (3.2). □

Combining this theorem with Theorem 1 gives the following corollary:

Theorem 4. *An outcome function ω is supportable as an equilibrium in a regular competing mechanism game with four or more players, if and only if it is supportable as an equilibrium in the reciprocal contracting game.*

The basic logic of reciprocal mechanisms is quite simple - competing mechanisms are complex, but ultimately, it is possible to understand quite a bit about them by using well understood logic that looks much like the logic in repeated games. As with the literature on repeated games, this means that many things can be supported as equilibrium outcomes. It is important to understand that there are two distinct reasons for multiplicity here. As always, any particular competing mechanism game can have many equilibrium outcomes. For example, the reciprocal contracting game we described above has a large number of equilibrium outcomes.

However, there are also many different ways to model competing mechanisms. Each model can have many equilibrium outcomes. The reciprocal contracting game described above can be used to understand all these outcomes. This is analogous to the fact that there are many different incentive compatible outcomes that can be described using the revelation principle. In practice, some kind of external selection criteria has to be applied to choose among these outcomes. For example, in a problem with collusion one could maximize the payoff of the colluding players across all outcomes that satisfy (3.1), and (3.2) in order to identify behavioral properties that could be used to identify collusion. Furthermore, the reciprocal contracting game provides a convenient contracting game (analogous to a direct mechanism in the usual revelation story) that can be used to think about strategic issues.

One example of a regular contracting game is any cheap talk extension of the basic Bayesian game described in Section 3 that does not allow players any additional commitment ability. In any such game, players are uncommitted in every information set. The set of outcome functions supportable as equilibrium in such a game is just the set of communications equilibrium ((Forges 1986)) of the original game, and can be described formally by setting $\rho_i = \omega$ for each player in (3.2). Any communications equilibrium outcome is supported as an equilibrium in the reciprocal contracting game in the obvious way (since $\rho_i = \omega$ for each i).

Constraints on Contracting

As we have mentioned, most competing mechanism models make very specific assumptions about what can and can't be contracted on. As we have shown, the outcome functions supportable as weak perfect Bayesian equilibrium in such a game must be contained in the set of outcome functions supportable as equilibrium in a reciprocal contracting game. This suggests that constraints on contracting can be translated into constraints on the set of supportable outcomes.

To put it another way, no indirect contracting game can support a bigger set of outcomes than the reciprocal contracting game unless it somehow modifies the strategic position of some player in the original game. In this sense, the reciprocal contracting game is analogous to a 'complete markets' model where, in the contracting sense, everything is working as it should. If it is possible to articulate a constraint on contracting, the natural approach is to impose that constraint on the reciprocal contracting game in order to isolate its impact from other implicit restrictions imposed in the indirect contracting game. This makes it possible to get a better idea of how the constraint works.

Any kind of complete analysis of this issue would go well beyond the scope of this paper, so we only sketch the way this might work for a special case. A very simple constraint on contracting would be to assume that some players simply couldn't commit at all. This is simple because this constraint on contracting is likely to look the same in every contracting game. Call a player a *no-commitment player* if he is uncommitted in the sense of Definition 4.1 in every one of his information sets. That is to say, that no matter where he finds himself in the game, he always has a behavioral strategy that attains a no-commitment information set with probability 1.[8]

Suppose that the first m players in a contracting game are no-commitment players. The other $n - m$ players will be referred to as commitment players. Suppose there are at least 3 commitment players. If we simply add the restriction that players 1 through m make no proposals in the reciprocal contracting game, then it is straightforward to show that equilibrium will impose a constraint on the outcome function similar to the second part of (3.2) for each of the no-commitment players. In other words, whatever action the no-commitment players are supposed to take in a supportable outcome, it had better be the case that that action is a best reply for them conditional on knowing they are supposed to take that action. A similar restriction has to be added to the punishment.

To see what is learned from this, suppose we had instead modeled an indirect contracting game with no commitment players and managed to characterize its equilibrium. There are two possibilities to consider - there is an outcome in the indirect game that can't be supported in the restricted reciprocal contracting game and conversely. In the first case, the outcome function can be supported in the unrestricted reciprocal contracting game by the theorems above, so there must be something about the indirect game that is making it possible to by pass the contracting restriction. Whatever it is might be interesting, but not because of the contracting restriction. Similarly if there is an outcome function that is supportable in the restricted reciprocal contracting game, but not in the indirect game, the conclusion is that the indirect game is implicitly imposing more restrictions on contracts.

