

RECIPROCAL CONTRACTING

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 bypass 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 agent's 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 compete 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 in competitive environments. What it means not to work is that there are many outcomes that can be supported as equilibrium in competing contract games that cannot be supported if contract designers

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 Yamashita¹—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 contract environments should normally be expected to support multiple equilibrium unless there are severe constraints on the contracts that players are allowed to offer.

This leads to a second complication, since the rules by which players interact in a competing contract environment are often obscured because the cooperative outcomes they support are illegal. For example, a bidding ring in an auction supports its cooperative 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 would be 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 provide a set of mechanisms called reciprocal mechanisms and a game, which is referred to as the reciprocal contracting game, in which players compete by choosing these reciprocal mechanisms. We provide a full characterization of the set of outcome functions that can be supported as weakly perfect Bayesian equilibria in this relatively simple competing mechanism game. We then show that the set of outcome functions supportable as equilibrium in some *regular*² competing mechanism game is equivalent to the set of outcome functions supportable in reciprocal contracting games.

The advantage of this theorem is twofold. 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.

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 a competing mechanism game, it will also characterize equilibrium outcomes when players' ability to commit is restricted. For example, in a common agency, the agent 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 these restrictions must be contained in the set characterized here. For some of these restrictions, the approach here makes it easy

²A 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.

to see what additional constraints are needed to understand the new set of equilibrium. 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.

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 cooperative 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 v_h . Each buyer has a private valuation, either v_l or v_h , realized the obvious way, with $v_l < v_h$. Payments to the seller are equal to the money he receives while payments 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 equal, i.e., 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 a bidder bids $q \in [v_l, v_h]$ independent of type. The best that sellers can do in any of these equilibrium outcomes is a payment 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 cooperative 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 cooperative 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 payment subject to a set of inequalities.

To illustrate, suppose the three prices are p_{hh} , p_{ll} and p_{hl} . The sellers have to find a way to share the trading responsibility and to make the trading outcomes incentive compatible. The scheme we implement will have one of the sellers submit a high bid which will result in a trade only if both buyers have high value. The other seller will submit a lower bid that will guarantee him a trade and determine

each other type reports and correcting messages in the second round. In the equilibrium path there is also a third round in which players pay a cheap talk game. This third round is trivial in the example, so we defer a more complete discussion. The correcting messages are numbers between 0 and 1. The public commitments are messages taken from a measurable set which is to be described. Each message is tied to a very specific commitment which depends on the public messages and the messages that each player privately sees at the end of the second stage.

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. The messages used by players in the first round are proposals about how each player should pay and what each player should do in case there is a disagreement. The rules of the reciprocal contracting game tie each proposal to a very specific commitment. Formally, proposals are lists of direct mechanisms. To describe these, we describe mechanisms that players might use to implement the cooperative outcome above.

Second stage reports are a private. The reason for private communication is that we might not want a defecting buyer or seller to know what these type reports are. However, it is important in this story that players make commitments based on messages they send. It is fundamental to mechanism design that messages be verifiable so that commitments can be based on them. This idea is at the heart of inscrutability (Myerson 1983) which allows an informed principal to hide his type information by making commitments based on a type report that he sends after agents have agreed to participate in his mechanism. One difference between the competing mechanisms framework and the usual one is that after the principal sends this message about his type, his agents may use his message as part of their own mechanisms. We are going to adopt the simple assumption that the principal can and does commit himself to send the same message to all of his agents.

Let D_i be the set of measurable mappings from n -tuples of correcting messages in $[0, 1]$ and pairs of type reports into the set of bids. Let $\delta_i = \{d_i, \{p_j\}_{j \neq i}\}$ be a list of 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 $\delta = \{\delta_i\}_{i=1,\dots,4}$ with δ the set of a δ . The set δ represents the set of proposals about how the double auction should be played.

Each proposal that the player makes in the first stage commits the player to a particular action that depends on the proposals of the other players. Specifically, if $\{\delta^i\}_{i=1,\dots,4}$ is a collection of proposals, then the commitment associated with these proposals for player i is assumed to be

$$\lambda^i(\delta^1, \delta^2, \delta^3, \delta^4) = \begin{cases} d_i & \delta^1 = \delta^2 = \delta^3 = \delta^4 \\ p_j & \delta^j \neq \delta^i, \delta^i = \delta^{i'} \forall i' \neq j \\ a_i \in A_i & \text{otherwise} \end{cases}$$

In words, this says that if player i makes a proposal δ^i and the other three proposals agree with his, then i is committed to carry out his part of the proposal. If one of the others' proposals agree with his, but there is a single player who disagrees, then i

is committed to carry out his punishment against that player. In all other cases, i is simply allowed to choose whatever action he desires in the third stage.

