Competing Mechanisms and Folk Theorems: Two Examples*

Andrea Attar[†] Eloisa Campioni[‡]
Thomas Mariotti[§] Gwenaël Piaser[¶]

September 25, 2020

Abstract

In competing-mechanism games under exclusivity, principals simultaneously post mechanisms, and agents then simultaneously participate and communicate with at most one principal. In this setting, we develop two complete-information examples that question the logic of the folk theorems for competing-mechanism games established in the literature. In the first example, there exist pure-strategy equilibria in which some principal obtains less than her min-max payoff, computed over all players' actions. Thus folk-theorem-like results must generally involve bounds on principals' payoffs that depend on the spaces of messages available to the agents, and not only on the players' actions. The second example shows that even this nonintrinsic approach is misleading: there exist incentive-feasible allocations in which principals obtain more than their min-max payoffs, computed over arbitrary spaces of mechanisms, but which cannot be supported in equilibrium. Key to these results is the standard requirement that agents' participation and communication decisions are tied together.

Keywords: Competing Mechanisms, Folk Theorems, Exclusive Competition.
JEL Classification: D82.

*We thank the editor and two anonymous referees for very thoughtful and detailed comments. We also thank Gorkem Celik and Takuro Yamashita for valuable feedback. Finally, we thank the seminar audience at Université de Cergy-Pontoise, as well as conference participants at the 2018 International Conference on Game Theory, the 2019 International Conference on Public Economic Theory, and the 2019 SAET Conference for many useful discussions. Financial support from the ANR (Programme d'Investissement d'Avenir ANR-17-EURE-0010), the Chaire SCOR–TSE, and the MIUR (PRIN 2015) is gratefully acknowledged.

[†]Toulouse School of Economics, CNRS, University of Toulouse Capitole, Toulouse, France, and Università degli Studi di Roma "Tor Vergata," Roma, Italy. Email: andrea.attar@tse-fr.eu.

[‡]Università degli Studi di Roma "Tor Vergata," Roma, Italy. Email: eloisa.campioni@uniroma2.it.

[§]Corresponding author. Toulouse School of Economics, CNRS, University of Toulouse Capitole, Toulouse, France, CEPR, and CESifo. Email: thomas.mariotti@tse-fr.eu.

[¶]IPAG Business School, Paris, France. Email: gwenael.piaser@ipag.fr.

1 Introduction

Competition in financial, labor, and monetary markets is often modeled by assuming that sellers noncooperatively design trade mechanisms. Final allocations are then determined by buyers' strategic participation and communication decisions in these mechanisms. In competing auctions (McAfee (1993), Peters (1997), Peters and Severinov (1997), Virag (2010)), privately informed buyers observe the posted mechanisms, choose the auction they want to participate to, and then bid according to their valuations. In competitive search (Moen (1997), Eeckhout and Kircher (2010), Wright, Kircher, Julien and Guerrieri (2019)), buyers apply to their preferred trade mechanism, meet a seller according to some meeting technology, and, once in a meeting, communicate with the seller they are matched with, possibly revealing private information to her. These interactions are instances of competing-mechanism games in which principals first simultaneously commit to mechanisms and agents then simultaneously participate and communicate with principals. Mechanisms are public and are not contingent on one another. That is, when designing a mechanism, a principal cannot directly condition her actions on the *market information* generated by her competitors' mechanisms, and held by the agents.

However, a principal can, in principle, condition her actions on reports by the agents about their market information. Since the seminal work of Epstein and Peters (1999), the literature on competing mechanisms has emphasized that providing agents with the opportunity to report both their exogenous private information and their endogenous market information can spectacularly enlarge the set of equilibrium allocations. Following Yamashita (2010), several contributions (Peters and Troncoso-Valverde (2013), Xiong (2013), Ghosh and Han (2018)) have indeed offered different versions of a folk theorem: in a nutshell, letting principals' mechanisms be sufficiently reactive to agents' reports about their market information allows one to support in equilibrium any incentive-compatible allocation that yields each principal a payoff above a well-specified min-max bound. These results are established under fairly general conditions on the primitives of the game, which questions the relevance of the equilibrium analyses provided in the applied literature.

The present paper further elaborates on this issue. We focus on situations in which agents' participation decisions are strategic, in line with the intended economic applications of the competing-mechanism paradigm, and we provide two examples that fundamentally challenge the logic of folk theorems in this context.

As a contrast to our results, it is useful to review the arguments leading to folk theorems for competing-mechanism games, as first developed by Yamashita (2010). Let each principal commit to punishing a unilateral deviation by any of her competitors, when reported by a majority of agents through appropriate messages, and assume that the agents' message spaces are rich enough to allow them to select a specific punishment for each of these deviations. Then, if there are at least three agents, any such punishment can be selected by majority voting in the message game played by the agents. When unilaterally deviating, each principal hence anticipates that her competitors will react by min-maximizing her: as a result, any incentive-compatible allocation yielding her a payoff above the corresponding bound can be supported in equilibrium by such mechanisms. The lower bound of a principal's equilibrium payoffs turns out to coincide with her min-max payoff computed over all mechanisms available to principals, which establishes the folk theorem. This equality, however, obtains by making agents able to select the worst punishment against every principal in every subgame, which effectively requires that *each* agent participates and communicates with *each* principal for *any* profile of posted mechanisms. We find this assumption hard to justify in the light of economic applications, in which agents' communication decisions are closely tied to their participation decisions. That is, an agent can communicate the information he possesses to a given principal only if he chooses to participate with her.

We show that taking into account agents' participation decisions in competing-mechanism games has dramatic implications for the possibility of deriving folk-theorem results. We focus on the situation in which agents participate and communicate with at most one principal, as is assumed both in competitive-auction and in competitive-search models. In this exclusive-competition scenario, we construct two examples for the complete-information case in which agents' types are degenerate, so that they can only take participation decisions and can only report about their market information.

Our first example exhibits equilibria of competing-mechanism games in which some principal obtains a payoff below her min-max payoff, computed over the set of principals' actions. In this example, the explicit consideration of agents' participation decisions leads to discontinuities and nonconvexities that prevent from applying the standard min-max logic, despite the fact that there are only two principals who are allowed to randomize over their decisions. The min-max payoff of the principal in question is strictly higher than her max-min payoff, and hence cannot be a relevant bound for equilibrium characterization. The result suggests that to establish a folk theorem in complete-information games, one may need to specify a nonintrinsic bound that depends on the agents' message spaces, which in turn limits the predictive power of the approach. This result is robust to several extensions. First, it does not depend on the assumption that the agents' communication decisions are tied to their participation decisions: as long as participation is strategic, the same characterization obtains even if communication is unrestricted. Second, it extends to situations in which each principal only *partially* delegates to agents the implementation of her final action, retaining the option to select it from a menu in the last stage of the game.

Our second example establishes that even the nonintrinsic approach is unsatisfactory when communication is tied to participation. In this example, inspired by the competing-hierarchy model of Myerson (1982), each agent has a dominant participation strategy, and principal-agents hierarchies are thus fixed in the game. In addition, the min-max payoff for each principal can be straightforwardly computed over arbitrary mechanisms, and it coincides with the corresponding max-min payoff. Yet, the fact that each agent can communicate with at most one principal makes it impossible to construct sophisticated equilibrium threats. As a consequence, although there exist many incentive-feasible allocations in which principals obtain more than their min-max payoffs, none of them can be supported in equilibrium. Each principal's equilibrium payoff thus coincides with her min-max payoff, even if the analysis is extended so as to allow for mixed-strategy equilibria.

