A FOLK THEOREM FOR COMPETING MECHANISMS

MICHAEL PETERS AND CHRISTIAN TRONCOSO VALVERDE
UNIVERSITY OF BRITISH COLUMBIA

Abstract. We prove a folk theorem for games in which mechanism designers compete in mechanisms and in which there are at least 4 players. All allocations supportable by a centralized mechanism designer, including allocations involving correlated actions (and correlated punishments) can be supported as Bayesian equilibrium outcomes in the competing mechanism game.

In a competing mechanism game, multiple principles 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 equilibria 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 possiblity was intially mentioned in (McAfee 1993), however, examples that illustrated this possiblity in common agency games were only offered later in (Peck 1994, Martimort and Stole 1998, Epstein and Peters 1999).

A characterization of the *contracts* needed to support all equilibrium outcomes was provided in (Epstein and Peters 1999). For the special case of common agency, simpler characterizations, again of the set of contracts needed to understand all equilibrium allocations, were provided 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 equilibria have been provided in special environments by(Han 2006), and (Andrea Attar and Portiero 2008).

Only recently has the literature focussed on trying to characterize the set of *allocations* that can be supported by equilibrium. The first paper to do this is (Yamashita 2007), who borrows an idea from social choice and uses *recommendation mechanisms* to support a large set of pure allocations as equilibria. The essential idea is that principals will carry out a action or implement a direct mechanism if a majority of the agents he communicates with tell him to. Since disagreeing with

Very Preliminary April 15, 2009.

the majority is never a strictly best reply, he uses this to show that a large set of allocations can be supported. Despite a restriction to non-random contracts and pure strategy equilibrium, the set of outcomes that can be supported is perhaps unreasonably large.

Our purpose in this paper is to provide a complete folk theorem for a richer environment than the one considered in Yamashita. Beyond an interest in randomization, we allow principals as well as agents to have private information. In our story, there need not be any agents at all, since principals are allowed to communicate directly with one another. Our purpose is to illustrate how every allocations that is supportable by a centralized mechanism designer can be supported as an equilibrium in competing mechanisms.

A notable feature of our contracts is that, though we allow arbitrary message spaces, and unlimited commitment, we do not allow principals to write random contracts - all contracts map messages into pure actions. Nonetheless, we are able to support not only random allocations, but correlated allocations. We borrow methods from the computer science literature ((A.T. Kalai and Samet 2007)) that allow us to convert random messages into arbitrary randomizations over actions. We then extend the Yamashita idea to show how a communication protocol can be used to implement arbitrary mixtures over joint action by having players send private messages over two rounds of communication.

The formalism that we use to characterize allocations is similar to, but much simpler than the formulation in (Peters and Szentes 2008) who considered contracts that conditioned on other contracts. There is an important difference between contracts that condition on other contracts and contracts that condition on arbitrary messages. Contracts determine actions while messages do not. So to the extent that a contract is expected to affect some other players action, it has properties that resemble a costly message. In our approach messages have no payoff consequences beyond those built into principals contracts.

1. Fundamentals

There are $n \geq 5$ 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 type θ_i drawn from a finite set Θ. Payoffs are given by $u_i : A \times \Theta^n \rightarrow \mathbb{R}$. 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 P occurs, where the set A is indexed in some arbitrary fashion.

Let B be a set with K elements indexed in some arbitrary way. Let π be a vector of K probabilities that sums to one. Let $\tilde{t}$ be a random variable uniformly distributed on $[0, 1]$. The function

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

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 $\left\{\alpha_1^{\pi}\left(\tilde{t}, A\right), \ldots, \alpha_n^{\pi}\left(\tilde{t}, A\right)\right\}$ implements the joint randomization π.

For any non-negative real number x, $\lfloor x \rfloor$ means the *fractional part* of x. Let $\tilde{x}_1, \ldots, \tilde{x}_n$ be a collection of n independent random variables, where each $\tilde{x}_i$ is uniformly distributed on $[0, 1]$.