We can now explain the mechanisms that support the cooperative outcome described above. Let $x = \{x_1, x_2, x_3, x_4\}$ be the vector of coordinating messages. For example, for seller 1, the coordinating message x_1 is the message that he sends to each of the other players. As mentioned above, we assume his mechanism commits him to send the same message to each of the other players. The message x_2 is what he receives from seller 1. Define $\gamma(x) = \lfloor \sum_i x_i \rfloor$, i.e., the fractional part of the sum of the coordinating messages. For seller 1, define the direct mechanism

$$d_1^*(v, x) = \begin{cases} p_h^* & \gamma(x) > \frac{1}{2} \\ v_l & \gamma(x) \leq \frac{1}{2} \end{cases}$$

while for seller 2 the mechanism

$$d_2^*(v, x) = \begin{cases} p_h^* & \gamma(x) \leq \frac{1}{2} \\ v_l & \gamma(x) > \frac{1}{2} \end{cases}.$$

For buyers, define direct mechanisms

$$d_i^*(v, x) = \begin{cases} p_h^* & v_i = v_h \\ v_l & v_i = v_l \end{cases}$$

For punishments, define the mechanism

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

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

$$1 \quad \{ d_1^*, \rho_s^*, \rho_b^*, \rho_b^*, d_2^*, \rho_s^*, \rho_b^*, \rho_b^*, d_3^*, \rho_s^*, \rho_s^*, \rho_b^*, d_4^*, \rho_s^*, \rho_s^*, \rho_b^* \}.$$

The outcome function supported by this equilibrium is the cooperative 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 coordinating message uniformly from the interval $[0, 1]$ and that buyers are expected to report their types to each mechanism truthfully. Since all players commit themselves to send the same coordinating message to each of the other players, each of the players observes the same value $\gamma(x)$ —just as if it were a public signal. If the coordinating 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 valuations are high.

The transformation γ has the property that if each of the players chooses his corresponding message uniformly then $\gamma\left(x_i + \sum_{j \neq i} x_j\right)$ is uniformly distributed independent on x_i .³ So it is sequentially rational to choose a corresponding 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 bidder 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_h^* despite the fact that they realize they could lower the trading price by bidding less. Their contracts compel them to make this bid. 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 p_h^* .

The upshot is that each player is better off proposing v_l than they are making some proposal because doing so results in the other players punishing them much in the manner of a repeated game. The appeal of this extensive form is two fold - first it makes the competing mechanism logic trivial. This ought to make it much easier to understand how various contracting restrictions work because the game theoretic logic is minimal. It illustrates the set of outcome functions that can be supported in some competing mechanism game. However, most important, the set of outcome functions that can be supported coincides exactly with the set of outcome functions that can be supported in some competing mechanism game. We turn to the argument 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. In this basic 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_{-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 $u_i : Q \times T \rightarrow \mathbb{R}$ where u_i is linear in q . An outcome

³See (Peters and Troncoso-Valverde 2009), who develop the idea from (A.T. Kalai and Samet 2010).

function is a mapping $\omega : T \rightarrow Q$. So player i 's payoff from this outcome function is $\mathbb{E}\{u_i(\omega(t), t | t_i)\}$.

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

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

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 recommendation. 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 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. Whatever this punishment is, the players should still want to report their types truthfully while the non-participating player should want to carry out the recommendation of the mechanism designer after updating his beliefs conditional on receiving the recommendation.

Let $\rho_i : T \rightarrow Q$ 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

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

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

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

and for each player other than i there is some belief about distribution of player types such that

$$(3.3) \quad \mathbb{E}\{u_j(\rho_j(t_j, t_{-j}), t | t_j)\} \geq \mathbb{E}\{u_j(\rho_i(t'_j, t_{-j}), t) | t_j\}.$$

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 and beliefs such that (3.1), (3.2) and (3.3) 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; it is the reciprocal contracting game described above. Notice that the reciprocal contracting game plays the role of a direct mechanism in the usual revelation principle. General competing mechanism games are analogous to indirect mechanisms.

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 competing mechanism games have enough structure to ensure that the constraints given above are satisfied in equilibrium. We defer this discussion until later in the paper.

