A FOLK THEOREM FOR COMPETING MECHANISMS

MICHAEL PETERS
UNIVERSITY OF BRITISH COLUMBIA
AND CRISTIAN TRONCOSO VALVERDE
UNIVERSIDAD DIEGO PORTALES MANUEL RODRÍGUEZ

Abstract. We extend a method of Yamashita (Econometrica 2010) to show that every allocation rule supportable by a centralized mechanism designer, including allocations involving correlated actions (and correlated punishments) can be supported as a Bayesian equilibrium outcome function in the competing mechanism game provided the game has four or more players.

In a competing mechanism game, multiple principals design contracts that commit them to actions that are conditional on messages they receive from other players. It is well known that static competing mechanism games can have equilibrium in which principals can support collusive outcomes by using mechanisms which require agents (or any other player with whom they communicate) to communicate *market* information along with information about their types. This possibility was initially mentioned in McAfee (1993), however, examples that illustrated this would work were offered later in Peck (1994),Martimort and Stole (1998), or Epstein and Peters (1999).

A characterization of a set of *contracts* that could be used to support all equilibrium outcome functions in a competing mechanism game was provided in Epstein and Peters (1999). These contracts were direct mechanisms in which agents' types were defined in a way that incorporated their market information. As the agents' types involved a hierarchical construction, these contracts were complex. For the special case of common agency, a much simpler set of contracts could be used to understand all equilibrium allocations. These were described in Martimort and Stole (2002),Peters (2001) and Pavan and Calzolari (2009). Similar attempts to describe a set of indirect mechanisms that can be used to support all competing mechanism equilibrium have been provided in special environments by Han (2006), and Andrea Attar and Portiero (2008).

August 19, 2011.

In a recent paper Yamashita (2010) suggests a method that could be used to characterize the set of supportable outcome functions. He proposes a class of mechanisms he calls *recommendation mechanisms*. The contracts resemble menus in that principals commit themselves to use whatever direct mechanism their agents recommend they use, provided the majority of agents make the same recommendation. Since disagreeing with the majority is never a strict best reply, a lot of outcome functions then become potential equilibrium candidates.[1]

Though Yamashita explains perfectly well how competing mechanisms can be used to support multiple outcomes, he doesn't provide an explicit theorem characterizing supportable outcomes. There are two reasons for this. First, he restricts players to use pure strategies both on and off the equilibrium path for expositional reasons. So a characterization isn't really the point of the paper. Secondly, as he describes what a characterization might look like, he describes a 'value' that imposes a lower bound on principals' payoffs for supportable outcomes. This 'value' is the lowest payoff that the principal attains from any mechanism he can offer in any continuation equilibrium against any array of mechanisms of the other players. Since the calculation of these values basically requires the calculation of all equilibrium strategy rules, the 'characterization' is really nothing more than a restatement of the definition of equilibrium. Our primary goal in this paper is to provide a (more useful) characterization. Our main result is that the set of Bayesian equilibrium outcomes in competing mechanism games is equivalent to the set of outcomes supportable by a centralized mechanism designer who collects information then controls the actions of all the players.

Proving this characterization theorem involves a number of issues that go far beyond the discussion in Yamashita. First of all, in order to provide a more useful characterization that can related to standard concepts in mechanism design, we focus on Bayesian equilibrium. This give us a simple relationship between supportable outcome functions and those that could be implemented by a centralized mechanism designer. For games of incomplete information, this brings us as close as possible to 'folk theorem' characterizations involving collections of inequalities that are familiar from the repeated game literature. A second advantage of the Bayesian approach arises from the fact that standard refinements are problematic in competing mechanism games.

---

[1] The argument is similar to, for example, Gerardi and Yariv (2007) who show the equivalence of voting rules that don't require unanimity. One key element of their logic is that when everyone else is expected to vote a certain way, it is always a best reply to do likewise (provided unanimity isn't part of the voting rule).

Since all refined equilibria are also Bayesian equilibrium, our approach gives some insight into these equilibria anyway (though the detail of the refinement will obviously provide more insight).

The role that mixed strategies play in has been widely discussed in the literature in common agency. In particular, papers like (Peck 1994) have shown how indirect mechanisms can be used to correlate principals actions. For this reason, it seems important to deal with randomization. when characterizing supportable outcomes. We incorporate a device that makes it possible to deal with mixing and correlation at the same time while still using familiar direct mechanisms to support outcomes. We borrow a some results from the computer science literature, in order to show how to implement correlated outcomes in an environment in which principals still communicate privately with their agents. For complete information games, our results show how the folk theorem like results in A.T. Kalai and Samet (2010) can be extended to arbitrary numbers of players. For games of incomplete information, our characterization involves the usual inequalities associated with incentive compatibility and individual rationality.

The customary way to model competing mechanisms, which the approach that Yamashita follows, is assume that each principal is uninformed and deals with the same set of informed agents. The agents have an interest in the outcome, but have no commitment ability. We want to extend the scope of our theorem to a much broader set of environments than is traditionally discussed in the literature. In particular, the extension to informed principals is valuable in problems like collusion Laffont and Martimort (1997) or Che and Kim (2006)) in which the problem of exactly who offers the collusive mechanism is often an issue. For example, we would like to provide a characterization of the set of outcomes that can be supported by a group of colluding agents who can make commitments to each other, but who otherwise don't have access to an uninterested outside coordinator. For this reason, we allow all players to communicate and make commitments.

The indirect mechanism that we add to achieve correlation with private communication is the only new conceptual idea in this paper. We exploit two properties of the competing mechanism environment that aren't typically considered in the existing literature. The first is that principals can communicate with one another after contracts are signed in much the same way agents do.[2] The second is the fact that

---

[2] In the principal agent setting, we can prove our characterization theorem without communication among principals provided there are enough agents. For example, in the environment considered by Yamashita, our theorem applies whether principals can communicate with each other or not. The one setting where we do

communication can be sequential. We use sequential communication to correlate principals' actions in a manner similar to the way that unmediated communication supports correlated equilibrium in games.[3]

Apart from the fact that we consider a much broader set of environments, we have been careful to adopt conventions that are standard in the literature. For example, we assume that principals communicate privately with their agents. We allow only non-random mechanisms and assume that players use pure strategies both on and off the equilibrium path for everything except for special messages we refer to as 'correlating messages' which are, in a sense we describe below, payoff irrelevant.

1. FUNDAMENTALS

There are $n \geq 4$ players. We sometimes write N to represent the set of players. Player i must choose an action a_i from a finite set A_i. Let $a = \{a_1, \ldots, a_n\}$ be an array of actions in $A = A_1 \times \cdots \times A_n$. $A_{-i} = \prod_{j \neq i} A_j$.