Related Literature

This paper is closely related to the recent literature on folk theorems in competing-mechanism games initiated by Yamashita (2010). The main contribution of Yamashita (2010) is to highlight the role of *recommendation mechanisms.* By offering such a mechanism, a principal commits to a direct mechanism if all but one agent recommend her to do so. Letting principals post recommendation mechanisms makes it possible to reproduce the equilibrium allocations associated to the *universal* space of mechanisms identified by Epstein and Peters (1999). Yamashita (2010) further assumes that each agent participates and communicates with all principals for any profile of mechanisms, which allows him to rely on recommendation mechanisms to derive an equilibrium characterization in terms of principals' min-max payoffs. His analysis has been extended in two important directions.

First, the bound for principals' equilibrium payoffs proposed by Yamashita (2010) is sensitive to the mechanisms available in the game. This makes it difficult to evaluate his contribution in the light of standard folk theorems. The recent work of Peters and Troncoso-Valverde (2013) provides an abstract framework for incomplete-information games and formulates the corresponding bounds through agents' incentive constraints in the spirit of Myerson (1979). The bounds are therefore defined in terms of the primitives of the model and, in contrast to Yamashita (2010), do not depend on the set of available (indirect) mechanisms. Ghosh and Han (2018) extend Yamashita (2010) to repeated interactions and reformulate the bounds on principals' equilibrium payoffs in these settings.

Yet, neither Peters and Troncoso-Valverde (2013) nor Ghosh and Han (2018) allow players to take any action after mechanisms are posted. Our first example shows that, under complete information, the bounds identified by Peters and Troncoso-Valverde (2013) and Ghosh and Han (2018) are no longer relevant for principals' equilibrium payoffs if agents' participation decisions are taken into account.

Second, Yamashita (2010) restricts principals to deterministic mechanisms and only considers pure-strategy equilibria of the agents' game. Szentes (2009) shows that this restriction is critical by constructing a simple complete-information game in which the equilibrium allocations supported by deterministic mechanisms yield a principal a payoff below Yamashita's (2010) min-max bound. Xiong (2013) provides a generalized version of Yamashita (2010), in which mixed strategies are allowed and a folk theorem is established. Crucially, he also assumes that each agent always communicates with all principals. We share with Szentes (2009) the focus on complete information, but we allow principals to post random mechanisms, and we do not restrict agents to play pure strategies. In contrast with these approaches, we explicitly model agents' participation decisions. Our second example then shows that recommendation mechanisms may not guarantee a system of punishments allowing each principal to min-maximize her opponents.

An alternative route to folk theorems in competing mechanism games is based on the notion of *contractible contracts*. Following Tennenholz (2004) and Kalai, Kalai, Lehrer, and Samet (2010), Peters and Szentes (2012), Peters (2015) and Szentes (2015) let principals design mechanisms that depend on the mechanisms of their opponents. This allows them to *directly* punish a deviator in a way that depends on the specific deviation she chooses, which can yield a folk theorem even in the absence of a strategic role for the agents. The observability requirements underlying this approach, however, are too demanding in the light of the economic applications we consider.

The paper is organized as follows. Section 2 introduces a general model of exclusive competition under complete information. Sections 3 and 4 present our examples. Section 5 discusses the robustness of our results to unrestricted communication. Section 6 concludes.

2 The Model

We consider a setting in which several principals, indexed by $j \in \mathcal{J} \equiv \{1, \ldots, J\}$, contract with several agents, indexed by $i \in \mathcal{I} \equiv \{1, \ldots, I\}$. Agents have no private information, and we denote each agent's single type by t.

Actions and Payoffs Agents only take participation decisions, and we denote by $a^i_j \in A^i_j \equiv \{Y, N\}$ agent i's decision to participate (Y) or not (N) with principal j. Such decisions are exclusive, in that each agent i can participate with at most one principal j. Hence, overall, each agent i takes an action a^i in the set $A^i \equiv \{(a^i_1, \ldots, a^i_J) : a^i_j = Y \text{ for at most one } j\}$. Each principal j in turn takes an action x_j in a finite set X_j. We let $v_j : A \times X \to \mathbb{R}$ and $u^i : A \times X \to \mathbb{R}$ be the payoff functions of principal j and of agent i, respectively, where $A \equiv A^1 \times \ldots \times A^I$ and $X \equiv X_1 \times \ldots \times X_J$.

Communication Communication takes place through the public mechanisms posted by the principals and the messages sent by the agents in these mechanisms. Formally, agent i sends a private message m^i_j to principal j in some Polish space M^i_j.[1] Each message space M^i_j includes the empty message $\emptyset$, which corresponds to agent i not communicating with principal j, as well as the trivial message t. Communication is tied to participation, in the sense that agent i sends a nonempty message to principal j if and only if he decides to participate with her. Hence, overall, each agent i sends messages m^i in the space $M^i \equiv \{(m^i_1, \ldots, m^i_J) : m^i_j \neq \emptyset \text{ for at most one } j\}$, and we say that a profile $(m^i, a^i) \in M^i \times A^i$ is *consistent for agent* i whenever $m^i_j \neq \emptyset$ if and only if $a^i_j = Y$ for all j. We denote by C^i the space of such consistent communication and participation profiles for agent i.

Mechanisms Each principal j can take an action contingent on the messages $m_j \in M_j$ she receives and the agents' decisions $a_j \in A_j$ to participate with her, where by definition $M_j \equiv M^1_j \times \ldots \times M^I_j$ and $A_j \equiv A^1_j \times \ldots \times A^I_j$. We say that a profile $(m_j, a_j) \in M_j \times A_j$ is *consistent for principal* j if $m^i_j \neq \emptyset$ if and only if $a^i_j = Y$ for all i. We denote by C_j the space of such consistent communication and participation profiles for principal j. Notice that, because M_j is Polish and A_j is finite, C_j is Polish. A mechanism for principal j is a Borel-measurable mapping $\gamma_j : C_j \to \Delta(X_j)$ that associates to every consistent communication and participation profile for principal j a lottery over her actions.

Admissibility Whereas most of our analysis focuses on situations in which principals play pure strategies in equilibrium, a general requirement for defining expected payoffs in our and related games is that the evaluation mapping $(\gamma_j, c_j) \mapsto \gamma_j(c_j)$ describing how the distribution

[1] Our first example only allows for minimal communication, and thus finite message spaces. Our second example, by contrast, allows for rich communication, and thus uncountable message spaces. Requiring these spaces to be Polish entails no loss of generality.

of principal j's action varies with her mechanism and the consistent communication and participation profile she observes be measurable. Thus, at a minimum, we must define a measurable structure on the set $\Gamma_j^{M_j}$ of mechanisms for principal j. Two cases can arise. If M_j and, hence, C_j, is countable, we can take $\Gamma_j^{M_j}$ to be the set of all Borel-measurable mappings $\gamma_j : C_j \to \Delta(X_j)$; a natural measurable structure on $\Gamma_j^{M_j}$ is then the product Borel σ-field on the product of at most countably infinitely many copies of $\Delta(X_j)$. If M_j is uncountable, however, there is no measurable structure on the set of all Borel-measurable mappings $\gamma_j : C_j \to \Delta(X_j)$ such that the evaluation mapping for principal j is measurable (Aumann (1961)); in that case, there is no other choice than to restrict the set of admissible mechanisms $\Gamma_j^{M_j}$. Admissibility can be shown to coincide with the requirement that $\Gamma_j^{M_j}$ be of bounded Borel class (Aumann (1961), Rao (1971)), allowing for a rich class of mechanisms for our analysis. With this caveat in mind, we hereafter fix an admissible space $\Gamma_j^{M_j}$ of mechanisms for every principal j, with associated σ-field $\mathcal{G}_j^{M_j}$.