3.1 Reciprocal Mechanisms. The reciprocal contracting game mimics the classic competing auction game in which players simultaneously announce commitments to direct mechanisms, then report their types to each other. However, the commitment phase of the game may not fully determine the actions of the players. In the competing auction game, for example, buyers are still free to choose which seller to visit. More important here, a deviating player might deliberately want to deviate from his action choice. To deal with this we add a final stage in which players whose actions are not fully committed play a cheap talk game which determines the final outcome. The basic logic is that players attempt to come to a reciprocal agreement about how to play the game. The cheap talk game determines the outcome when they cannot find such an agreement. To capture all this, the reciprocal contracting game involves three parts - commitments, type reports and the cheap talk game.

To make commitments, the players simultaneously make proposals. The proposal is taken from a common message space Δ for each player. We describe these proposals and the commitments that are attached to them momentarily, though they are simply a generalization of the proposals described above.

In the reporting part of the game, each player i privately sends a signal in the set $[T_i \times [\cdot, 1]]$ to each of the other players. The first element of this signal is a type report, the second is a correlating message used to support randomization. We refer to this message as a type report, though the player reports more than just type. In particular, for reasons that will be apparent momentarily, we focus in what follows on situations in which the player chooses his correlating message randomly using a uniform distribution on $[\cdot, 1]$. To keep things conceptually simple, we say that a player reports his type truthfully if he announces his true type *and* chooses a correlating message uniformly.

In the cheap talk game, players tell each other what messages they heard in the reporting stage, then based on what they hear from the others, they recommend an action to any player whose actions are not yet fully committed. We refer to the reports about other players' type reports in the reporting stage as *recollections*. Of course, the advice that players give other players are just called recommendations. Formally, the messages available to player i in the cheap talk stage are taken from a finite set $M_i = T_{-i} \times A_{-i}$ - just the product of his recollections and his recommendations. Once the cheap talk messages have all been sent, any player who has not yet fully committed himself chooses an action. Since i sends a message in M_i to every other player in the cheap talk stage, the action space for every uncommitted player in the cheap talk game is $\Delta \times [T_i \times [\cdot, 1]]^{n-1} \times M_i^{n-1} \times A_i$.

We need first to define the proposals in Δ . Define $T_i \equiv [T_i \times [\cdot, 1]]^{n-1} \times \prod_{j \neq i} [T_j \times [\cdot, 1]]$. This is the set of type reports that player i observes at the end of the reporting stage. This is all the type reports i sent to other players, as well as all the type reports he received from other players. Let Div_i

measurable mappings $d_i: T \rightarrow A_i$. We refer to each of these mappings as a *direct mechanism*.

There are two differences between these mechanisms and more conventional direct mechanisms. First, a fundamental part of mechanism design with informed principals is Myerson's inscrutability principle (Myerson 1983) in which a principal is allowed to choose a mechanism from a set specified beforehand at the same instant that his agents send their type reports. This allows the principal to hide his type from agents at the point where they submit their type reports to him. We are going to give players the same option here. In the competing mechanism environment, however, players may want to hide the mechanism they are using from some other players. We let them do this formally by allowing them to condition their actions on the messages they *send* to other players. In particular, if a direct mechanism commits to an arbitrary action unless the message the player sends to each of the other players is the same, then we refer to the mechanism as a *homogeneous direct mechanism*. A direct mechanism that commits to send the same message only to some subset of the other players will be referred to as a *partially homogeneous direct mechanism*. Observe that committing to send the same message to a group of players is not the same thing as making the messages publicly precise because players can limit the set of the opponents who see the same message.⁴

The second sense in which these mechanisms differ from standard direct mechanisms is that they allow players to send one another coordinating messages. The role of such messages was first pointed out by Pech (195). Apart from the fact that such messages can be used to support coordination, they will also make it possible to implement what are effectively random mechanisms while still writing mechanisms as mappings from messages into actions rather than mixtures over actions.

A *punishment mechanism* p_i^j for player i to use against player j is a partially homogeneous direct mechanism in which i 's action is independent of player j 's type, and in which i commits to an arbitrary action unless the signals he sends to all players other than j coincide.

We can now describe the set of commitment messages Δ along with the commitments associated with each of them. Let $\delta_i = \left\{ d_i, \{p_i^j\}_{j \neq i} \right\}$ be a list consisting of a homogeneous direct mechanism along with $n - 1$ punishment mechanisms for player i to use against each of the other players. Heuristically, the first mechanism in this list is a cooperative mechanism that player i will use when the other players are cooperative. The punishment mechanisms represent what i will do when each of the other players unilaterally refuses to cooperate. The notation $\delta = \{\delta_1, \dots, \delta_n\}$ represents a list of such lists, one for each of the n players. Δ is the set of all such δ .