Each player i has a privately observed payoff type θ_i drawn from a finite set Θ. Payoffs are given by $u_i : A \times \Theta^n \to \mathbb{R}$. Players have expected utility preferences over actions.

Let P_i, P_{-i}, and P be the set of probability distributions on A_i, A_{-i}, and A respectively. A typical element $p \in P$ is a vector with p_k equal to the probability that the k^{th} element in A occurs, where the set A is indexed in some arbitrary fashion.

Let $q : \Theta^n \to P$ be an allocation rule. In what follows we slightly abuse notation by writing $u_i\left(q, \theta\right)$ instead of $\sum_{a \in A} q_a u_i\left(a, \theta\right)$. We are interested in allocation rules that are incentive compatible and individually rational. Incentive compatibility means

$$\mathbb{E}\left\{u_i\left(q\left(\theta\right), \theta\right) | \theta_i\right\} \geq \mathbb{E}\left\{u_i\left(q\left(\theta_i', \theta_{-i}\right), \theta\right) | \theta_i\right\} \tag{1.1}$$

for each $i \in N$, and $\theta_i' \in \Theta_i$. Individual rationality means that for each player i there is a punishment $p^i : \Theta_{-i} \to P_{-i}$ such that for every θ_i

$$\mathbb{E}\left\{u_i\left(q\left(\theta_i, \theta_{-i}\right), \left(\theta_i, \theta_{-i}\right)\right) | \theta_i\right\} \geq$$

$$\max_{a_i} \mathbb{E}\left\{\mathbb{E}\left\{u_i\left(a_i, p^i\left(\theta_{-i}\right), \left(\theta_i, \theta_{-i}\right)\right) | \theta_i\right\}\right\}. \tag{1.2}$$

---

need to allow principals to communicate is common agency. We return to this issue below.

[3] As in (Gerardi 2004), or (Forges 1990)for example. The difference here is that principals can commit to the way they respond to messages.

With complete information, an allocation is individually rational if and only if it provides each player with an expected payoff that exceeds his or her *minmax value,* defined for player i as

$$u_i^* \equiv \min_{p_{-i} \in P_{-i}} \max_{a_i \in A_i} u_i \left(a_i, p^i \right). \tag{1.3}$$

Again, with complete information the punishment

$$p_{-i}^* \in \arg \min_{p_{-i} \in P_{-i}} \max_{a_i} u_i \left(a_i, p^i \right)$$

can be used to support all implementable allocations.

Notice that when constructing a punishment, or a minmax value, punishers are allowed to correlate their punishments. This is appropriate for a mechanism designer who can enforce contracts and correlate actions among agents who have agreed to participate.

2. Competing Mechanism Game

Players determine their actions by writing contracts that restrict their actions conditional on messages they (privately) receive from other players. Feasible messages and mechanisms are different for different models. Sometimes, the set of feasible mechanism is very restricted. For example, in common agency, the only communication that is allowed to occur is that between a principal and a single agent. The messages the agent can send are also quite restricted. For example, the agent may be asked to pick an element of a non-linear price schedule. The characterization theorem we prove below requires a richer set of mechanisms. Rather than trying to be very general about this, we describe a relatively simple set of mechanisms that will work.

Let Γ_i be the set of all measurable mappings from $(\Theta \times [0,1])^{n(n-1)} \rightarrow A_i$. We are going to let the message space $\mathcal{M}_i$ for player i be $\Gamma_i \times (\Theta \times [0,1])^n$. The set of mechanisms for player i is then going to be the set of all measurable mappings $\mathcal{R}_i$ from $(\mathcal{M}_i)^{n-1}$ into A_i.

Let us explain. We are going to have each player i offer a mechanism that asks the other players to make a recommendation and then to report types and correlating message. Up to some qualifications that we specify below, we describe equilibrium in which each principal commit himself to follow the other players' recommendation provided they all agree. Recommendations are mappings that look pretty much like direct mechanisms save for the additional correlating messages.

The complication we need to address is to make sure that the type report (and correlating message) that another player provides him is the same type report that player sent to every third player. The way we are going to deal with this is to have players send messages in

$(\Theta \times [0,1])^n$ over two rounds. Each player will report his own type and correlating message in the first round, then report the messages that he heard from the other types on the second round.

As the process of reporting on reports means that our mechanisms explicitly involve sequential communication, so we refer to them as *sequential communication mechanisms*. It might also be noted at this point that these mechanisms do not on their own induce any randomization - they map into pure actions. We will induce randomization when we allow agents to randomize over the messages they send in $[0,1]$.

An equilibrium for the competing mechanism game is a Bayesian equilibrium of the usual sort. The players' strategies specify for each of their types, a mechanism and a rule that specifies the messages they send in each round as a function of the mechanisms offered by the other players, and the messages they received in previous rounds. When we need it, we use the notation Σ_i to refer to the set of strategy rules available to player i. The notation σ_i refers to a specific strategy. A Bayesian equilibrium is a collection of strategy rules that are jointly best replies to one another.

3. Theorem

At this point we can state our main theorem:

Theorem 1. *If there are* 4 *or more players, then an allocation rule can be supported as a Bayesian equilibrium in the competing mechanism game if and only if it is incentive compatible and individually rational.*

It is important to point out what this theorem adds to the logic in Yamashita. Most obviously it covers random, and even correlated outcomes that could not be captured because of the pure strategy non-random mechanism assumptions in Yamashita. Secondly, it extends the characterization from the uninformed principal informed agents framework to an environment in which there are informed principals. It covers common agency provided there are three or more principals, which is ruled out by Yamashita's approach. It also admits problems in which bargaining power is evenly distributed among players. At the most fundamental level, it provides a characterization in the form of a set of inequalities, which Yamashita's paper does not do, as we explained above.

Of course, Yamashita's point is not to provide a characterization in the first place. It is simply to show how recommendation mechanisms work. Our model goes beyond this. We start with the set of incentive compatible individually rational allocation rules, then show how to implement all of them.[4]

One interpretation of Yamashita's theorem is that permitting too big a set of feasible mechanisms makes the competing mechanism paradigm uninteresting since it potentially loses most of its predictive power. Since he doesn't provide a characterization, this interpretation of his theorem may be going too far. However, the theorem we present here looks more like a standard folk theorem. There is a sense in which this theorem can be used in a positive way. To explain it, we take a moment to describe the collusion proof-ness problem. The gist of the idea is from McAfee and McMillan (1992) who imagined a mechanism designer like an auctioneer designing a mechanism while realizing that bidders might want to collude against him. They suggested that the colluding bidders should be subject to the same incentive compatibility and individual rationality constraints as the primary mechanism designer. Ultimately, this led to the idea that the primary mechanism designer's mechanism should be designed to rule out the possibility that bidders could themselves designing a mechanism that is at once profitable, but also incentive compatible and individually rational.