Remark 1. For any $\tilde{x}_i$,

$$\Pr\left\{ \lfloor \tilde{x}_i + \sum_{j \neq i} \tilde{x}_k \rfloor \leq x \right\} =$$

$$\Pr\left\{ \sum_{j \neq i} \tilde{x}_k \leq x - \tilde{x}_i \right\} + \sum_{i=1}^{n-2} \Pr\left\{ i - \tilde{x}_1 \leq \sum_{j \neq i} \tilde{x}_k \leq i + x - \tilde{x}_i \right\} +$$

$$+ \Pr\left\{ n - 1 - \tilde{x}_1 \leq \sum_{j \neq i} \tilde{x}_k \leq n - 1 \right\}$$

$$= \frac{(x - \tilde{x}_i)}{n-1} + \frac{(n-2)\, x}{n-1} + \frac{\tilde{x}_i}{n-1} = x.$$

What this simple calculation illustrates is that the random variable $\lfloor \sum_i \tilde{x}_i \rfloor$ has a uniform distribution on $[0, 1]$ independent of $\tilde{x}_i$. We use this property later in our proof.

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

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

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

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

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

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

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 the 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. As always, the pair $\{\gamma_i, M_i\}$ is a mechanism for player i with $\gamma_i : M_i \to A_i$. We presume in what follows that message spaces are always measurable. We don't need random mechanisms for the folk theorem we are about to prove, which is why we restrict the mapping to have image in A_i instead of P_i. However, we are going to exploit the fact that messages are sent sequentially. In particular, we imagine the message space M_i is a cross product space $M_{i0} \times M_{i1} \times \cdots \times M_{ik}$ and that messages in M_{i0} are sent first. After all the players have observed their first round messages from M_{i0} from all the other players, they send messages in M_{i1}, which can be conditional on the first round messages they received, and so on. This process continues until all required messages have been sent. We refer to mechanisms like this as *sequential communcations mechanisms.* [1]

[1]In the usual single principal mechanism design problem, multiple rounds of messages between agents are handled by letting messages be contingent plans that

We prove our folk theorem using two basic ideas. Following (Yamashita 2007), we start by having all players offer recommendation mechanisms. To support desired allocations, we then have all players recommend a relatively simple sequential communication mechanism involving only two rounds of communication. In the first round, players send each other recommedations, type reports and signals. In the second round, players tell each other what they heard from the other players.

A recommendation mechanism involves a message space consisting of two parts, a function space R^i and a more standard measurable message space M^i. The space R^i is the set of all measurable mappings from $(M^i)^{(n-1)}$ into A_i. The mechanism $(\gamma_i, \{R^i, M^i\})$ is called a *recommendation mechanism* if

$$\gamma_i\left(r_{-i}, m_{-i}\right) = \begin{cases} r'\left(m_{-i}\right) & \text{if } \{\exists! j : r_j \neq r_k \equiv r' \forall k \neq j\} \vee \{r_k = r_j \equiv r' \forall j, k\} \\ \overline{a}_i & \text{otherwise.} \end{cases}$$

In words, if all but possibly one of the others make the same recommendation about how to translate the messages into actions, then the recommendation mechanism commits player i to carry out that common recommendation. Otherwise, the mechanism takes an arbitrary action.

The set R^i is a set of mechanisms. We will have all players make the same recommendation on the equilibrium path. This recommendation is a special kind of sequential communication mechanism based on something called a *confirmation process.*

A confirmation process for a set of $n-1$ players with message space S is a function that compresses $(n-1)^3$ messages from S into n elements of S. The process itself is simply a function from $S^{(n-1)^3}$ into S^n. However, the confirmation process is part of a sequential communcations mechanism in which each player sends a message from S to each of the other players in the first round, then, in the second round, tells each of the other players what messages he heard in the first round. In the notation above, $M^i = S^n$. To emphasize the sequential nature of this process we adopt the following notation: each player i sends player j a report from S in the first round. In the second round he sends j a report in $S^{n-1} \equiv T$.

that explain how the agent will respond to the messages of the other agents. The messages that the agents hear from other agents are specified by the principal, so that this makes sense. Here, the agents respond to external messages that they themselves specify when the construct their own mechanism. So the usual formalism won't work here.

Let s^i_j be the first round signal that player i makes to player j for each $j \neq i$. For each $i \neq j, k$, let t^i_{jk} be the report that player i makes to player j about the signal he received from player k in the first round. In this notation t^i_{jj} is the report that i makes to j about the signal he received from him. On the first round, each player then sends out $n-1$ signals.