Strategies and Timing The competing-mechanism game G^M induced by $M \equiv M_1 \times \ldots \times M_J$ unfolds in three stages:

1. The principals simultaneously post mechanisms.

2. The agents simultaneously take consistent communication and participation decisions.

3. The principals' mechanisms are implemented, lotteries realize, and all payoffs accrue.

A strategy for principal j is a probability measure $\mu_j \in \Delta(\Gamma_j^{M_j})$ over the σ-field $\mathcal{G}_j^{M_j}$. A strategy for agent i is a measurable mapping $\lambda^i : \Gamma^M \to \Delta(C^i)$ that associates to every profile of mechanisms a probability measure over consistent communication and participation profiles for agent i, where $\Gamma^M \equiv \Gamma_1^{M_1} \times \ldots \times \Gamma_J^{M_J}$ is endowed with the product σ-field $\mathcal{G}_1^{M_1} \otimes \ldots \otimes \mathcal{G}_J^{M_J}$ and $\Delta(C^i)$ with the Borel σ-field. The allocation $z(\gamma, \lambda)$ induced by the mechanisms $\gamma \equiv (\gamma_1, \ldots, \gamma_J)$ and the strategies $\lambda \equiv (\lambda^1, \ldots, \lambda^I)$ is the probability measure over outcomes in $A \times X$ uniquely defined by the marginal of $\lambda^1(\gamma) \otimes \ldots \otimes \lambda^I(\gamma)$ over A and the probability transitions γ_j from C_j to X_j. Notice that the mapping $z(\cdot, \lambda) : \Gamma^M \to \Delta(A \times X) : \gamma \mapsto z(\gamma, \lambda)$ is measurable. Hence we can define the allocation $z(\mu, \lambda)$ induced by the strategies $\mu \equiv (\mu_1, \ldots, \mu_J)$ and $\lambda \equiv (\lambda^1, \ldots, \lambda^I)$ by

$$z(\mu, \lambda)(a, x) \equiv \int_{\Gamma^M} z(\gamma, \lambda)(a, x)\, \mu_1(\mathrm{d}\gamma_1) \otimes \ldots \otimes \mu_J(\mathrm{d}\gamma_J)$$

for all $(a, x) \in A \times X$.

Equilibrium The strategy profile (μ, λ) is a subgame-perfect Nash equilibrium (SPNE) of

G^M whenever:

(i) For each $\gamma \in \Gamma^M$, $(\lambda^1(\gamma), \ldots, \lambda^I(\gamma))$ is a Nash equilibrium in the subgame γ played by the agents.

(ii) Given the continuation equilibrium λ, μ is a Nash equilibrium of the game played by the principals.

We denote by $\Lambda^*(\gamma)$ the set of Nash equilibria of the subgame γ. Following Epstein and Peters (1999) and Han (2007), we will mostly focus on SPNEs of G^M in which principals play pure strategies (SPNE-PSP). That is, each principal deterministically posts a mechanism; notice, however, that this mechanism may involve randomization over her final actions.

Direct Mechanisms In our complete-information setting, a *direct* mechanism for principal j is a mechanism such that the message space M_j is restricted to the trivial message t and the empty message $\emptyset$. Because in this case communication decisions are redundant relative to participation decisions, such a mechanism can be identified to a mapping $\tilde{\gamma}_j : A_j \to \Delta(X_j)$ that associates to every participation profile for principal j a lottery over her actions. We denote by $\tilde{\Gamma}_j$ the set of direct mechanisms for principal j, and we let $\tilde{\Gamma} \equiv \tilde{\Gamma}_1 \times \ldots \times \tilde{\Gamma}_J$ and $\tilde{G}$ be the competing-mechanism game in which principals are restricted to direct mechanisms. Notice that, because A_j and X_j are finite, $\tilde{\Gamma}_j$ is a compact and convex subset of a Euclidean space. A strategy for agent i in $\tilde{G}$ can be identified to a Borel-measurable mapping $\tilde{\lambda}^i : \tilde{\Gamma} \to \Delta(A^i)$ that associates to every profile of direct mechanisms a lottery over participation profiles for agent i. We say that the allocation $z(\tilde{\gamma}, \tilde{\lambda})$ is *incentive-feasible* if $\tilde{\lambda} \in \Lambda^*(\tilde{\gamma})$.

A Team Interpretation Our complete-information setting with observable contracts is reminiscent of the environment studied by Fehrstman, Judd, and Kalai (1991), with the difference that we do not impose a priori restrictions on mechanisms and allow each principal to delegate the choice of her actions to several agents. More concretely, our setting has a natural interpretation as a team-formation model. In this interpretation, principals are teams—such as, for instance, football teams or law firms—while agents are players—football players or lawyers—who can each join at most a single team. Each team can commit to possibly random actions contingent on the identities of the players who choose to join it; by contrast, the composition of other teams is noncontractible, or simply not observable at the time the team chooses its actions. The teams' action sets may include, for instance, task allocation and remuneration schemes. Teams' and players' payoffs can depend in complex ways on the composition of each team and on actions taken within the teams. This team-formation model is in line with the competing-hierarchy model of Myerson (1982). The difference is that the composition of teams is endogenous in our model, in that it depends on the mechanisms offered by each team, whereas it is taken as exogenous in the standard competing-hierarchy model.

3 The First Example

Our first example emphasizes the impact of agents' strategic participation decisions on equilibrium outcomes. Specifically, we show that, as long as these decisions are payoff-relevant, payoffs for the principals below the min-max-min bounds identified by Peters and Troncoso-Valverde (2013), Peters (2014), and Ghosh and Han (2018) can be supported in equilibrium.

3.1 The Physical Environment

Let $I = J \equiv 2$ and let $X_1 \equiv \{x_{11}, x_{12}\}$ and $X_2 \equiv \{x_{21}, x_{22}\}$ be the sets of actions of principal 1 (P1) and principal 2 (P2), respectively. Let $A^1 = A^2 \equiv \{YN, NY, NN\}$ be agent 1's (A1) and agent 2's (A2) sets of actions, where YN, for instance, refers to the agent participating with P1 but not with P2. Payoffs are represented in the matrix in Table 1 below, in which agents choose in the external box and principals choose in the internal 2×2 cells. Each array represents the payoffs to P1, P2, A1, and A2, respectively.