One important problem in doing this is to decide who exactly will come up with the mechanism that guides the colluders. Our approach illustrates one possible answer. It suggests that modeling the collusive process as if it were designed by an outside coordinator can be 'decentralized' into a model in which all the potential colluding players can offer commitments. This seems more reasonable in many ways than arbitrarily choosing one of the players and giving them this responsibility.

We now turn to the proof of this result which is completely constructive.

4. SOME PRELIMINARY IDEAS.

Our proof combines a number of ideas. We borrow methods from computer science to implement correlated and random outcomes. We then develop a sequential communication mechanism that effectively converts private communication into a public correlating device. We explain each of these methods before we proceed to the proof of the main theorem.

---

[4] Even if we don't know what these allocation rules are, it seems a far easier problem to calculate them for some environment than it does to find all mechanisms which have pure strategy continuation equilibrium.

4.1. **Implementing random outcomes with non-random contracts.** Let B be a set with K elements indexed in some arbitrary way. Let π be a vector of K probabilities that sum to one. Let $\tilde{t}$ be a random variable uniformly distributed on $[0,1]$. The *randomizing function* $\alpha^{\pi}(\cdot, B)$ for mixture π on the set on set B is defined by

$$\alpha^{\pi}\left(\tilde{t}, B\right) = \left\{ b_k : k = \min_{k \in \{1,\dots,K\}} \sum_{l=1}^{k} \pi_l \geq \tilde{t} \right\}. \tag{4.1}$$

This randomizing function takes value b_k with probability π_k. To see how this device will be used, suppose that player i can observe a verifiable random device $\tilde{t}$ which is uniformly distributed on $[0,1]$. Then the contract $\alpha^{\pi}\left(\tilde{t}, A_i\right)$ which maps from the randomizing device into pure actions implements the mixture π on A_i. More broadly, $\alpha^{\pi}\left(\tilde{t}, A\right)$ implements joint action a^k with probability π_k. Let $\alpha_i^{\pi}\left(\tilde{t}, A\right)$ be the projection of α onto A_i. If each player writes a contract based on $\tilde{t}$ that commits them to take action $\alpha_i^{\pi}\left(\tilde{t}, A\right)$, then the set of contracts $\left\{\alpha_1^{\pi}\left(\tilde{t}, A\right), \dots, \alpha_n^{\pi}\left(\tilde{t}, A\right)\right\}$ implements the joint randomization π.

4.2. **A property of uniform distributions.** For any non-negative real number x, $\lfloor x \rfloor$ means the *fractional part* of x (sometimes the terminology is x mod 1). Let $\tilde{x}_1, \dots, \tilde{x}_n$ be a collection of n independent random variables, where each $\tilde{x}_i$ is uniformly distributed on $[0,1]$. For $n \geq 2$, fix $\tilde{x}_i = \overline{x}$ for some i. Then $\lfloor \overline{x} + \sum_{j \neq i} \tilde{x}_j \rfloor$ is a random variable. This random variable turns out to be uniformly distributed on $[0,1]$ independent of $\bar{x}$. [5] Since this argument proves very useful below, we give a simple proof in the Appendix section 8.1..

4.3. **Confirmation Process.** Now we describe how we will turn private messages into public messages. There are two issues here - the first is to create what amounts to a public correlating device. Perhaps as important, each player will convey type information to the other players. Since the player's type report must be the same in each mechanism to which he reports, we have to provide some kind of incentive for players to say the same thing to all players. We do this using a special sequential communication mechanism that we call a 'confirmation process'. Players send messages in the first round, but commit

---

[5] This appears to be conventional wisdom in statistics. The theorem is referred to in Deng and E.Olusegun (1990). A proof that the sum mod 1 of a pair of random variables on $[0,1]$ is uniform as long as at least one of the random variables is uniform is given in Deng, Lin, Wang, and Yuan (1997), Theorem 3.1 (see especially the comment after the theorem).

themselves to react only when other players confirm these messages in the second round.

Recall that in our mechanisms, a player expects to receive a message in $(\Theta \times [0,1])^n$ from each of the other players along with some recommendation about how to deal with these messages. The recommendation is a mapping $\gamma_i : (\Theta \times [0,1])^{n(n-1)} \to A_i$ suggesting to player i how he should convert the messages in $(\Theta \times [0,1])$ into actions. As they aren't mechanisms per se, we will refer to the mappings γ_i as *processes*. To simplify notation a bit, define $S \equiv (\Theta \times [0,1])$. Let τ be a mapping from $S^{n(n-1)} \to S^n$.

Now adopt the following notation: let s_j represent the first round message that player i receives from player j. Let t_j^k represent the second round report that player j makes to i about player k's first round report. Observe that the message t_j^i would be j's report about what i told him in the first round. This message is important in what follows despite the fact that i already knows what report he made to j in the first round. Player i's mechanism will deal with $n-1$ reports like this and convert them into a single vector in S^n.

Definition 2. The mapping $\gamma_i : S^{n(n-1)} \to A_i$ is called a *confirmation process* for player i if there is a mapping $\tau : S^{n(n-1)} \to S^n$ and a mapping $\gamma_i^\tau : S^n \to A_i$ such that $\gamma_i(\tilde{s}) = \gamma_i^\tau(\tau(\tilde{s}))$ for every $\tilde{s} \in S^{n(n-1)}$, and such that the j^{th} component of this transformation τ is given by

$$\tau_j(s_{-i}, t_{-i}) =$$

$$\begin{cases} s' & \text{if } j = i; \left\{\exists! j' : t_{j'}^i \neq t_k^i \equiv s' \forall k \neq j', i\right\} \vee \left\{t_k^i = s' \forall k \neq i\right\} \\ s' & \text{if } j \neq i; \left\{t_k^j = s' \forall k \neq j\right\} \vee \left\{\exists! k \neq i, j : t_k^j \neq s_j = s'\right\} \\ \underline{s} & \text{otherwise.} \end{cases} \tag{4.2}$$

In the expression above, the notation ∃! means there 'exists a unique...', the notation ∨ stands for 'or'.

We explain. A confirmation process breaks the mapping γ_i into two parts. One part, the γ_i^τ looks much like an inscrutable direct mechanism in which the action the player takes is a function of the types of his agents as well as his own type. His 'agents' in this case, are just the other players. However, he doesn't simply ask the others to report their types or choose his own type. Instead, he derives these types from a larger set of messages. This is the 'confirmation' part τ, which reduces the $n(n-1)$ messages sent over two rounds into the n messages required by γ_i^τ.