Definition 2. A *confirmation process* for player i with message space S takes any array in $S^{(n-1)} \times S^{(n-1)^2}$ into an array in S^n according to (2.1)

$$\tau^i_j\left(s_{-i}, t_{-i}\right) = \begin{cases} t' & \text{if } j = i; \left\{\exists! j' : t^{j'}_{ii} \neq t^k_{ii} \equiv t' \forall k \neq j, i\right\} \vee \left\{t^{j'}_{ii} = t^k_{ii} = t' \forall j', k \neq i\right\} \\ t' & \text{if } j \neq i; \left\{t^k_{ij} = t' \forall k \neq j\right\} \vee \left\{\exists! k \neq i, j : t^k_{ij} \neq s^j_i = t'\right\} \\ 0 & \text{otherwise.} \end{cases}$$

A sequential communications mechanism based on a confirmation process is a is one that has the property that actions only depend on the τ^i_j. The number τ^i_j is computed slightly differently depending on whether $j = i$. Player i computes τ^i_i by looking at the element of S that each of the other players say they heard from him in the first round. If all these reports, or all but one of these reports are the same, then τ^i_j is set equal to this common report. Otherwise, it is set to zero. τ^i_j is computed by looking at j's report in the first period, and what the others said that j reported to them in the first period. If all these messages, or all but one of these messages agree, then τ^i_j is set equal to that common message. Otherwise, it is set to zero.

We need n to be at least 4 in the construction because the uniqueness restriction (the left hand one in the first line of (2.1), and the right hand one in the second line) requires that a strict majority of players make the same report. Since player i communicates with $n-1$ others, $n-1$ has to be at least 3. A confirmation process is useful when players report truthfully.

When a collection of n players all use an identical confirmation process, we say that player i uses a *truthful revelation strategy* if his first round report to every player is the same, and his second round report about what he heard in the first round is always truthful.

Lemma 3. *Suppose a set of $n \geq 4$ players all use an identical confirmation process $\tau(\cdot)$ with message space S, and that all players other than i are expected to follow truthful revelation strategies. Then whatever the realizations (s_{-i}, t_{-i}) of the others' reports, $\tau_k(s_{-i}, t_{-i}) = \tau_k(s_{-j}, t_{-j})$ for each k, and for every pair i and j no matter what reports that player i makes (be they truthful or not). Player i can follow a strategy* *that assigns any value in S to $\tau_i(s_{-j}, t_{-j})$ for $j \in N$, but otherwise he cannot affect $\tau_k(s_{-j}, t_{-j})$ for any $k \neq i$.*

Proof. First, suppose all players do what they are supposed to, sending the same message to every other player on the first round, and truthfully reporting the messages they received. Suppose that i's report on the first round to every other player is s'. Then

$$\tau_i(s_{-i}, t_{-i}) = s'$$

because each of the others will send back the same report s' that they received from i in the first round. On the other hand, player j, hears the report s' from player i, and has each of the players other than i say that they heard s' as well. So $\tau_i(s_{-j}, t_{-j}) = s'$. On the other hand $\tau_k(s_{-i}, t_{-i})$ is computed from a report $\tilde{s}'$ that k makes to i in the first round. Since k is expected to make that same report to each of the other players, $t^j_{ik} = \tilde{s}'$ as well. So by the second line of (2.1), $\tau_k(s_{-i}, t_{-i}) = \tilde{s}'$. Player j is expected to receive the same report $\tilde{s}'$ from player k that i did, since k is expected to send the same first round report to everyone. The others, including player i are expected to report to j that k sent them the report $\tilde{s}'$ in the first round, so $\tau_k(s_{-j}, t_{-j}) = \tilde{s}'$. Even if i deviates and lies to the others on the second round, $\tau_k(s_{-j}, t_{-j}) = \tilde{s}'$, since by (2.1), i's unilaterally different report will be ignored.

Player i can try to manipulate these numbers by sending different messages to the others in the first round, then lying about the reports he received when he makes his second round reports. Suppose i sends out 3 or more distinct messages in the first round. The others are expected to truthfully report these on the second round. Then $\tau_i(s_{-i}, t_{-i})$, which is based only on the second round reports of the others, must fail both conditions in the first line of (2.1), and $\tau_i(s_{-i}, t_{-i}) = 0$. Player j receives a first round report from i. Player i expects player j to hear truthful reports of the messages he sent to the others. Since i sent 3 or more distinct messages on the first round, there must be at least two distinct messages that player j receives from the others about i's report. Then from the second line of (2.1), $\tau^j_i(s_j, t_j) = 0$. These two conclusions will be true no matter what i reports to the others on the second round since neither of these numbers depends on i's second round reports. If k reports $\tilde{s}'$ as above, $\tau_k(s_{-j}, t_{-j}) = \tilde{s}'$ whether i lies on the second round or not.