The function $\lambda_i: \Delta^n \rightarrow D_i$ describes the commitment associated with each array of signals. In what follows, the functions d_i and p_i use second round signals as their arguments. In that case, we can abuse notation slightly and write the commitments

⁴Myerson's inscrutable mechanisms implicitly assumes that the principal 'sends' the same message about his type to every agent. What is different here is that agents can also make commitments that are allowed to depend on this common message. It is possible to design mechanisms that check that a principal has sent the same message about his type to each of the other players. The way this is done is to have each of the players tell each other what messages they heard from the other players, and to commit themselves to arbitrary actions unless these reports agree. The details are in (Peters and Troncoso-Valverde 2009).

associated with each array of first round signals as

$$(3) \quad \lambda_i(\delta_i, \delta_{-i}) = \begin{cases} d_i & \exists \delta^* \quad \delta_j = \delta^* \forall j \\ p_i^j & \exists \delta^* \quad \exists j \quad \delta_j \neq \delta^* \wedge j \neq i \\ a_i \in A_i & \text{otherwise} \end{cases}$$

In this notation, p_j^i and d_i represent the corresponding elements of δ_i and the notation \exists means there exists a unique

By (3), each array of reciprocal mechanisms announced by the players in the first stage publicly commits that player to either a homogeneous direct mechanism, a partially homogeneous punishment mechanism, or allows the player to choose whatever action he likes after listening to cheap talk messages. If the commitments agree, all the players can do according to the first line of (3), is to send type reports to the other players. These reports fully determine all the actions, so the cheap talk game at the end is irrelevant. If there is a single dissenting player (second line on the right hand side of (3)), that player simply listens to the cheap talk messages then takes an action. The non-dissenting players send one another type reports, send their recommendations to each other, then make recommendations to uncommitted players. In every other history, players play a cheap talk game in the third and fourth rounds. A weak Perfect Bayesian equilibrium of this Bayesian game is defined in the usual way.

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), (3.2) and (3.3), 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 announce a common signal δ^* in the first stage, report their types truthfully, and choose a correlating message uniformly from the interval $[\underline{\cdot}, \bar{\cdot}]$ in the second stage.*

The full proof is contained in Section 4. The important part and more difficult part of this theorem is the part that shows how to support outcome functions ω that satisfy the restrictions given by (3.1), (3.2) and (3.3) as equilibria. The essential logic is straightforward - if all the other players match a commitment \bar{d} then all of them are obliged to carry it out. If some player decides to deviate, then the others implement a punishment. The outcome function ω along with the punishments $\{\rho_i\}$ and inequalities provided by the centralized mechanism provide the basic building blocks that are used to construct the strategies in the reciprocal contracting game.

There are two basic complications. The first stems from the fact that the outcome function ω typically involves a joint randomization, possibly involving correlation over the actions of all the players, while any commitment \bar{d} 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 Aikawa and Samet (1997) and a slight generalization of it in Peters and Troncoso-Vera Verde (2016) 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 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 pay supports an outcome in which the punishing players send the same recommendations because each of them expects the others to send that recommendation.

The theorem is stated here for weak Perfect Bayesian equilibrium. It is possible to modify the story in a fairly straightforward way so that beliefs on the equilibrium path are consistent in the sense of sequential equilibrium. We restrict to arbitrary equilibrium paths here because the reciprocal contracting game isn't finite, so we are able to avoid a significant technical discussion about how to define sequential equilibrium in continuous games.⁵

3 Competing Mechanism Games. The reciprocal contracting game is but one example of a large set of potential games. It has a number of properties that seem unusual relative to the standard literature on common agency and competing auctions. For example, all players are allowed to communicate and to commit. In a common agency, for example, principals can't communicate with each other, and the agent has no commitment ability. In a competing auction game, bidders aren't allowed to send correlating messages. In a problem with collusion, the grand mechanism designer is competing with a group of colluding players who often go to great pains to hide the method they are using to organize. Papers like Yamashita (2007) or Peters and Troncoso-Vaquer (2012) show how the set of supportable outcomes explodes when players are allowed to communicate messages more complex than their payoff types. Sequential commitments support different outcomes than simultaneous commitments. Pavan and Calzolari (2009)

This vast array of modeling choices makes it difficult to come to specific conclusions about 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 the 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 that


⁵The only part of reciprocal contracts that isn't finite is the correlating messages. However, if mechanisms are Lebesgue measurable functions of correlating messages, then the set of mechanisms is actually fairly simple. The functions map into a finite set of actions. If the functions are required to be Lebesgue measurable, then they can all be characterized by a finite set of numbers in $[0, 1]$ where the functions jump from one action profile to another.