The number $\tau_j(s_{-i}, t_{-i})$ is the type (and correlating message) in S that i will use for player j in the 'direct mechanism' γ_i^τ. First consider how player i derives the type that he uses for himself - i.e., the value τ_i. The logic is described in the first line of (4.2). He asks the others to tell him what type he reported to them in the first round. If they all agree, or all but one of them agrees he uses whatever type they agree on. Otherwise, he uses some arbitrary type.

For the others, the computation differs only slightly - the logic is described in the second line of (4.2). To find a type (and correlating message) for player j, he asks the players other than j what type j reported to them in the first round. If they all agree, he uses that type. If there is a single dissenting message, he compares the messages that do agree with the type that j reported to him on the first round. If those agree, then he uses that type. Otherwise, he uses an arbitrary type.

We give more structure to the processes γ_i below. For the moment, we focus on a very specific property of confirmation processes. Fix an array of mechanisms for the players. Every such array of mechanisms indexes a subgame of the original game in which players send recommendations and reports. Generally players don't know what recommendations other players make to each other, so they can't predict exactly how their reports in S^n are being converted into actions by any other player.

However, if we fix a set of strategy rules, then each player j believes that the relationship between reports and actions for player i is given by some mapping $\tilde{\gamma}_i^j : S^{n(n-1)} \rightarrow P_i$. This mapping depends on j's type, though we suppress this in the notation to make it a bit simpler. Implicit in this mapping is a presumption that j follows the strategy σ_j when he makes his recommendation to i.

In many subgames, this uncertainty will disappear. For example, if player i uses a mechanism which makes his action independent of players' recommendations. The case we are interested in here is one in which the strategy rules that players are using are such that player j knows what recommendations the others will make to player i. Given some subgame and some array of strategy rules, we say that player j believes that player i is using process γ_j if $\tilde{\gamma}_j^i = \gamma_j$.[6]

Lemma 3. *Suppose $n \geq 4$. Consider any subgame and set of strategy rules such that some player j believes that player i is using a confirmation process. Suppose further that all the players other than j are using strategy rules that involve a consistent revelation strategy. Then whatever the realizations (s_{-j}, t_{-j}) of the others' reports, $\tau_k^i(s_{-i}, t_{-i})$ is*

---

[6] Equality here means that the mixture on the left is a degenerate mixture with support that coincides with the values on the right hand side.

independent of what j reports if $k \neq j$, while there are reports that j can send to i such that $\tau_j^i(s_{-j}, t_{-j})$ takes any value in S.

The important thing about this Lemma is that when j considers what messages to send to player i, he is strategically in exactly the same position he would be in participating in a standard direct mechanism. All he can affect is his type report (and correlating message) to player i.

We give a full proof here, but the logic is straightforward. There are 4 or more players, so i expects messages from at least three players. By (4.2) i will ignore a message from one player unless it agrees with all the others. When player j is considering what value player i will use for the type and correlating message of some player k, he expects the other players to report a common value in S to i as they are all using consistent reporting strategies. As a consequence, he expects his own message will be ignored. Observe that this logic applies to both the first and second round message. Also note that player i is also expected to use a consistent reporting strategy here, so k could have the value i in this paragraph.

On the other hand, if j considers what message i will use for him, the logic is different. As the others are using consistent reporting strategies, he simply needs to send the same message in S to each of the other players to ensure that i uses that message. None of this argument works if n is three or less because i cannot tell which of two different messages he should ignore.

4.4. Consensus Mechanisms.

The idea that principals should ask their agents for recommendations about how to process information is due to Yamashita (2010). His idea was to have the principal commit himself to carry out the recommendations of the agents provided an outright majority of the agents make the same recommendation. We simply adapt this idea here. In our context there may or may not be agents, so players ask other players for recommendations. If the principal's mechanism commits him to carry out the recommendation when all the other players, or all but one of the other players agree, then we say that the principal's mechanisms is a *consensus mechanism.*

Formally, a mechanism $r_i : (\Gamma_i)^{(n-1)} \times S^{n(n-1)} \to A_i$ is a consensus mechanism if

$$r_i(\gamma_{-i}, s_{-i}) =$$

$$\begin{cases} \gamma'(s_{-i}) & \text{if } \{\exists! j : \gamma_j \neq \gamma_k \equiv \gamma' \forall k \neq j\} \vee \{\gamma_k = \gamma_j \equiv \gamma' \forall j, k\} \\ \overline{a}_i & \text{otherwise.} \end{cases}$$

Now the proof of our folk theorem can be done constructively. On our equilibrium path all players will offer a consensus mechanism independent of their type. If all players do this, then each of them will recommend a confirmation process to each of the other players. The details of the confirmation process will depend on the allocation rule we are trying to support. If some player deviates and offers something other than a consensus mechanism, then the other players will recommend to each other a confirmation process than penalizes the deviator.

It should mostly apparent why this construction will work. Players can see whether or not everyone has offered a confirmation process after mechanisms are announced. They then believe that they know what recommendations the others will make. The nature of a consensus process is such that unilateral disagreement is ignored, so going along with the majority is at least a weak best reply. The confirmation process has been constructed specifically so that all players can accomplish when they participate is to lie about their type. As long as the confirmation process implements an incentive compatible allocation rule, there is no incentive to do this. The only real complication in the proof is to show that it is an equilibrium for players to send correlating messages that correctly implement randomized outcomes.

5. THE PROOF OF THE MAIN THEOREM

Proof. The proof is constructive. Let $q(\theta)$ be the randomization that is to be supported when types are θ. Since the allocation rule is individually rational, there is a collection of punishments that ensure participation by each player. Let $\{p_i(\theta_{-i})\}_{i\in N}$ be the type contingent randomization that is to be carried out by the players other than i when i is being punished.

We first describe the recommendations we want players to make.

Let τ be a confirmation process with message space $S^n = (\Theta \times [0,1])^n$. Write (θ, x) as a typical element of S. The function $\tau_j(s_{-i}, t_{-i}) \in \Theta \times [0,1]$, so write $\tau_j(s_{-i}, t_{-i}) = \left\{\tau_j^\theta(s_{-i}, t), \tau_j^x(s, t_{-i})\right\}$. The equilibrium path recommendation by other players to player i is given by

$$\gamma_i(s_{-i}, t_{-i}) = \alpha_i^{q\left(\tau^\theta(s_{-i},t_{-i})\right)}\left(\lfloor\sum_{j\in N}\tau_j^x(s_{-i}, t_{-i})\rfloor, A\right) \tag{5.1}$$

where α_i^q is the projection of the randomizing function for mixture $q\left(\tau^\theta(s_{-i}, t_{-i})\right)$ on the set A onto the set A_i. The randomizing function is defined by (4.1) above.