		YN			NY	
		x_{21}	x_{22}		x_{21}	x_{22}
YN	x_{11}	$(0, \zeta, 5, \frac{25}{2})$	$(0, \zeta, 5, \frac{25}{2})$	x_{11}	$(0, 0, 5, 12)$	$(0, 10, 5, 8)$
	x_{12}	$(0, \zeta, 10, \frac{9}{2})$	$(0, \zeta, 10, \frac{9}{2})$	x_{12}	$(0, 0, 10, 12)$	$(0, 10, 10, 8)$
		x_{21}	x_{22}		x_{21}	x_{22}
NY	x_{11}	$(0, 10, 8, 12)$	$(0, 8, 9, 12)$	x_{11}	$(0, \zeta, 40, 7)$	$(0, \zeta, 4, 13)$
	x_{12}	$(0, 10, 8, 8)$	$(0, 8, 9, 8)$	x_{12}	$(0, \zeta, 40, 7)$	$(0, \zeta, 4, 13)$

Table 1: The payoff matrix.

Observe that P1's payoff is constantly equal to 0, and that, for any profile of participation decisions by the agents, P1's choice of action has no impact on P2's payoff. Hence there are no direct payoff externalities between the principals, and P1 can affect P2's payoff only insofar as she can influence the agents' participation decisions through her choice of a mechanism. We assume, in addition, that the no-participation decision NN is strictly dominated for every agent, and we let $\zeta < 0$ be an arbitrarily large loss for P2 if neither A1 nor A2 participate with her, or if both A1 and A2 participate with her. Notice also that there exists at least one incentive-feasible allocation yielding P2 her maximal payoff of 10. To see this, consider the simple direct mechanisms in which P1 chooses x_{11} and P2 chooses x_{21} for any participation decisions of the agents. The resulting subgame played by the agents only admits the Nash equilibrium (NY, YN), which yields P2 a payoff of 10.

3.2 The Game $\tilde{G}$

We first consider the game $\tilde{G}$ in which each principal does not ask for private messages and associates to every profile of agents' decisions to participate with her a lottery over her actions. Let us first observe that a direct mechanism for P1, say $\tilde{\gamma}_1$, is represented by the following list of participation-contingent probability distributions over X_1:

$$\tilde{\gamma}_1(Y,Y) = (\delta_{Y,Y}, 1-\delta_{Y,Y}), \qquad \tilde{\gamma}_1(Y,N) = (\delta_{Y,N}, 1-\delta_{Y,N}),$$
$$\tilde{\gamma}_1(N,Y) = (\delta_{N,Y}, 1-\delta_{N,Y}), \qquad \tilde{\gamma}_1(N,N) = (\delta_{N,N}, 1-\delta_{N,N}),$$

where $\delta_{a_1^1,a_1^2}$ denotes the probability of P1 choosing action x_{11} given a participation profile $(a_1^1, a_1^2) \in A_1$ for P1. Thus, for instance, $\delta_{Y,N}$ is the probability that P1 chooses x_{11} if only A1 chooses to participate with her. Similarly, a direct mechanism for P2, say $\tilde{\gamma}_2$, is represented by participation-contingent probability distributions over X_2, and we let $\sigma_{a_2^1,a_2^2}$ denote the probability that P2 chooses x_{21} given a participation profile $(a_2^1, a_2^2) \in A_2$ for P2. Our first result is that the principals' payoffs are uniquely pinned down in any SPNE-PSP of $\tilde{G}$.

Proposition 1 *The principals obtain payoffs* $(0, 10)$ *in any SPNE-PSP of* $\tilde{G}$.

Proof. Recall that P1's payoff is constantly equal to 0. We thus only need to show that for every direct mechanism posted by P1, that is, for every family of transition probabilities $\delta_{a_1^1,a_1^2}$, there exists a direct mechanism for P2, that is, a family of transition probabilities $\sigma_{a_2^1,a_2^2}$, inducing a unique Nash equilibrium in the subgame played by the agents, in which P2 achieves her maximal payoff of 10.

Case 1 Suppose first that $\delta_{Y,Y} > \frac{2}{5}$. In this case, let P2 post a mechanism such that $\sigma_{Y,Y} = \sigma_{N,Y} = \sigma_{Y,N} = 1$, which induces the subgame in Table 2.

	YN	NY
YN	$(10-5\delta_{Y,Y}, \frac{9}{2}+8\delta_{Y,Y})$	$(10-5\delta_{Y,N}, 12)$
NY	$(8, 8+4\delta_{N,Y})$	$(40, 7)$

Table 2: The subgame of $\tilde{G}$ induced by $\delta_{Y,Y} > \frac{2}{5}$ and $\sigma_{Y,Y} = \sigma_{Y,N} = \sigma_{N,Y} = 1$.

Because $\delta_{Y,Y} > \frac{2}{5}$, we have $10 - 5\delta_{Y,Y} < 8$. Thus NY is a strictly dominant strategy for A1 in this subgame, which guarantees that (NY, YN) is the unique Nash equilibrium. Because $\sigma_{Y,N} = 1$, P2 obtains a payoff of 10.

Case 2 Suppose next that $\delta_{Y,Y} \leq \frac{2}{5}$. In this case, let P2 post a mechanism such that $\sigma_{Y,Y} = \sigma_{N,Y} = \sigma_{Y,N} = 0$, which induces the subgame in Table 3.

	YN	NY
YN	$(10 - 5\delta_{Y,Y}, \frac{9}{2} + 8\delta_{Y,Y})$	$(10 - 5\delta_{Y,N}, 8)$
NY	$(9, 8 + 4\delta_{N,Y})$	$(4, 13)$

Table 3: The subgame of $\tilde{G}$ induced by $\delta_{Y,Y} \leq \frac{2}{5}$ and $\sigma_{Y,Y} = \sigma_{Y,N} = \sigma_{N,Y} = 0$.

Because $\delta_{Y,Y} \leq \frac{2}{5}$, we have $\frac{9}{2} + 8\delta_{Y,Y} < 8$. Thus NY is a strictly dominant strategy for A2 in this subgame, which, as $10 - 5\delta_{Y,N} \geq 4$, guarantees that (YN, NY) is the unique Nash equilibrium. Because $\sigma_{N,Y} = 0$, P2 obtains a payoff of 10. Hence the result. ■

The proof of Proposition 1 actually shows the stronger result that, given any direct mechanism $\tilde{\gamma}_1$ posted by P1, P2 can defend her maximal payoff of 10 by posting a direct mechanism $\tilde{\gamma}_2$ that induces a unique Nash equilibrium in the subgame $\tilde{\gamma} \equiv (\tilde{\gamma}_1, \tilde{\gamma}_2)$. That is, the following min-max-min payoff for P2:

$$\overline{V}_2 \equiv \min_{\tilde{\gamma}_1 \in \tilde{\Gamma}_1} \max_{\tilde{\gamma}_2 \in \tilde{\Gamma}_2} \min_{\tilde{\lambda} \in \Lambda^*(\tilde{\gamma})} \mathbf{E}_{z(\tilde{\gamma}, \tilde{\lambda})}[v_2(a, x)], \tag{1}$$

is equal to 10 in the game $\tilde{G}$. Notice that the definition of $\overline{V}_2$ only allows principals to offer direct mechanisms; as such, it is specified only in terms of the primitives of the model, that is, the actions available to the players and the resulting payoffs. In this respect, it is analogous to the min-max bounds introduced by Peters and Troncoso-Valverde (2013), Peters (2014), and Ghosh and Han (2018) in competing-mechanism games with complete information, the only difference being that agents can take real participation decisions. By contrast, $\overline{V}_2$ differs from the min-max-min bound introduced by Yamashita (2010), in which the set of admissible mechanisms for each principal includes a recommendation mechanism, that is, a specific indirect mechanism committing her to asking agents to recommend her a direct mechanism and to following the majority recommendation.