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 [1998], Yamashita [2000] or Peters and Troncoso-Vaquer [2005] or the literature on common agency: Pavan and Ca佐ari [2009] or Martimort and Stokey [2006] or Bernheim and Whinston [1996] 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 modes in which mechanisms are offered sequentially as in Pavan and Ca佐ari [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 1 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 2 sends signal θ^* , and to use d otherwise.

Any set of behavioral strategies specifies 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 $\{\sigma_i, b_i\}_{i=1, \dots, n}$ 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 beliefs 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 $\{\sigma_i, \sigma_{-i}\}$.

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 strategy σ'_i that attains some no-commitment information set with probability 1.⁶

⁶This definition is inspired by a similar assumption in (Peters and Szentes 2008).

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' payoffs are zero, which makes it pretty predictive by game theoretic standards. We already now 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 1 is allowed to choose one of two contracts. The first commits him to tails, the second to heads. We now allow player 2, still moving simultaneously with player 1, to commit himself in a manner that depends on the commitment made by player 1.⁷ 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 1. 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 1 of the ability to select his action simultaneously with player 2.

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 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.

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 a regular contracting game. Then there is a collection of punishments $\{\rho_i\}_{i=1,\dots,n}$ such that (3.1), (3.2) and (3.3) 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

$$1 \quad \rho|t|\sigma'_i \equiv \mathbb{E}_i\{\rho|t|\sigma^*, t|t, \sigma'_i\}.$$

In words, ρ 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

$$\mathbb{E}\{u_i|\rho|t|\sigma'_i, t|t_i\}.$$

⁷This is how the meet the competition argument works.

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 i of type t who does this is

$$\mathbb{E}\{u_i | \rho_{t'_i, t_{-i}} | \sigma'_i, t_i, t_{-i} | t_i\}.$$

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

$$\begin{aligned} & \mathbb{E}\{u_i | \rho_{t | \sigma'_i}, t | t_i\} = \\ & \sum_{A_i} \mathbb{E}\{u_i(a_i, \rho_{A_{-i} | t | \sigma'_i}, t) | t_i, a_i\} \mathbb{E}\{\rho_{a_i | t | \sigma'_i} | t_i\} \geq \\ & \max_{t_i \in T_i} \sum_{A_i} \max_{a_{i'} \in A_i} \mathbb{E}\{u_i(a'_i, \rho_{A_{-i} | t'_i, t_{-i}} | \sigma'_i, t) | t_i, a_i\} \mathbb{E}\{\rho_{a_i | t'_i, t_{-i}} | \sigma'_i | t_i\}. \end{aligned}$$

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 $\rho_{\cdot | \sigma'_i}$ satisfies 3.

The punishment $\rho_{\cdot | \sigma'_i}$ satisfies 3.3 for the players other than i because of the fact that each such player is using behavioral strategy σ^* , which is sequentially rational as part of a weakly perfect Bayesian equilibrium. \square

Combining this theorem with Theorem 1 gives the following corollary.

Theorem 4. *If an outcome function ω is supportable as an equilibrium in a regular competing mechanism game with four or more players, then 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 goes 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 compatibilities.

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 3.3 and 3.3. 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 cannot be contracted on. As we have shown, the outcome functions supportable as weak perfect Bayesian equilibrium in such 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 identical to the same in every contracting game. Call a player a *no-commitment player* if he is uncommitted in the sense of Definition 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 3.3 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

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

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 payers 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 bypass 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 [1] 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 mechanisms. Nonetheless, they were able to show that provided mechanisms 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 [2, 3] 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, 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 qualify as common agency equilibrium. Martimort and Stokey and Peters [4] 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 Ca'zolari [5] show a similar result for common agency using what they call the set of extended direct mechanisms. All robust pure equilibria 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

others is considerably simpler than the universal set of mechanisms; they are not sufficiently structured to allow a characterization of supportable outcomes.⁹

Yamashita [1] 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 [10] 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 [8], Fershtman and Kalai [9]), a number of papers have shown large equilibrium sets for special cases (Katz [6], Tennenholtz [11], Yamashita [1], Peters and Troncoso-Valverde [10]). Our paper differs from these in two ways. First we impose no restrictions on the environment (Katz [6], Tennenholtz [11], for example assume complete information; Yamashita [1] assumes that players who offer contracts have no private information).