When player k unilaterally deviates in the mechanism design stage and offers something other than a recommendation mechanism, the non-deviators will recommend

$$\gamma_i^k\left(s_{-ik},t_{-ik}\right)=\alpha_i^{p_i\left(\tau^\theta\left(s_{-ik},t_{-ik}\right)\right)}\left(\lfloor\sum_{j\neq k}\tau_j^x\left(s_{-ik},t_{-ik}\right)\rfloor, A_{-k}\right) \tag{5.2}$$

to each non-deviating player i, where $\left(s_{-ik},t_{-ik}\right)$ is an array of messages from the other *non-deviating* players. In words, the non-deviators will recommend to each other the (projection of the) randomizing function associated with the punishment.

In any information set in which all players have offered a consensus mechanism, player i should recommend γ_j to each other player j, truthfully report to each player $k\neq j$ the message received from player j, truthfully report his type to every other player, and send every other player a correlating message x drawn uniformly from $[0,1]$. In any history in which a single player, say player k, has deviated and offered some mechanism other than a consensus mechanism, player i should recommend the punishment mechanism γ_j^k to each player $j\neq k$, truthfully report the private message received from each player $j\neq k$ to each player $j'\neq k,j$, send the same correlating message s' to each of the other players where s' is chosen using a uniform distribution on $[0,1]$, and report his type truthfully to each of the players other than k. Any action unspecified here can be chosen arbitrarily.

Now we proceed to prove that the strategies specified constitute a Bayesian equilibrium. First, it is immediately a best reply for each player to offer a consensus mechanism. If he does that, he should expect the allocation rule $q\left(\theta\right)$. If he deviates, he should expect the others to implement the punishment $p_i\left(\theta\right)$. Since the allocation rule satisfies (1.2), this can't increase his payoff.[7]

On the equilibrium path, all players offer consensus mechanisms, and each player recommends a confirmation processγ_j as given by (5.1). We have already explained that for any confirmation process, it is a best reply for each player to report the same type and correlating message to each of the other players provided they believe that the others are doing the same. For the correlating message, the others are expected to send a correlating message that is uniformly distributed on $[0,1]$. As we have explained in Remark 5 above, this implies that $\lfloor\sum_{j\in N}\tau_j^x\rfloor$ has a uniform distribution independent of what signal x_i i chooses to send. Then for each θ_{-i} and each report θ_i' and signal $\tilde{x}_i$ that i chooses

---

[7] Notice that because we are only interested in Bayesian equilibrium at this point, this particular argument works even if there are only three players in the game. We still need the fourth player to support equilibrium in the confirmation process.

to send to the others on the first round

$$\alpha_k^{q\left(\tau^{\theta}\right)}\left(\lfloor\sum_{j\in N}\tau_j^x\rfloor, A\right) =$$

$$\alpha_k^{q\left(\theta_i',\theta_{-i}\right)}\left(\tilde{x}, A\right)$$

where $\tilde{x}$ has a uniform distribution on $[0,1]$. Since this rule implements the incentive compatible rule q, player i has no incentive to misrepresent his type. It is also a best reply for player i to choose a signal uniformly from $[0,1]$.

The other direction for this theorem is straightforward. If an allocation rule is supported as a Bayesian equilibrium, it is obviously incentive compatible. It is individually rational because a player can always deviate to a mechanism that makes his action depend only on his own correlating message. The payoff associated with this deviation is given by the right hand side of (1.2), where the punishment is whatever outcome function is supported by continuation play when the deviator uses this mechanism. Since the left hand side of (1.2) represents the payoff before a deviation, inequality (1.2) must hold whenever the allocation rule $q(\theta)$ is supported as a Bayesian equilibrium. ☐

6. REMARKS.

The approach above shares many of the methods of the literature on communication in games(Gerardi 2004),(Forges 1986),(Barany 1992). (Gerardi 2004), for example, uses the majority rule approach to ensure that players all send the 'correct' message in his communication protocols. This is exactly the idea behind a consensus mechanism. He also uses the randomization idea in (4.1), albeit restricted to two players.[8] The important difference between our paper and all this literature is the fact that we are doing mechanism design - players can make commitments based on messages. So the allocation rules we support aren't typically communication equilibrium (or correlated equilibrium with complete information). The important difference in the design of communication protocols is that when trying to support communication equilibrium, it is critical that players only learn what they themselves are supposed to do after they receive and send all their messages. In our context, this doesn't matter because once messages

---

[8] He has two players publicly announce numbers in the interval $[0,1]$ then uses the fractional part as a public correlating device. As there are only two players, he doesn't need our Remark 5. In our model, there are no public messages at all, beyond the mechanisms that players announce at the beginning of the game.

have been exchanged, every player's action is committed through the mechanism that he offered at the beginning of the game. It doesn't matter whether or not he learns what the other players are committed to.

There is a literature on mechanism design in communication networks ((J. Renault and Tomala 2010) or (Renou and Tomala 2010)) which considers sequential communication schemes like the one we described in Section 4.3. In this literature, a centralized mechanism designer can communicate with only a subset of all the agents. However, the agents can communicate among themselves according to some exogenously fixed communication protocol. The papers cited above provide communication protocols which allow agents to communicate their type information secretly to the principal. The essence of their result is to show that, provided the communications network is right, there is a way for agents to encode their own information along with the information they have received from others, and pass it along in such a way that only the the mechanism designer can decode it.

In order to ensure that players pass along encoded information truthfully, their papers use a method that resembles our confirmation process. A protocol that transmits player 1's type (assumed here to be a positive number) to the mechanism designer is repeated, say, 3 times. Player 1 chooses at random one of the three repetitions and transmits his type on that repetition. On each of the other two repetitions he transmits the number 0 as his type. If the mechanism designer decodes 2 zeros and one positive number, he responds as if the type is a positive number. If he decodes any other sequence, he implements a punishment. The purpose of this is to ensure that the other players transmit messages from player 1 truthfully. They don't know which of the three messages from player 1 contain his type report. So they have a 1 in three chance of changing the outcome in a way that they might like, and a 2/3 chance of inducing the punishment when they lie. Assuming that there is a punishment that is strictly worse than any outcome the mechanism designer might otherwise implement, no matter the type of any other player, then repeating the protocol enough times will ensure that players other than player 1 transmit messages truthfully.

The details of the argument differ, but the spirit is the same as the confirmation process - if other players are hearing the same message that you are, then there are sometimes ways to to check whether they are transmitting the information truthfully. Our communication process is structured to do this, so we don't need a 'worst outcome' that a mechanism designer can use to enforce truthtelling. In our framework, deviating messages are simply ignored. Of course, the context of our result is quite different since we don't have a mechanism designer in the first place - we deal with decentralized competition.