3.3 Indirect Mechanisms with Minimal Private Communication

In the game $\tilde{G}$, private communication between the agents and the principals can only take place through the trivial message t, which an agent sends to a principal if he decides to participate with her, and the empty message $\emptyset$, which he sends to her otherwise. Such messages are redundant relative to participation decisions and, hence, essentially trivial. We now consider a game G^M in which a minimal degree of meaningful private communication is allowed for. Specifically, we allow A1 to send *one* additional message m to P1. That is, $M_1^1 \equiv \{t, m, \emptyset\}$, while $M_1^2 = M_2^1 = M_2^2 \equiv \{t, \emptyset\}$ as in the game $\tilde{G}$.

In the game G^M, P2 can only offer direct mechanisms $\tilde{\gamma}_2$ described as above by transition probabilities $\sigma_{a_2^1, a_2^2}$. By contrast, P1 can also offer indirect mechanisms γ_1 contingent on the message m sent by A1, allowing her to generate additional threats. Extending our previous notation, such a mechanism is represented by the following list of message- and participation-contingent probability distributions over X_1:[2]

$$\begin{aligned} \gamma_1((t,Y),Y) &= (\delta_{(t,Y),Y}, 1-\delta_{(t,Y),Y}), & \gamma_1((t,Y),N) &= (\delta_{(t,Y),N}, 1-\delta_{(t,Y),N}), \\ \gamma_1((m,Y),Y) &= (\delta_{(m,Y),Y}, 1-\delta_{(m,Y),Y}), & \gamma_1((m,Y),N) &= (\delta_{(m,Y),N}, 1-\delta_{(m,Y),N}), \\ \gamma_1(N,Y) &= (\delta_{N,Y}, 1-\delta_{N,Y}), & \gamma_1(N,N) &= (\delta_{N,N}, 1-\delta_{N,N}). \end{aligned}$$

The subgame $(\gamma_1, \tilde{\gamma}_2)$ is represented in Table 4 below.

	YN	NY
$(t,Y)N$	$(10-5\delta_{(t,Y),Y}, \frac{9}{2}+8\delta_{(t,Y),Y})$	$(10-5\delta_{(t,Y),N}, 8+4\sigma_{N,Y})$
$(m,Y)N$	$(10-5\delta_{(m,Y),Y}, \frac{9}{2}+8\delta_{(m,Y),Y})$	$(10-5\delta_{(m,Y),N}, 8+4\sigma_{N,Y})$
NY	$(9-\sigma_{Y,N}, 8+4\delta_{N,Y})$	$(4+36\sigma_{Y,Y}, 13-6\sigma_{Y,Y})$

Table 4: The subgame $(\gamma_1, \tilde{\gamma}_2)$ of G^M.

The following result shows that this minimal enlargement of a single agent's message space compared to the game $\tilde{G}$ has a dramatic impact on P2's SPNE-PSP payoff set.

Proposition 2 *If the loss ζ incurred by P1 when both A1 and A2 participate with her is large enough, then the principals can obtain any payoffs in $\{(0,\pi) : \pi \in [0,10]\}$ in an SPNE-PSP of G^M.*

Proof. For the sake of clarity, all equilibrium objects will be indexed by a $*$. We prove that, for each $\sigma^* \in [0,1]$, there exists an SPNE-PSP of G^M in which P1 posts a mechanism γ_1^* and P2 posts a direct mechanism $\tilde{\gamma}_2^*$ such that $\sigma_{N,Y}^* = \sigma^*$ and, on the equilibrium path, A1 participates with P1 and A2 participates with P2 with probability 1, which yields P2 a payoff $10(1-\sigma^*)$.

[2] To alleviate the notation, and when no confusion can arise, we hereafter only indicate the nonempty messages t and m sent by A1 to P1.

Thus fix some $\sigma^* \in [0,1]$. To construct an SPNE-PSP in which P2 posts a direct mechanism $\tilde{\gamma}_2^*$ such that $\sigma^*_{N,Y} = \sigma^*$, we proceed as follows. First, let P1 post a mechanism γ_1^* in which $\delta^*_{(t,Y),Y}$ is such that $\frac{9}{2} + 8\delta^*_{(t,Y),Y} = 8 + 4\sigma^*$, that is, $\delta^*_{(t,Y),Y} = \frac{7+8\sigma^*}{16} \in (0,1)$. In addition, let $\delta^*_{(t,Y),N} = 0$, $\delta^*_{(m,Y),Y} = 0$, $\delta^*_{(m,Y),N} = 1$, and $\delta^*_{N,Y} = 1$. Second, let P2 post a mechanism $\tilde{\gamma}_2^*$ in which $\sigma^*_{Y,Y} = \frac{1}{6}$ and $\sigma^*_{N,Y} = \sigma^*$. One can check that the subgame $(\gamma_1^*, \tilde{\gamma}_2^*)$ has an equilibrium in which A1 and A2 play $((t,Y)N, NY)$. This yields P2 a payoff $10(1-\sigma^*)$.

Suppose next that P2 deviates to some direct mechanism $\tilde{\gamma}_2$. The agents play the game in Table 5.

	YN	NY
$(t,Y)N$	$(10 - \frac{5}{16}(7+8\sigma^*), 8+4\sigma^*)$	$(10, 8+4\sigma_{N,Y})$
$(m,Y)N$	$(10, \frac{9}{2})$	$(5, 8+4\sigma_{N,Y})$
NY	$(9-\sigma_{Y,N}, 12)$	$(4+36\sigma_{Y,Y}, 13-6\sigma_{Y,Y})$

Table 5: The subgame $(\gamma_1^*, \tilde{\gamma}_2)$ of G^M.

The analysis of such subgames consists of three steps.

Step 1 Consider first the subgames $(\gamma_1^*, \tilde{\gamma}_2)$ for $\sigma_{Y,Y} \leq \frac{1}{6}$ and $\sigma_{N,Y} \geq \sigma^*$. Our candidate for an SPNE-PSP of G^M has A1 and A2 playing $((t,Y)N, NY)$ in any such subgame, which is indeed a Nash equilibrium if $\sigma_{Y,Y} \leq \frac{1}{6}$ and $\sigma_{N,Y} \geq \sigma^*$, because NY is then weakly dominant for A2 and $(t,Y)N$ is then a best response of A1 to A2 playing NY. The corresponding payoff for P2 is $10(1-\sigma_{N,Y})$, which is strictly decreasing in $\sigma_{N,Y}$. By construction, P2 does not want to deviate to a mechanism $\tilde{\gamma}_2$ such that $\sigma_{Y,Y} \leq \frac{1}{6}$ and $\sigma_{N,Y} > \sigma^*$, which would make her strictly worse off, or to a mechanism $\tilde{\gamma}_2$ such that $\sigma_{Y,Y} < \frac{1}{6}$ and $\sigma_{N,Y} = \sigma^*$, which would leave her indifferent.