Secondly, like the papers by A.T. Kalai and Samet [9] and Peters and Szentes [8], 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 [9], that players have complete information.

In this sense, the paper closest to this one is Peters and Troncoso-Valverde [10] which provides a characterization of outcome functions supportable as equilibrium using Yamashita's consensus mechanisms. The difference between the two is the equivalence result given here for reciprocal contracting games and regular competing mechanism games. Generally the set of outcome function supportable by consensus mechanisms is strictly contained in the set of outcome functions supportable as equilibrium with reciprocal contracting. The reason for the difference is that the reciprocal contracting game allows players who are punishing a deviator to send information to the deviator during the cheap talk part of the game.¹¹

Finally, this paper in many ways relies on the arguments developed in Peters and Szentes [8]. 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 in Peters and Szentes [8] 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' and 3'' 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 debatable contracts. Absent the theorem in their paper there is no way to now

⁹Characterizations 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.

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

¹¹To see why this might make a difference see (Celik and Peters 2008) or (Dequiedt 2006).

whether the equilibrium outcomes described here are robust to the introduction of new contracts¹²

The second difference is that Peters and Szentes restrict players ability to communicate once contracts are announced. Here we allow players to continue to communicate about the whole making commitments about how this communication will be handled. This is what is used to expand the set of equilibrium outcomes supportable with reciprocal contracts to include outcome functions satisfying 3.1 and 3.3 rather than just a subset, as was the case in Peters and Szentes.

Finally and least important, Peters and Szentes 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.

6 CONCLUSION

We have shown that a 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.

APPENDIX: PROOFS

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 and only if ω satisfies (3.1), (3.2) and (3.3). Furthermore, along the equilibrium path of this game, all players announce a common signal δ^* in the first stage, report their types truthfully, and choose a correlating message uniformly from the interval $[\cdot, 1]$ in the second stage.*

Proof. We begin by showing that any outcome function ω satisfying 3.1, 3.2 and 3.3 can be supported as a weak Perfect Bayesian equilibrium in the reciprocal contracting game. 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 and the corresponding equilibrium path reciprocal mechanisms.

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

¹²Also 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.

a_i^k For $t, x \in T_i$ let t_{ij} and x_{ij} be the type and corresponding message that player i sends to player j . The variables t_j and x_j are the corresponding messages that i receives from $j \neq i$. The vector of messages that i receives from the other players are t_{-i} and x_{-i} . Now define a homogeneous direct mechanism such that

$$d_i^\omega(t, x) = \begin{cases} a_i^k & k = \min_{k'} \sum_{\tau=1}^{k'} \end{cases}$$

The logic of the proof is then straightforward. If a player makes commitment δ^ω , he should anticipate the outcome function ω . If he deviates, he should anticipate players to implement a punishment. During the punishment he will hear cheap talk messages that provide him the same information that he would have received from the mechanism designer, and so he will anticipate the same punishment outcome he would have faced against the mechanism designer. Since the construction of the outcome function makes this unprofitable, he won't want to deviate.

The complication in proving this is to specify sequentially rational continuation pay that supports the outcome.

Continuation Play when all players make the same commitment: If all players make the same commitment δ^ω , each player should choose a correlating message using a uniform distribution on $[0, 1]$, then send that correlating message and their true type to every other player.

The cheap talk game is irrelevant in this case, so strategies for the continuation can be defined arbitrarily. As we explained above, if players report truthfully and choose their correlating messages uniformly, these mechanisms will jointly implement ω . Since ω is incentive compatible by 3.1 and the correlating messages are non-manipulable, these strategies are immediately sequentially rational.

We need to specify strategies for information sets in which players make a common commitment δ which does not coincide with δ^ω . In this case, we assume all players report correlating messages using the uniform distribution as above. This reporting strategy is non-manipulable, so no player has an incentive to deviate from this reporting strategy no matter type reports are being submitted. Conditional on this, players choose messages from a finite set, and each array of messages determines a payoff. Since this defines a finite Bayesian game, there exists at least one set of strategies that constitute a Bayesian equilibrium. For each common commitment, let these Bayesian equilibrium strategies constitute the continuation strategies of the players. They will evidently be sequentially rational given players' interim beliefs. Since these information sets play no role in supporting equilibrium, we will not discuss them again.