Very similar arguments apply when i sends two distinct messages in the first round, say s' to all but one of the other players, and s'' to the other, say player k. Then using (2.1), $\tau_i(s_{-i}, y_{-i}) = s'$ since only player k disagrees about what i reported. For j, $\tau_i(s_{-j}, t_{-j}) = s'$. Player k

reports that i told him s'' on the first round, but i reported s' to j and the others confirm that is what they heard. For the rest the arguments are as above. □

The jist of the argument is that if the others participate truthfully in a confirmation mechanism, then i can go along and be truthful, or not. If not, he always looks like a unilateral dissenter and is ignored. However, he can manipulate the value of $\tau_i^k(\cdot,\cdot)$ using a truthful strategy simply by sending the same message to all players.

Theorem 4. *Suppose there are* 4 *or more players. Then any incentive compatible and individually rational allocation rule can be supported as a Bayesian equilibrium in the competing mechanism game.*

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 on the first round. This construction will explain both how the messages are used to implement a randomization with contracts that map into pure actions, and also how the two rounds of messages can be used to correlate actions without a public randomizing device while still satisfying the single-deviation perfection requirement.

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

$$r_i(s_{-i}, t_{-i}) = \alpha_i^{q\left(\tau^\theta(s_{-i},t_{-i})\right)}\left(\lfloor \sum_{j\in N} \tau_j^x(s_{-i}, t_{-i})\rfloor, A\right). \tag{2.2}$$

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

$$r_i^k(s_{-ik}, t_{-ik}) = \alpha_i^{p_i\left(\tau^\theta(s_{-ik},t_{-ik})\right)}\left(\lfloor \sum_{j\neq k} \tau_j^x(s_{-ik}, t_{-ik})\rfloor, A_{-k}\right) \tag{2.3}$$

to each non-deviating player i, where (s_{-ik}, t_{-ik}) is an array of messages from the other *non-deviating* players.

In any history in which all players offer a recommendation mechanism, player i should truthfully report to each player $k \neq j$ the message received from player j, should report his type truthfully to every other player and send every other player a signal x drawn uniformly from $[0, 1]$, and should recommend to player j that he should use the sequential communication mechanism based on a confirmation process $r_j(\cdot, \cdot)$ as defined in (2.2) above. In any history in which a single player, say player k, has deviated and offered some mechanism other than a recommendation mechanism, player i should truthfully report the private message received from each player $j \neq k$ to each player $j' \neq k, j$, should choose a message s' using a uniform distribution on $[0, 1]$ and should send that message along with his true type to all players other than k along with a recommendation to use the sequential communciation mechanism r_j^k as defined in (2.3).

Now we proceed to prove that the strategies specified constitute a Bayesian equilibrium by showing that no player alone can affect the allocation except by sending false information about his type. First, it is immediately a best reply for each player to offer a recommendation mechanism. The reason is that no matter which continuation equilibrium is played in response to the deviation, the deviator should expect the others to implement joint action $p_i(\theta)$. As a consequence, i's payoff cannot exceed $\underline{u}_i$ which is less than the payoff associated with the recommendation mechanism.

The rest of the argument is similar. Each player i should recommend that player j use the sequential communication mechanism based on a confirmation process $r_j(\cdot, \cdot)$. The reason is that since there are at least 4 players in the game by assumption, there are at least two other players who are expected to recommend $r_j(\cdot)$ to player j. So j is going to implement r_j no matter what i recommends. As a result it is a best reply for i to recommend r_j as well.

On the equilibrium path, all players offer recommendation mechanisms, and each player recommends $\{r_j\}_{j \neq i}$. The r_j are sequential communication mechanisms based on a confirmation process, and other players are expected to use truthful reporting strategies when they participate in these mechanisms. By Lemma 3, each player's action is based on the same collection of numbers $\left\{\tau_j^\theta, \tau_j^x\right\}_{j \in N}$, and i can only affect the value of $\left\{\tau_i^\theta, \tau^x\right\}$. Since the others are expected to report their types truthfully $\tau_j^\theta = \theta_j$ for each $j \neq i$. Since others are expected to send the same signal x to the others uniformly distributed on $[0, 1]$, τ_j^x has a uniform distribution. As we have explained in Remark 1 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 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 $[0,1]$. □

A few remarks. The traditional model of competing mechanisms has a number of uninformed mechanism designers competing to influence a number of privately informed agents. In the model here, there are only players. However, the traditional formulation is captured by noting that the difference between principals and agents is whether or not they have some action that they control. Some subset of the players in our formulation could be players who have private information that other players care about, but who take no actions of their own. What these 'agents' do is to make recommendations to the active players about the mechanisms they should use. They also collect type reports and signals from all other agents, which they pass on to all the other players. So our model differs from the traditional model only in so far as communication takes place sequentially. From the proof above, the role of sequential communication is to make it possible to implement correlated actions.