Nonetheless, the method they describe illustrates how the results presented here can probably be extended to games with fewer than four players. Our communication mechanism requires agreement among all but one of the players who are participating in a mechanism. If there are only two players, and the messages they send are different, then the player who is interpreting them does not know which message is the correct one, and which is a deviation. Each of our players has to have at least three others sending him messages for our method to work. The method above illustrates how a player might detect deviations with messages from only two players provided the sequential communication mechanism goes on for long enough. It may also be possible to extend our results if there is a public correlating device using methods like those in (Forges and Vida 2011) who show that communications equilibrium outcomes can be implemented with long cheap talk in games with only two players using a public device.

The use of a second round of communication to provide a mechanism designer with additional information is similar to the argument in (Mezzetti 2004), who shows how a mechanism designer can improve outcomes by using a second round of information in which players provide information about their values. When players' payoffs are interdependent, each player's value contains information about everyone else's type in much the same way the first round reports do here. Of course, the method we use to get players to reveal this information is quite different than it is in that reference.

Folk theorem like results for competing mechanism games have been provided by (Tennenholtz 2004), (A.T. Kalai and Samet 2010) and (Peters and Szentes 2008). The essential difference between these papers and our result here is that they assume contracts can condition directly on the contracts of other players. So players do not directly convey type information or make recommendations to one another. As mentioned above, our primary contribution here is to show how to structure the sequential communication mechanism among players (including agents) so that private communications can be used in the same way as a public correlating device (which not only makes the type public, but also ensures player send the same type reports to all other players). At a more mundane level, the first two papers above deal with games of complete information. The paper by (Peters and Szentes 2008)deals with incomplete information games but imposes two restrictions which play no role here. First, it restricts to non-random contracts and focuses on pure strategy equilibria. Second, it does not allow any direct communication between players - all type information is conveyed through information revealed by contract offers themselves. For this reason the set of outcome functions that can be supported as equilibrium is a strict subset of the set that can be supported here.[9]

The paper by (Peters 2010) provides a similar characterization result to the one we give here. In fact it borrows the correlating device we have, in turn, borrowed here. However, the point of that paper is quite different. It revisits the question in (Epstein and Peters 1999) and provides a modified set of direct mechanisms that can be used to mimic equilibrium outcomes in any competing mechanism game - effectively providing a revelation principle for competing mechanisms. Since every equilibrium in mechanisms supports an outcome function that can be supported as an equilibrium in the direct mechanisms provided by (Peters 2010), the same must be true for all the outcome functions supportable here. This is why the direct mechanisms in that paper have to provide the same folk theorem as the one presented here for environments with four or more players.

At the other extreme, the model of competing mechanisms that is most often discussed in the literature imposes strong restrictions on the environment and on players' ability to contract. For example, a standard approach is to distinguish a number of uninformed mechanism designers and have them compete to influence a number of privately informed agents. Principals are typically not allowed to communicate at all beyond telling agents what their mechanisms are. Agents are viewed to have a very limited set of actions and no ability at all to commit. For example, in a typical delegated common agency (e.g. (Martimort and Stole 2008)), agents simply decide how much, if any, output they should produce on behalf of each of the principals. Principals choose output contingent transfers.[10]

We use finite action spaces here, so outputs would have to be constrained to a finite set. Transfers in turn would have to be chosen from a finite set, so that the set of output contingent transfers remains finite. However, this difference is inessential. The more important distinction is that limiting communication and commitment ability of agents

---

[9] Note that it is only for games with four or more players that the papers are comparable.

[10] Just to be explicit about the relationship, in (Yamashita 2010), agents take no actions at all (i.e., there is nothing for them to commit to) and simply transmit recommendations and type reports to principals. Agents are privately informed, while principals have no information. Those are all special restrictions on the environment which we don't impose. He does prevent principals from communicating with one another.

will add additional constraints to (1.1) and (1.2). For example, if an outcome function is supportable as an equilibrium in recommendation mechanisms, it will still satisfy (1.1) and (1.2). However, the agents' inability to commit their participation decisions will impose additional restrictions on the punishment. Conversely additional restrictions will have to be added to the participation decisions involved in a punishment when players cannot commit their participation decisions with contracts.

7. Refinements

Our theorem uses Bayesian equilibrium as a solution concept. Refinements, like restrictions on commitments, are going to impose additional restrictions on supportable allocations. For example, if we want to impose a refinement like sequential equilibrium, then allocations supportable with consensus mechanisms will still satisfy (1.2). However, the punishment mechanism will also be incentive compatible. Conversely, a sufficient condition for an allocation to be supportable as a sequential equilibrium will be the existence of an incentive compatible punishment. The nature of these additional restrictions depends on the environment, the refinement, and whatever additional restrictions are imposed on contracting, as we discussed above. We illustrate here. However, a useful discussion of these issues goes well beyond the scope of this paper..

To understand the biggest problem with refinements, consider a game with complete information. There are four players (simply so that the assumptions of our theorem above are satisfied). Suppose that player 1 has three possible actions, $\{a, b, c\}$. None of the other players controls any actions at all. Player 1 offers a mechanism, and the solution concept requires that after seeing the mechanism, continuation play constitutes a Nash equilibrium (subgame perfection). Obviously, player 1 simply chooses his favorite action in any Bayesian equilibrium. However, player 1 could deviate and offer a mechanism which invites players 2 and 3 to send a message in $[0, 1]$. He commits to translate the messages m_2 and m_3 into actions the following way:

$$\gamma(m_2, m_3) = \begin{cases} a & \text{if } m_2 < m_3 < m_2 + \frac{1}{2}, \\ b & \text{if } m_2 = m_3 \text{ or } m_3 = m_2 + \frac{1}{2}, \\ c & \text{otherwise.} \end{cases}$$

Now imagine payoffs for player 2 are$u(a) = -1$, $u(b) = 0$, and $u(c) = 1$. Player 3's payoff is $-u$. This is simply the Sion Wolfe Sion and Wolfe (1957) example of a game that has no equilibrium in either pure or mixed strategies. This is a feasible mechanism in our framework, and a reasonable looking mechanism in any framework. So in this simple setting, there can be no subgame perfect equilibrium to the mechanism game unless mechanisms like the one above are ruled out.

There are (at least) two ways to get around this problem - restrict the set of mechanisms that players are allowed to use, or use refinements other than sequential equilibrium. For example, one could restrict players to mechanisms that involve only finite message spaces. As long as actions and types are both finite, it ought to be possible to show that the overall mechanism design game must have a sequential equilibrium. Alternatively, we could require that equilibrium satisfy the weaker property that players' strategies be sequentially rationalizable in every information set.

As mentioned above, a characterization under either of these approaches would require that we strengthen (1.2) by adding constraints requiring the punishment to be either incentive compatible, or incentive 'rationalizable'.