Continuation Play when there is a unilateral deviation in commitments: If there is a unilateral deviation to an alternative commitment by some player j , then each of the players other than j is committed to the partially homogeneous punishment mechanism $p_i^{\rho_j}$. Their actions are fully committed, however they need to send reports and possibly recommendations and recommendations during the cheap talk part of the game. If the common commitment is different from δ^ω , then we proceed as above and use any Bayesian continuation equilibrium to define strategies. So we focus on the case in which the common mechanism is δ^ω . For the rest of the process, the sequential rationality of the various strategy rules depend on what players expect to occur later in the game, so we specify the strategy rules then return to provide beliefs that make these strategy rules sequentially rational.

Reporting Strategies (A): If all players except player j make commitment δ^ω , then every player, including j , should report his type truthfully to each of the players other than j . We take it for granted here and from now on that this includes choosing the correlating message uniformly etc. The players other than j should babble¹⁴ when they are reporting to player j .

¹⁴Babble, means to send every feasible message with the same probability.

When j is a unilateral deviator, choose as part of the equilibrium three arbitrary players other than j who will communicate with j during the cheap talk game. We refer to these three players as j 's *communication partners*.

Recollections (B): If all players except j make commitment δ^ω , then each of j 's communication partners should truthfully report to each other the type report they heard from j .

Recommendation strategies are more complex since there are some off-path issues to deal with. To deal with on-path recommendations, observe that in the centralized implementation described in 3.1, 3.2 and 3.3, when player j refuses to participate, he nonetheless sends a type report to the mechanism designer, who then acts in the role of a mediator, and sends a recommendation to him that depends on the type t_j he reported, and on the reports t_{-j} the mechanism designer has received from the other players. The probability distribution over this recommendation is the marginal distribution $\rho_{A_j}(t_j, t_{-j})$ since j is supposed to want to carry out this recommendation. Let $\rho_{A_j}^{a_j}(t_j, t_{-j})$ be the probability assigned to action a_j by this marginal distribution.

Recommendations (C): For each communication partner i , if at least one of the recollections of the other two communication partners coincides with the report \tilde{t}_j that i received from j , then i should send j the recommendation

$$3 \quad a_j = j = \min_{k'} \sum_{\tau=1}^{k'} \rho_{A_j}^{a_{k'}}(\tilde{t}_j, t_{-j}) \geq \lfloor \sum_{i' \neq j} x_{i'} \rfloor,$$

where t_{-j} and $x_{i'}$ are i 's recollection of the type reports from the players other than j . Otherwise, each of the communication partners should babble to j . Each of the players who is not a communication partner with j should babble to j .

Note that, as above, if players correlate messages are chosen uniformly and the recollections of all three players coincide, then from j 's perspective, each of the three players will send him the same recommendation, and each such recommendation a_j will occur with the same probability as it would have if the centralized mechanism designer were making the recommendation after hearing reports \tilde{t}_j and t_{-j} .

The strategy of the unilateral deviator depends on the recommendations he receives. To define strategies, choose some action that maximizes his expected payoff given his interim beliefs and given that he expects the other players to report their types truthfully to one another. In other words, he should anticipate the punishment $\rho_{A_{-j}}(t_{-j})$. We will refer to this action as player j 's interim best reply.

Player j now expects to hear the same recommendation from exactly three players. Furthermore, this recommendation should have positive probability given the strategy defined by 3.3 and his interim beliefs. The other players are babbling in our construction, so j should ignore any recommendations that they make. The information sets in which Bayes rule does not admit a well-defined posterior belief are those in which the deviator j sends a common type report to at least two of his three communication partners but does not get a common recommendation back. Alternatively, Bayes rule does not give a posterior belief in information sets in which the three communication partners make a common recommendation, but this recommendation does not lie in the support of 3.3. We refer to these as *off-path* information sets.

Information sets in which Bayes rule is well defined are those in which j sent a common type report to at least two of his communication partners and receives a common recommendation, say a_j , in the support of β from each of them. What j does in these information sets should depend on the common type report. If this common type report was t_j , then he should update his beliefs about t_{-j} given the recommendation a_j using β with $t_j = t_j$. He should then choose a best reply to $\rho_{A_{-j}}(t_{-j})$. Whichever action this happens to be will be referred to as his best reply given the report t_j and recommendation a_j . We refer to these information sets as *on-path information sets with type report t_j and recommendation a_j* .