Step 2 We next show that P2 does not want to deviate to any $\tilde{\gamma}_2$ such that $\sigma_{Y,Y} > \frac{1}{6}$. To see why, observe first that, if $\sigma_{Y,Y} > \frac{1}{6}$, then $(t,Y)N$ is strictly dominated in $(\gamma_1^*, \tilde{\gamma}_2)$ for A1 and $(\gamma_1^*, \tilde{\gamma}_2)$ has no pure strategy Nash equilibrium. It is easy to check that $(\gamma_1^*, \tilde{\gamma}_2)$ has a unique mixed-strategy Nash equilibrium in which A1 plays $(m,Y)N$ with probability p and NY with probability $1-p$, where

$$p \equiv \frac{6\sigma_{Y,Y} - 1}{6\sigma_{Y,Y} + 4\sigma_{N,Y} + \frac{5}{2}},$$

and A2 plays YN with probability q and NY with probability $1-q$, where

$$q \equiv \frac{36\sigma_{Y,Y} - 1}{36\sigma_{Y,Y} + \sigma_{Y,N}}.$$

Then P2's payoff is

$$pq\zeta + p(1-q)10(1-\sigma_{N,Y}) + (1-p)q(8+2\sigma_{Y,N}) + (1-p)(1-q)\zeta,$$

which is negative when the loss ζ is large enough as p is bounded away from 1 and q is bounded away from 0 and 1 no matter the values of $\sigma_{Y,Y} > \frac{1}{6}$ and $\sigma_{Y,N}$.

Step 3 We finally show that P2 does not want to deviate to a mechanism $\tilde{\gamma}_2$ such that $\sigma_{Y,Y} \leq \frac{1}{6}$ and $\sigma_{N,Y} < \sigma^*$. Consider first the subgames such that $\sigma_{Y,Y} = \frac{1}{6}$ and $\sigma_{N,Y} < \sigma^*$. Our candidate for an SPNE-PSP of G^M has A1 and A2 playing (NY, NY) in any such subgame, which is indeed a Nash equilibrium. The corresponding payoff for P2 is ζ, so that she has no incentive to deviate from her postulated mechanism. Consider next the subgames such that $\sigma_{Y,Y} < \frac{1}{6}$ and $\sigma_{N,Y} < \sigma^*$. Observe that none of the resulting subgames $(\gamma_1^*, \tilde{\gamma}_2)$ has a Nash equilibrium in which A2 plays a pure strategy. Thus, A1 must play $(t, Y)N$ with positive probability, for, otherwise, the unique best response of A2 would be NY. Moreover, A1 must play $(m, Y)N$ or NY with positive probability, for, otherwise, the unique best response of A2 would be YN. We distinguish three cases.

Case 1 Suppose first that A1 randomizes over $(t, Y)N$ and $(m, Y)N$. For A2 to be indifferent between YN and NY, it must be that A1 plays $(t, Y)N$ with probability p' and $(m, Y)N$ with probability $1 - p'$, where

$$p' \equiv \frac{\frac{7}{2} + 4\sigma_{N,Y}}{\frac{7}{2} + 4\sigma^*}.$$

Similarly, for A1 to be indifferent between $(t, Y)N$ and $(m, Y)N$, it must be that A2 plays YN with probability q' and NY with probability $1 - q'$, where

$$q' \equiv \frac{1}{1 + \frac{1}{16}(7 + 8\sigma^*)}.$$

These strategies form a Nash equilibrium if A1 is not tempted to deviate to NY, which is the case if and only if $q' \geq q$, with q as defined in Step 2. Then P2's payoff is

$$q'\zeta + p'(1-q')10(1-\sigma_{N,Y}) + (1-p')(1-q')\zeta,$$

which is negative when the loss ζ is large enough as q' is bounded away from 0 no matter the value of σ^*.

Case 2 Suppose next that A1 randomizes over $(t, Y)N$ and NY. For A2 to be indifferent between YN and NY, it must be that A1 plays $(t, Y)N$ with probability p'' and NY with probability $1-p''$, where

$$p'' \equiv \frac{1-6\sigma_{Y,Y}}{1-6\sigma_{Y,Y}+4(\sigma^*-\sigma_{N,Y})}.$$

Similarly, for A1 to be indifferent between $(t,Y)N$ and NY, it must be that A2 plays YN with probability q'' and NY with probability $1-q''$, where

$$q'' \equiv \frac{6(1-6\sigma_{Y,Y})}{6(1-6\sigma_{Y,Y})+\frac{1}{16}(19+40\sigma^*)-\sigma_{Y,N}}.$$

These strategies form a Nash equilibrium if A1 is not tempted to deviate to $(m,N)Y$, which is the case if and only if $q \geq q'$, with q as defined in Step 2. Then P2's payoff is

$$p''q''\zeta + p''(1-q'')10(1-\sigma_{N,Y}) + (1-p'')q''(8+2\sigma_{Y,N}) + (1-p'')(1-q'')\zeta,$$

and we must show that, when the loss ζ is large enough, this does not exceed the candidate equilibrium payoff $10(1-\sigma^*)$ uniformly in $\sigma_{Y,Y} < \frac{1}{6}$, $\sigma_{N,Y} < \sigma^*$, and $\sigma_{Y,N}$, given the above expressions for p'' and q''. For conciseness, let us write

$$\begin{aligned}
\varepsilon &\equiv 1-6\sigma_{Y,Y},\\
\eta &\equiv \sigma^*-\sigma_{N,Y},\\
\xi &\equiv \frac{1}{16}(19+40\sigma^*)-\sigma_{Y,N},
\end{aligned}$$

and notice that ξ, unlike ε and η, is bounded away from 0. An upper bound for P2's payoff from deviating is

$$B \equiv [p''q''+(1-p'')(1-q'')]\zeta + p''(1-q'')10(1-\sigma_{N,Y}) + 10(1-p'')q'',$$

and, with the above notation, we have

$$\begin{aligned}
B-10(1-\sigma^*) &\leq [p''q''+(1-p'')(1-q'')]\zeta + p''(1-q'')10(\sigma^*-\sigma_{N,Y}) + 10(1-p'')q''\\
&\propto (6\varepsilon^2+4\eta\xi)\zeta + \varepsilon\xi\eta + 240\eta\varepsilon\\
&< \eta(4\xi\zeta+\varepsilon\xi+240\varepsilon),
\end{aligned}$$

which, as ξ is bounded away from 0 and $\varepsilon \in [0,1]$ as $\sigma_{Y,Y} < \frac{1}{6}$, is negative when the loss ζ is large enough, uniformly in η and ε.

Case 3 Finally, if $q = q' = q''$, A1 is ready to randomize over $(t,Y)N$, $(m,Y)N$, and NY. Instead of considering a completely mixed Nash equilibrium of $(\gamma_1^*, \tilde{\gamma}_2)$, we can however select either of the equilibria constructed in Cases 1 and 2, which ensures that P2 has no incentive to deviate. Hence the result. ∎

The proof of Proposition 2 shows how P1 uses her communication with A1 to construct additional threats. This flexibility is exploited to deter P2's deviations and to support a large number of equilibrium allocations. The example shares with Yamashita (2010) the idea that, following any deviation of P2, P1 delegates to the agents the implementation of her actions. In particular, it should be noted that selecting the appropriate punishments for P2 is part of the agents' equilibrium strategies. However, this effect does not transit in our example through the selection of a particular action of P1, as there are no direct payoff externalities between the principals for any given profile of participation decisions for the agents. Rather, punishments take place through the agents' participation decisions and, specifically, through their willingness to take such decisions at random given the mechanism posted by P1. This, from the perspective of P2, creates the possibility of a miscoordination whereby both A1 and A2, or neither of them, participate with her, two outcomes that are equally detrimental to her. As a consequence, equilibrium multiplicity obtains without having to assume at least three agents, which would be necessary to punish deviators via majority voting as in the recommendation mechanisms of Yamashita (2010). It should also be noted that the example could be made generic by allowing for direct externalities between principals, as long as the game is perturbed so as to make punishments more severe; in fact, direct externalities make it easier to engineer such punishments. Such a perturbation can be designed to preserve 10 as the min-max-min payoff for P2 and to make sure that all payoffs for P2 between 0 and 10 can be supported in equilibrium.