Yet, with respect to providing a 'useful' characterization, there is another issue. Consider an allocation rule satisfying (1.1) and (1.2), and suppose that in addition that there is some belief about the action the deviator will take when he deviates that makes the punishment p^i incentive compatible. It doesn't follow from this that the allocation rule can be supported as a sequential equilibrium. The reason is that the deviator can announce commitments - even commitments that depend on type reports from the players who are supposed to be punishing him. The non-deviators see what these commitments are, and can make recommendations that depend on these commitments. A 'characterization' of the kind above is then going to require a constraint like (1.2) for every possible commitment the deviating player can make. We could try to add all these constraints to (1.2). However, this would generally be little more than a tautological restatement of what an equilibrium is.[11]

If refinements are important, then the issue is not to provide a general characterization, as we have done here, but to find restrictions on mechanisms and the environment in which a useful characterization is possible with refined equilibrium. As an example of how this might

---

[11] This is the kind of 'characterization' provided in Yamashita's paper. He gives a value for a deviating principle which is described as the lowest payoff the principle can attain in any continuation equilibrium associated with any deviation. So it is more difficult (since one needs to find all possible continuation equilibria) to find these values than it is to check whether there is a sequential equilibrium supporting a particular outcome.

work, consider a delegated common agency with many principals and many agents. Index the principals as players $1, \ldots, m$ and the agents as players $m+1, \ldots, n$.

As an illustration, consider a delegated agency problem with many agents. Index the principals as players $1, \ldots, m$ and the agents as players $m+1, \ldots, n$. We make the usual assumption that principals have no private information, while agents are privately informed. Suppose each principal wants to buy some input from one of the agents. Each principal can then choose to make a transfer from a finite set of transfers to one or more of the agents. We assume this set contains a zero transfer and no negative transfers. The agents choose a quantity of the input to produce for each principal (again restricted to a finite set of non-negative quantities to fit the formalism here). This set is assumed to include an action that involves producing no output for any or all principals. We assume principals' payoffs are given by $u(x, \theta) - t$ where x is the total quantity produced for him, $u(x, \theta) \geq 0$, t is the sum of the transfers the principal makes, and θ is the vector of agents' types. Agents' payoff are assumed to be $t - c(x, \theta)$, where again t is the total transfer they receive, x is the total output they produce and θ is the vector of agent types. Then the no-externalities condition from (Peters 2003) holds, and conditional on agents' actions, agents' preferences over any menu of one principal's actions are independent of the actions of any other principal. This assumption about preferences along with the assumption that agents can offer consensus contracts together make it possible to give a useful characterization of supportable outcome functions.[12]

The set of contracts available to the players is the same as we have described above. As our characterization result applies only in this environment, refer to this particular competing mechanism game as the 'delegated agency game'.

Theorem 4. *Suppose there are 4 or more players. Let $q(\theta)$ be an incentive compatible allocation rule. The q can be supported as a sequential equilibrium in the delegated agency game if and only if each player's payoff is non-negative.*

---

[12] In this formulation the mechanism that a principal offers is a mapping from the types of all the agents into his own payment. It might seem more natural to model the principal's contract as a mapping from agent types to transfers and quantities. As we have illustrated above, consensus contracts involving confirmation processes allow players to correlate their actions, so the two approaches will give the same set of supportable allocations. The approach we use allows a very simple comparison with Theorem 1.

Most of the details of the proof are straightforward applications of the reasoning given above. Consider the sufficiency argument. Reporting the same type and correlating message to each of the others, and randomizing over the correlating message are sequentially rational no matter what players believe about the type of a deviator, provided there are at least 3 non-deviators left to support the majority arguments used in the confirmation process. So as in the proof of Theorem 1, suppose that when all players offer a consensus mechanism, player unanimously recommend the direct mechanisms that implement q (see the construction in the proof of Theorem 1). Sequential equilibrium becomes an issue after a deviation. However, if some player deviates and offers something other than a confirmation process, there are many mechanisms that implement a zero payoff for the deviating player. Sequential rationality simply requires the existence of some mechanism which is incentive compatible no matter what the deviator offers and which implements zero trade with the deviator. The non-deviators are going to recommend this punishment to all the non-deviators who are still committed to consensus mechanisms. Going along with the majority is sequentially rational independent of beliefs by the nature of a consensus mechanism. There are clearly many mechanisms that are incentive compatible (any constant mechanism for example). The necessity argument is immediate because every player can choose a mechanism that yields non-negative payoff.

It should be apparent that we can go even further in this simple game. It is uncommon in common agency to allow agents to commit themselves. In existing competing mechanism games it is normally assumed that principals can't communicate with each other. It should be apparent that Theorem 4 will be true in a modified version of the game in which only agents send messages (recommendations in this case) provided there are at least four agents. Four agents are needed because they are making commitments and there have to be three agents left after one deviates to enforce a punishment.

Notice that theorem 4 allows correlation in types and interdependent preferences.

Finally, we can extend this kind of reasoning even further by assuming that agents can't make commitments. To deal with this we need to make the output choice part of the principal's mechanism. Then if there are three agents, a theorem like Theorem 4 will be true if there is an outcome function for the non-deviators which is incentive compatible

CONCLUSION

Our basic contribution is to complement the Yamashita theorem by characterizing the set of outcome functions supportable as Bayesian Nash equilibrium in the competing mechanism game. Whether or not there is a sensible equilibrium refinement for competing mechanism games that will narrow the of supportable outcomes is an open question.

8. Appendix

8.1. **Uniform Distributions and independence.**

Remark 5. $\lfloor \bar{x} + \sum_{j \neq i} \tilde{x}_j \rfloor$ is uniformly distributed on $[0,1]$ independently of $\bar{x}$ provides each $\tilde{x}_j$ is uniformly distributed on $[0,1]$.

Proof. Suppose that $n = 2$. Then $\sum_{j \neq i} \tilde{x}_j = \tilde{x}_j$, and $\lfloor \bar{x} + \tilde{x}_j \rfloor$ is obviously uniform. Let both $\tilde{x}_1$ and $\tilde{x}_2$ be uniform on $[0,1]$. Then the probability density function of $\tilde{z} = \tilde{x}_1 + \tilde{x}_2$ is[13]

$$f(z) = \begin{cases} z & 0 \leq z \leq 1 \\ 2 - x & \text{otherwise.} \end{cases}$$

The probability that $\lfloor \tilde{z} \rfloor \leq w$ is then given by

$$\int_0^w z dz + \int_1^{1+w} (2 - z)\, dz = w.$$

So $\lfloor \tilde{x}_1 + \tilde{x}_2 \rfloor$ is uniformly distributed. So when $n = 3$, $\lfloor \bar{x} + \sum_{j \neq i} \tilde{x}_j \rfloor$ is uniformly distributed. Then the argument follows by induction. If for $n - 1$ players $\lfloor \bar{x} + \sum_{k \neq j} \tilde{x}_k \rfloor$ is uniformly distributed, then for n players

$$\lfloor \bar{x} + \tilde{x}_j + \sum_{k \neq i,j} \tilde{x}_k \rfloor =$$

$$\lfloor \tilde{x}_j + \lfloor \bar{x} + \sum_{k \neq i,j} \tilde{x}_k \rfloor \rfloor$$

and uniformity follows from the result for $n = 3$. □

---

[13] (Hall 1927).