Actions (D): If every player but j commits to δ^ω , in any on-path information set with type t_j and report a_j , j should use the action that is a best reply to t_j and a_j . In particular if t_j is j 's true type, then he should obey the recommendation a_j of his communication partners. In any off-path information set j should use his interim best reply.

Finally in all the histories for which there are three or more distinct commitments, players simply play a cheap talk continuation game using their interim beliefs. Any strategy rules for the players that are part of a Perfect Bayesian equilibrium of this cheap talk game will act as appropriate sequentially rational continuation strategies. Such an array of strategy rules will exist because the game is finite. Again, since these information sets arise after unilateral deviations, we don't need to discuss them further.

This completes the description of the strategy rules. We can now verify their sequential rationality. We already explained why strategies are sequentially rational when all players make the same commitment, and why they are sequentially rational when players make 3 or more different commitments. So we begin with reporting strategies after a unilateral deviation. Case A above. In this case, each non-deviating player expects j 's communication partners to conform each others' recommendations by B, then make by C, the same recommendation to j that the centralized mechanism designer would have made. By D above, j is expected to carry out this recommendation. So players anticipate the outcome function ρ_j . It is then sequentially rational for players other than j to report their types to one another truthfully by 3.3.

The deviating player on the other hand, is not committed, so he can send different reports to each of the non-deviators. If he sends the same report t_j to at least two of his communication partners, then by B and C, he will face the punishment ρ_j , anticipating by 3 to receive the same distribution of recommendations as what he would have received from the mechanism designer. If he sends three different reports, then by the second part of C, he will receive an uninformative report. By 3.3, it is sequentially rational then for him to report the same type to all players.¹⁵

The sequential rationality of the non-deviators' reports to the deviator are immediate since the others don't pay any attention to them.

By B and the fact that j reports truthfully if i is one of j 's communication partners, he expects each of the others to hear a recommendation that coincides with the type report they heard from j , which they expect to coincide with the type report they heard. As a result, they believe that their own recommendation will be ignored.

¹⁵Observe the role that the recollections player here, since they prevent the deviating player from sending different messages to each of his partners in order to try to get more information out of them.

It is therefore sequentially rational for i to report j 's type truthfully no matter what he believes.

A similar reasoning is used for recommendations, though there are a couple of path possibilities. Again, let i be one of j 's communication partners who heard j report t_j . If i hears at least one report that coincides with t_j and believes on path that j reported truthfully to each of his communication partners, and that if there is a report from one of the communication partners that doesn't coincide with t_j , that this report was the result of a tremble, then he should believe that both his communication partners have reported t_j to each other. He will then expect both of them to make the same recommendation determined by β to j no matter what recommendation he sends. By D, any recommendation he makes will be ignored in this case. This establishes that it is sequentially rational for him to send a report determined by β .

On the other hand, if both of the reports he receives differ from t_j , then i should believe that j trembled from truthfully reporting to babbling independently to each payer. By D, this means that j is going to use his interim best reply no matter what recommendations he receives. Under these circumstances it is sequentially rational for i to babble when making his recommendation.

Finally, if j sends three different reports to his communication partners, he should expect them to truthfully report these to each other. By C, this would induce them to babble when making recommendations. Thus it is sequentially rational for j to use his interim best reply. If he makes two or more common reports t_j and receives a common recommendation a_j as a result, then it is sequentially rational for him to use his best reply to t_j and a_j by definition. If he makes two or more common reports but doesn't get a common recommendation, then any sequentially rational choice of action will do.

Following a β , β guarantees that it does not pay any payer to deviate from the commitment message δ^ω .

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

$$\mathbb{E}\{u_i(\omega(t), t) | t_i\}.$$

Let δ' made by some payer i that ensures that with probability 1, i chooses his action without commitment. Let $\rho'(t)$ be the outcome function associated with a deviation to δ' by payer i . The various conditions 3.1, 3.2 and 3.3 then follow immediately from the fact that ρ' is continuation pay associated with a weakly perfect Bayesian equilibrium, and from the fact that i cannot benefit by deviating to δ' . \square

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 (2008): "Equilibrium Rejection of a Mechanism," Working paper, University of British Columbia.
- DEQUIEDT, V. (2006): "Ratification and Veto Constraints in Mechanism Design," Discussion paper, Grenoble Applied Economics Laboratory no. 2006-06.

- 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.
- MC AFEE, 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 (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 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.