It is worth remarking that a simple common agency with two principals and one agent isn't covered by our theorem (because there are only three players the confirmation mechanism doesn't work). This simple two principal common agency is the model that is used most widely to illustrate properties of competing mechanism equilibrium.

A single mechanism designer problem in which all but one of the players are agents without action does fit within our folk theorem. However, in that case, the only individually rational allocations are those that are optimal for the single player who controls all actions.

Note also that the theorem uses Bayesian equilibrium as a solution concept. The methods used in the theorem won't work for refinements like perfect Bayesian equilibrium. When there is a deviation, the confirmation process ensures that it will still be a best reply for players to send the same signal and type report to all of the non-deviating players provided there are at least four non-deviators. So it is still possible to support correlated actions. However, punishments aren't necessarily incentive compatible. As a result, a confirmation process by itself won't ensure that players report their types truthfully. Also it isn't generally possible to tighten the individual rationality constraint to make punishments incentive compatible. The reason is that whether or not something is incentive compatible off the equilibrium path depends on the deviation that gets you there.

One way to provide a trdrefinements is to restrict attention to complete information. A simple application of Theorem 4 to the case of complete information gives:

Corollary 5. *Let $p \in P$ be a joint randomization in which each player receives at least his minmax payoff $\underline{u}_i$. If there are four or more players, then p can be supported as a Nash equilibrium of the competing mechanism game.*

If one player deviates and offers something other than a recommendation mechanism, it will still be a best reply for every player to recommend the minmax punishment to each of the non-deviators, then send the same signal to all players, effectively correlating the punishment. If there are at least four non-deviators, then deviating from these rules has no effect on the outcome when all the others are using them, making the rules best replies to one another.

The complication in formalizing this is that subgame perfection requires that continuation play in all proper subgames be Nash. There is a proper subgame associated with *every* array of mechanisms offered in the first round, including mechanisms that might not seem to make sense from a modelling perspective. For example, suppose 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, and none of the others bothers to offer any mechanism. Player 1 offers a mechanism that asks player 2 and 3 to report a number between 0 and 1. He ignores player 4. The mechanism he offers is

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

Payoffs for player 2 are$u(a) = -1$, $u(b) = 0$, and $u(c) = 1$. Player 3's payoff is $-u$. This game has no Nash equilibrium in either pure or mixed strategies (Sion and Wolfe). Since this is a possible mechanism for player 1 to use, the competing mechanism game has no subgame perfect Nash equilibrium.

We mention this example to show that there is no simple way to refine Nash equilibrium in the competing mechanism game. However, the following result follows from our argument above:

Corollary 6. *Let p be a joint randomization supported as a Nash equilibrium in recommendation mechanisms as above in a competing mechanism game with at least 5 players. Then in any proper subgame in which player i has deviated and offered something other than a recommendation equilibrium, the strategies specified for the non-deviating players constitute a best reply no matter what the deviating player does.*

Proof. This follows immediately from Lemma 3 since there are at least 4 non-deviating players in a game with 5 players. □

Observe that this is not enough to establish the existence of a subgame perfect equilibrium, since some variant of the Zion-Wolfe mechanism described above may exist for some deviating player. This can't be resolved by ruling out mechanisms for which there is no equilibrium as is typically done in standard mechanism design, since whether or not the bad mechanism exists may depend on the mechanisms that the other (non-deviating) players offer.

Conclusion

The folk theorem provides a complete characterization of allocations that are supportable as Bayesian equilibrium in competing mechanism games. Ultimately, models based on competing mechanisms have no content unless they impose restrictions on the set of mechanisms. One obvious way to restrict mechanisms is to rule out sequential communication. Then the theorem in (Yamashita 2007) applies. His 'pure strategy' folk theorem illustrates that in most situations, there are still far too many equilibria.

A second way to limit the set of outcomes is to argue that the equilibria which are used to support them are unconvincing. For example, they require a high degree of coordination (of expectations), are subject to collusion (since players agree to do whatever the others want). However, for the reasons mentioned above, refining equilibrium in competing mechanism games is difficult.

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 (2007): "A Commitment Folk Theorem," Working paper, Tel-Aviv University.

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

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.

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

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

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

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.

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

Yamashita, T. (2007): "A Revelation Principle and A Folk Theorem without Repition in Games with Multiple Principles and Agents," manuscript, Stanford University.