The point is simply to illustrate that the reciprocal contracting approach can be used to study the implications of restrictions on contracting even though the game itself supports a large number of equilibria.

5. 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

[8] An example would be an agent in an intrinsic common agency who can supply effort to any subset of principals that he wants.

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 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) tries 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.[9]

(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 Bayesian equilibrium using the Yamashita method is given by (Peters and Troncoso-Valverde 2009) for competing mechanism games with at least four players.

However, the method here is not based on the Yamashita approach, but on the approach in (Peters and Szentes 2012). What that paper does is to try to provide a description of what the broadest set of such contracts would look like. Since contracts can condition on whether contracts condition on other contracts, etc, these contracts can be self referential, so the definition of this set is not at all straightforward. The key result in that paper (at least from the perspective of this one) is the fact that if contracts are restricted to be definable, then the competing contracting game is regular (as defined above). It is this fact that makes it possible to represent very complex contracting environments with simple reciprocal contracts (which aren't self referential at all). The simplicity of reciprocal contracts makes it easy to see how to modify the information structure to eliminate the restrictions imposed by the (Peters and Szentes 2012) contract game, and to show a complete equivalence between competing mechanism games and simple mechanism design. The implicit restriction imposed in (Peters and Szentes 2012) is that players can only communicate their types to one another publicly (by their contract offer). The reciprocal contracting game eliminates this restriction, which is why it supports more equilibrium outcomes. Perhaps a less important difference between the two papers is that (Peters and Szentes 2012) restrict attention to pure strategy equilibrium.

[9]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.

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.

Of course, 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.

Secondly, like the papers by (A.T Kalai and Samet 2010)[10] and (Peters and Szentes 2012) we provide a complete characterization of supportable equilibrium outcomes rather than simply illustrating that a large number of equilibrium outcomes can be supported.

One remaining question in all this is what is the set of equilibrium outcomes that is supportable in (Yamashita 2010). The important difficulty in providing such a characterization is that the Yamashita contracting game is not regular. When a player deviates in the Yamashita game, the agents all see what the deviation is. In principle they can tailor the punishment they recommend to depend on what this deviation is. This supports a punishment structure that is more like maxmin than the minmax punishments described in this paper. A characterization of supportable Bayesian equilibrium outcomes for the Yamashita game is given in (Peters and Troncoso-Valverde 2009). What the set of perfect Bayesian equilibrium outcomes looks like in that game is unknown.

6. 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.

7. Appendix: Proofs

7.1. Proof of Theorem 3.

Theorem. *If the reciprocal contracting game has four or more players, there is a weak Perfect Bayesian equilibrium that supports the outcome function ω if ω satisfies (3.1), (3.2). Furthermore, along the equilibrium path of this game, all*

[10]We borrowed the randomizing trick in (7.1) from this paper.

players announce a common proposal δ^* *, declare their types truthfully, and choose a correlating message uniformly from the interval* $[0,1]$ *in the first stage.*

Proof. We start by showing that an outcome function that satisfies (3.1) and (3.2) can be supported as a weak Perfect Bayesian equilibrium. The reciprocal contracts that are announced along the equilibrium path of the game require a direct mechanism for each player and a list of punishment mechanisms. We begin by describing these mechanisms. Then we describe the equilibrium path proposals.

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. For $(t,x) \in \hat{T}_i$, let t_i and x_i be the type and correlating message declared by player i in the first round. The following mapping defines a direct mechanism:

$$d_i^\omega(t,x) = \left\{ a_i^k : k = \min_{k'} : \sum_{\tau=1}^{k'} \omega^\tau(t_i, t_{-i}) \geq \lfloor x_i + \sum_{j \neq i} x_j \rfloor \right\} \tag{7.1}$$

The notation $\lfloor y \rfloor$ means the fractional part of the real number y. This function aggregates the correlating messages into a number between 0 and 1, then uses this to choose an action profile in A. The mechanism then commits i to carry out his part a_i^k of the corresponding action profile a^k.