The key insight for equilibrium characterization is that, in competing-mechanism games in which agents can take real participation decisions, instead of mere reporting decisions, the min-max-min payoff (1) is not relevant for describing principals' equilibrium payoffs. Indeed, Proposition 2 shows that even minimal communication allows one to support equilibrium payoffs for P2 below this value, which is equal to 10 in our example. We further discuss this issue in the next section.

3.4 Discussion

We first relate our findings to the failure of the min-max theorem in our example. We next argue that Proposition 2 is robust to an alternative extensive form due to Szentes (2009).

3.4.1 Min-Max > Max-Min

In analogy with (1), let us define the max-min-min payoff for P2 as follows:

$$\underline{V}_2 \equiv \sup_{\tilde{\gamma}_2 \in \tilde{\Gamma}_2} \min_{\tilde{\gamma}_1 \in \tilde{\Gamma}_1} \min_{\tilde{\lambda} \in \Lambda^*(\tilde{\gamma})} \mathbf{E}_{z(\tilde{\gamma},\tilde{\lambda})}[v_2(a,x)]. \tag{2}$$

It will become clear below why, in general, the max in (1) may have to be replaced by a sup in (2). Our discussion revolves around the following lemma, for which we provide a proof in the Appendix.

Lemma 1 $\overline{V}_2 > \underline{V}_2$.

To see why this holds, notice that, by Proposition 1, we have $\overline{V}_2 = 10$, while, by Proposition 2, we have

$$\sup_{\tilde{\gamma}_2 \in \tilde{\Gamma}_2} \min_{\gamma_1 \in \Gamma_1^{M_1}} \min_{\lambda \in \Lambda^*(\gamma_1, \tilde{\gamma}_2)} \mathbf{E}_{z(\gamma_1, \tilde{\gamma}_2, \lambda)}[v_2(a, x)] \leq 0. \tag{3}$$

Thus, to show that $\underline{V}_2 < \overline{V}_2$, we only need to establish that the payoffs in (2)–(3) coincide. We provide a proof of this fact in the appendix. The wedge between the max-min-min and min-max-min payoffs $\underline{V}_2$ and $\overline{V}_2$ reflects the fact that agents take strategic participation decisions. This contrasts with the min-max theorem for one-shot complete-information games established by Ghosh and Han (2018, Theorem 3) in line with the general discussions of Peters and Troncoso-Valverde (2013) and Peters (2014).[3]

There are indeed two reasons why Sion's (1958) min-max theorem does not apply in the present context. The first is that the Nash correspondence $\Lambda^* : \tilde{\Gamma}_1 \times \tilde{\Gamma}_2 \twoheadrightarrow \Delta(A^1) \times \Delta(A^2)$ in our example is not single-valued and is only upper hemicontinuous. This implies that the mapping $(\tilde{\gamma}_1, \tilde{\gamma}_2) \mapsto \min_{\tilde{\lambda} \in \Lambda^*(\tilde{\gamma})} \mathbf{E}_{z(\tilde{\gamma}, \tilde{\lambda})}[v_2(a, x)]$ is not upper semicontinuous in $\tilde{\gamma}_2$;[4] for instance, it is easy to check that it exhibits a downward discontinuity at $\sigma_{Y,Y} = \frac{1}{6}$.[5] The second is that, whereas the expectation $\mathbf{E}_{z(\tilde{\gamma}, \tilde{\lambda})}[v_2(a, x)]$ is multilinear in $(\tilde{\gamma}, \tilde{\lambda})$, there is no reason for its minimum with respect to $\tilde{\lambda} \in \Lambda^*(\tilde{\gamma})$ to be quasiconvex in $\tilde{\gamma}_1$ and quasiconcave in $\tilde{\gamma}_2$. Thus letting the agents take real participation decisions naturally leads to discontinuities and nonconvexities that prevent the usual min-max logic from applying, even if we allow principals to randomize over their actions. Given the importance of participation decisions in the applied literature on competing mechanisms, this casts serious doubt on the general relevance of this logic.

[3] In two-principal settings, Theorem 3 in Ghosh and Han (2018) follows directly from Sion's (1958) min-max theorem.

[4] By contrast, it follows from the second half of Berge's maximum theorem that this mapping is lower semicontinuous in $\tilde{\gamma}_1$ (Aliprantis and Border (2006, Lemma 17.30)), and thus attains its minimum over the compact set $\tilde{\Gamma}_1$. This minimum, however, need not be upper semicontinuous in $\tilde{\gamma}_2$, and thus may not attain its maximum over the compact set $\tilde{\Gamma}_2$. Hence the sup instead of the max in the definition of $\underline{V}_2$.

[5] Indeed, if $\sigma_{Y,Y} = \frac{1}{6}$, (NY, NY) is a pure-strategy Nash equilibrium in the game played by the agents that yields P2 a payoff of ζ. By contrast, if $\sigma_{Y,Y} < \frac{1}{6}$ and, for instance, $\delta_{Y,Y} = 1$ and $\delta_{Y,N} = \delta_{N,Y} = 0$, then the game between the agents admits only a completely mixed Nash equilibrium in which the probability p''' with which A1 plays YN is bounded away from 0, and the probability q''' with which A2 plays YN tends to 0 as $\sigma_{Y,Y}$ tends to $\frac{1}{6}$. The limit payoff for P2 is then $(1 - p''')\zeta + p'''10(1 - \sigma_{N,Y}) > \zeta$.

Admittedly, Proposition 2 does not provide a full characterization of equilibrium payoffs for principals; specifically, we have not checked whether $\underline{V}_2 < 0$ and whether equilibria with payoff for P2 down to that level can be sustained. Our second example will show that there is no hope for obtaining a general result along these lines: neither the min-max-max payoff (1) nor the max-min-min payoff (2) are relevant bounds for the characterization of equilibrium payoffs of competing-mechanism games with communication.

3.4.2 An Alternative Extensive Form

By showing the existence of equilibrium allocations yielding a principal a payoff below her min-max-min payoff, Proposition 2 stands in contrast with the folk-theorem results derived by Yamashita (2010) and Peters and Troncoso-Valverde (2013). Key to our results is the existence of a wedge between principals' min-max-min and max-min-min payoffs, as emphasized above. In this context, one may, however, argue with Szentes (2009) that a folk theorem relying on min-max-min bounds could still be established by considering more general mechanisms, whereby principals do not fully delegate the implementation of their final actions to the agents.