8.2. **Proof of Lemma 3.**

Lemma 6. *Suppose $n \geq 4$. Consider any subgame and set of strategy rules such that some player j believes that player i is using a confirmation process. Suppose further that all the players other than j are using strategy rules that involve a consistent revelation strategy. Then whatever the realizations (s_{-j}, t_{-j}) of the others' reports, $\tau_k^i(s_{-i}, t_{-i})$ is independent of what j reports if $k \neq j$, while there are reports that j can send to i such that $\tau_j^i(s_{-j}, t_{-j})$ takes any value in S.*

Proof. Fix the first round reports s_{-j} of the players other than j. We write in the obvious way s_{-jk} for the subvector consisting of reports in s_{-j} by players other than k. Suppose that j's strategy is consistent and he sends the message s' to each of the other players in the first round. Then since every other player is using a consistent strategy, the value that i uses for player j will be based on first round messages (s', s_{-ij}), second round message (s', s_{-kj}) from each player $k \neq j$ since each such player is using a consistent reporting strategy, and second round message s_{-j} from player j which doesn't depend on s'. Since the first round message from player j agrees with the second round reports of each of the other players, we conclude by (4.2) that

$$\tau_j^i\left((s', s_{-ij}), \prod_{k \neq i,j} (s', s_{-kj}), s_{-j}\right) = s'.$$

Notice that this verifies the last part of the theorem - j can induce any value for τ_j^i in S.

Player j can deviate from this consistent strategy by sending different messages to the other players on the first round. He could also send different messages to i on the second round, but τ_j^i doesn't depend on i's second round messages, so we defer discussion of this second kind of deviation. Let s_k be the message he sends to player k on the first round, and $\tilde{s}_{-j}$ the vector of $n-1$ messages he sends to i on the second round. In this case there are two possibilities. If the players $k \neq j$ all report $s'_k = s'$, or if all but one of the others reports $s' = s'_i$ then by (4.2),

$$\tau_j^i\left((s'_i, s_{-ij}), \prod_{k \neq i,j} (s'_k, s_{-kj}), \tilde{s}_{-j}\right) = s',$$

which is an outcome j could have obtained by using a consistent reporting strategy and reporting s' in the first round to everyone. Otherwise,

$$\tau_j^i\left((s'_i, s_{-ij}), \prod_{k \neq i,j} (s'_k, s_{-kj}), \tilde{s}_{-j}\right) = \underline{s},$$

which is an outcome he could also accomplish with a consistent strategy by sending the message $\underline{s}$ to each player then reporting accurately to i in the second round.

To complete the proof of the theorem, observe that since player k is using a consistent reporting strategy, he will make the same first round report s_k to each of the other players. With the possible exception of player j, each of the others will then report s_k to player i. Since at least two second round reports will agree with k's first round report, we have

$$\tau_k^i \left(\left(s_i', s_{-ij} \right), \prod_{k' \neq i,j} \left(s_{k'}', s_{-k'j} \right), \tilde{s}_{-j} \right) = s_k$$

independent of $\tilde{s}_{-j}$. □

References

Andrea Attar, D. Majumdar, G. P., and N. Portiero (2008): "Common Agency Games: Separable Preferences and Indifference," *Mathematical Social Science*, 56(1), 75–95.

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

Barany, I. (1992): "Fair Distribution Protocols or How the Players Replace Fortune," *Mathematics of Operations Research*, 17, 327–340.

Che, Y., and J. Kim (2006): "Robustly Collusion Proof Implementation," *Econometrica*, 74, 1063–1107.

Deng, L.-Y., and G. E.Olusegun (1990): "Generation of Uniform Variates from Several Nearly Uniformly Distributed Variables," *Communications in Statistics-Simulation and Computation*, 19(1).

Deng, L.-Y., D. K. Lin, J. Wang, and Y. Yuan (1997): "Statistical Justification of Combination Generators," *Statistica Sinica*, 7, 993–1003.

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

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

Forges, F. (1990): "Universal Mechanisms," *Econometrica*, 58(6), 1341–1364.

Forges, F., and P. Vida (2011): "Implementation of Communication Equilibria by Correlated Cheap Talk: The Two Player case," Discussion paper, CESifo Working Paper No. 3360.

Gerardi, D. (2004): "Unmediated communication in games with complete and incomplete information," *Journal of Economic Theory*, 114(1), 104–131.

Gerardi, D., and L. Yariv (2007): "Deliberative Voting," *Journal of Economic Theory*, 134, 317–338.

Hall, P. (1927): "The Distribution of Means for Samples of Size N drawn from a Population in which the Variate takes Values between 0 and 1, All Such Values Being Equally Probable," *Biometrika*, 19(3), 24–245.

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

J. RENAULT, L. R., AND T. TOMALA (2010): "Secret Information Transmission," Discussion paper, HEC Paris.

LAFFONT, J., AND D. MARTIMORT (1997): "Collusion under Asymmetric Information," *Econometrica*, 29, 280–305.

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

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

——— (2008): "Market Participation under Delegated and Intrinsic Common Agency," mimeo, University of Chicago and Toulouse School of Economics.

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

MCAFEE, P., AND J. MCMILLAN (1992): "Bidding Rings," 82, 579–599.

MEZZETTI, C. (2004): "Mechanism Design with Interdependent Valuations: Efficiency," *Econometrica*, 72, 1617–1626.

PAVAN, A., AND G. CALZOLARI (2009): "Sequential Contracting with Multiple Principals," *Journal of Economic Theory*, 50, 503–531.

PECK, J. (1994): "Competition in Transactions Mechanisms: The Emergence of Competition," Ohio State University mimeo.

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

——— (2003): "Negotiation versus Take It or Leave It in Common Agency," *Journal of Economic Theory*, 111, 88–109.

——— (2010): "On the Revelation Principle and Reciprocal Mechanisms in Competing Mechanism Games," Discussion paper, University of British Columbia.

PETERS, M., AND B. SZENTES (2008): "Definable and Contractible Contracts," Working paper, University of British Columbia.

RENOU, L., AND T. TOMALA (2010): "Mechanism Design and Communication Networks," Discussion paper, HEC Paris no. 926/2010.

SION, M., AND P. WOLFE (1957): *On a Game without a Value*Princeton:Princeton University Press.

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.