This will implement outcome a^k with probability $\omega^k(t_i, t_{-i})$ provided each of the x_j are uniformly distributed on $[0,1]$. The property of this construction that will be especially useful below, is the fact that as long as each of the other players is choosing x_j uniformly, the random variable $\lfloor x_i + \sum_{j \neq i} x_j \rfloor$ is uniform on $[0,1]$ for each value of x_i.[11] What this means is that the probability distribution over i's actions is independent of x_i. As a consequence, it is sequentially rational for i to choose his correlating message uniformly from $[0,1]$ (no matter what his commitment) provided he thinks the other players are doing the same.

By (3.2), there is a collection of punishments $\{\rho_i\}_{i=1,\ldots,n}$ associated with ω. For every such punishment, define for each of the players other than i the direct mechanism

$$p_j^{\rho_i}(t,x) = \left\{ a_i^k : k = \min_{k'} : \sum_{\tau=1}^{k'} \rho^\tau(t) \geq \lfloor \sum_i x_i \rfloor \right\}. \tag{7.2}$$

As above, this will implement j's part of the punishment providing reports are truthful and correlating messages are uniform.

We are now ready to give the strategies associated with the Perfect Bayesian equilibrium that supports ω. The proposals of each player are identical and are given by

$$\delta^* = \left\{ d_i^\omega, \{p_i^{\rho_j}\}_{j \neq i} \right\}_{i=1,\ldots,n}$$

By (3.3), if all players announce this same proposal, then each player i is committed to use the direct mechanism $d_i^\omega(t,x)$.

If all players make proposal δ^*, declare their types truthfully, and choose a correlating message using a uniform distribution, then every player should anticipate

[11] This device is from the paper (A.T Kalai and Samet 2010) who do this for two players. A proof of this last property when there are more than two players is given in (Peters and Troncoso-Valverde 2009).

the outcome function ω. Since ω is incentive compatible, it cannot pay for any player to deviate by announcing a false type. We explained above, a player cannot improve his payoff by choosing correlation messages from some other distribution than the uniform. Since all players actions are committed in this case, the cheap talk game in the second stage is irrelevant. So if there is a profitable deviation, it must include announcing a proposal other than δ^*.

By (3.3), a player who makes an alternative proposal learns nothing about the type declarations or correlating messages of the others, and simply chooses his action in the second stage. However, his type declaration and correlating message are still observed by the others. Furthermore, he should expect to receive private messages from each of the other players before he takes his action.

If it is player j who makes an alternative proposal, then each of the other players should send a message to j consisting of a recommended action $\tilde{a}_j$ that j is supposed to take in the default game. This action is based on j's type declaration in the first stage, and is given by

$$\left\{ a_j^k : k = \min_{k'} \sum_{\tau=1}^{k'} \rho_j^{\tau}\left(t_j, t_{-j}\right) \geq \lfloor \sum_{i'} x_{i'} \rfloor \right\}. \tag{7.3}$$

This is just j's part of the punishment outcome

Player j's strategy is to play the recommended action if he declared his type truthfully at the first stage, and if all or all but one of the recommendations coincide. If his initial declaration was false, then he should choose any best reply conditional on his beliefs after seeing the recommendation. If there are three or more distinct recommendations, then he should ignore the recommendations and play a best reply to $\rho_{A_{-j}}\left(t_{-i}|\tilde{t}_j\right)$ based on his interim beliefs where $\tilde{t}_i$ is his first stage declaration.

It is sequentially rational for the players other than j to make the equilibrium recommendation in this case because the don't believe they can change any outcome by deviating, given that the others make this recommendation. Sequential rationality is built into the continuation strategy for for j in every information set except those in which j declared his type truthfully and received no more than two distinct recommendations. Sequential rationality follows in this case from (3.2).

Finally, a player who unilaterally deviates and makes an alternative proposal should nonetheless declare his type truthfully in the first stage by the first part of (3.2). □

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.

Celik, G., and M. Peters (2011): "Reciprocal Relationships and Mechanism Design," Discussion paper, microeconomics.ca.

Epstein, L., and M. Peters (1999): "A Revelation Principle for Competing Mechanisms," *Journal of Economic Theory*, 88(1), 119–160.

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.

Forges, F. (1986): "An Approach to Communication Equilibria," *Econometrica*, 54, 1375–1385.

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 (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.

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

Pavan, A., and G. Calzolari (2009): "Sequential Contracting with Multiple Principals," *Journal of Economic Theory*, 50, 503–531.

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

Peters, M., and B. Szentes (2012): "Definable and Contractible Contracts," *Econometrica*, 80(1), 363–411.

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.