Szentes (2009) models partial delegation via a slightly different extensive form. First, each principal posts a mechanism specifying a menu of final actions contingent on any profile of messages she may receive. Second, each agent privately sends messages to each principal. Third, each principal selects an action from the resulting menu. Theorem 1 in Szentes (2009) states that, in this game, *in which the agents' participation decisions are not modeled*, an allocation can be supported in equilibrium if and only if it yields each principal a payoff above her min-max payoff, computed over the principals' final actions. It should be noted that lotteries are not allowed in this construction.

We now show that, when agents' participation decisions are explicitly taken into account, Proposition 2 is robust to partial delegation. The extended competing-mechanism game $G^{M,S}$ unfolds in four stages:

1. The principals simultaneously post extended mechanisms that specify menus of lotteries over final actions contingent on any consistent communication and participation profile.

2. The agents simultaneously take consistent communication and participation decisions.

3. The principals simultaneously select lotteries from the menus determined by their mechanisms and the agents' decisions.

4. Lotteries realize, and all payoffs accrue.

In this framework, an extended mechanism for principal j is a measurable correspondence $\gamma_j^S : C_j \twoheadrightarrow \Delta(X_j)$ that associates a menus of lotteries over final actions to any profile of consistent communication and participation decisions by the agents, and an extended direct mechanism is a correspondence $\tilde{\gamma}_j^S : A_j \twoheadrightarrow \Delta(X_j)$. To eschew trivial nonexistence problems, we suppose that these correspondences are compact-valued.

We now argue that the allocations characterized in Proposition 2 can also be supported in an equilibrium of this enlarged game. On the equilibrium path, P1 posts γ_1^* for some $\sigma^* \in [0,1]$, while P2 posts $\tilde{\gamma}_2^*$, as assumed in the proof of Proposition 2. Now, suppose that P2 deviates and posts an extended direct mechanism $\tilde{\gamma}_2^S$; that is, we assume, as in Proposition 2, that agents cannot send messages to P2. Then the agents' payoffs in the subgame $(\gamma_1^*, \tilde{\gamma}_2^S)$ are as in Table 5, where $\sigma_{a_2^1, a_2^2}$ now denotes the agents' equilibrium belief that P2 will choose x_{21} given a participation profile $(a_2^1, a_2^2) \in A_2$ for P2. In particular, these beliefs are now constrained to be consistent with P2's sequential rationality in the relevant subgames; for instance, it is clear from Table 1 that we must have $\sigma_{Y,N} = 0$ and $\sigma_{N,Y} = 1$ in the full-discretion case where $\tilde{\gamma}_2^S$ does not constrain the set of lotteries available to P2. As the analysis in Proposition 2 shows that P2 cannot profitably deviate when she can commit to any profile $(\sigma_{Y,Y}, \sigma_{Y,N}, \sigma_{N,Y}, \sigma_{N,N})$, it a fortiori implies that P2 cannot profitably deviate to an extended direct mechanism, which may not allow for such a commitment. This shows that any payoff between 0 and 10 for P2 can be supported in equilibrium in $G^{M,S}$. In particular, equilibrium payoffs for P2 below her min-max payoff can be supported under partial delegation. The logic of Szentes' (2009) Theorem 1 breaks down because P2 is constrained by the equilibrium of the game played by A1 and A2.

4 The Second Example

We now argue that the combined role of agents' participation and communication decisions fundamentally challenges the logic of folk-theorem results. Specifically, we show that, when communication is tied to participation, it may be impossible to construct equilibrium threats based on the agents' incentives to detect a deviation and report it to nondeviating principals. As a consequence, a folk theorem may not obtain even if the min-max bounds are specified in terms of arbitrary message spaces for the agents. Our second example is an instance of the competing-hierarchy model introduced by Myerson (1982), and first developed by Martimort (1996) and Gal-Or (1997).[6]

[6] In competing-hierarchy models, agents' participation decisions are not explicitly modeled, and hierarchies are taken as exogenous. This can be seen as a particular case of our competing-mechanism model, in which

4.1 The Physical Environment

Let $J \equiv 2$ and $I \equiv 3$, and let $X_1 \equiv \{x_{11}, x_{12}, x_{13}\}$ and $X_2 \equiv \{x_{21}, x_{22}, x_{23}\}$ be the sets of actions of principal 1 (P1) and principal 2 (P2), respectively. The agents' payoffs are as follows. Each agent's payoff is independent of the other agents' participation decisions. Moreover, agent 1 (A1) strictly prefers to participate with P1, while agent 2 (A2) and agent 3 (A3) strictly prefer to participate with P2. Thus, this setting can be interpreted as a competitive game played between the hierarchy formed by P1 and A1, and that formed by P2, A2, and A3, in which principals' payoffs are not affected by agents' participation decisions. The matrix in Table 6 below represents the payoffs to P1 and P2, respectively.

	x_{21}	x_{22}	x_{23}
x_{11}	$(1,1)$	$(-1,2)$	$(0,0)$
x_{12}	$(1,-1)$	$(-1,0)$	$(0,0)$
x_{13}	$(0,0)$	$(0,0)$	$(0,0)$

Table 6: The payoff matrix for principals.

There is no need to fully specify the agents' payoffs. Indeed, for the sake of our argument, it is enough to require that A1 payoffs are such that

$$u^1(x_{11}, x_{22}) > \max\{u^1(x_{12}, x_{22}), u^1(x_{13}, x_{22})\}, \tag{4}$$

so that, if P2 chooses x_{22}, A1 would strictly prefer P1 to choose x_{11} rather than x_{12} or x_{13}.

4.2 Min-Max = Max-Min

Unlike in our first example, we assume from the outset that the agents can communicate with the principals through rich message spaces, as done by Yamashita (2010). The following min-max theorem is an easy consequence of the fact that agents' actions have no impact on principals's payoffs.

Lemma 2 *For each $j = 1, 2$, and for arbitrary message spaces M^i_j,*

$$V_j \equiv \min_{\gamma_{-j} \in \Gamma_{-j}^{M_{-j}}} \max_{\gamma_j \in \Gamma_j^{M_j}} \min_{\lambda \in \Lambda^*(\gamma)} \mathbf{E}_{z(\gamma,\lambda)}[v_j(a,x)] = \max_{\gamma_j \in \Gamma_j^{M_j}} \min_{\gamma_{-j} \in \Gamma_{-j}^{M_{-j}}} \min_{\lambda \in \Lambda^*(\gamma)} \mathbf{E}_{z(\gamma,\lambda)}[v_j(a,x)] = 0. \tag{5}$$

Proof. Because principals' payoffs do not depend on agents' participation decisions, we can disregard the $\min_{\lambda \in \Lambda^*(\gamma)}$ operator in the min-max-min and max-min-min payoffs in (5). To

each agent has a dominant participation strategy.

show that they are equal to 0, observe that, for any mechanism $\gamma_{-j} \in \Gamma_{-j}^{M_{-j}}$, principal j can defend a 0 payoff by committing to choose x_{j3} with probability 1, regardless of the agents' participation and communication decisions. Likewise, principal $-j$ can bring down principal j's payoff to 0 by committing to choose x_{-j3} with probability 1, regardless of the agents' participation and communication decisions. The result follows. ■

An alternative, if less direct, proof of this result consists to notice that, because the sets of direct mechanisms $\tilde{\Gamma}_1$ and $\tilde{\Gamma}_2$ are compact and convex subsets of a Euclidean space and the payoff function $\mathbf{E}_{z(\tilde{\gamma},\tilde{\lambda})}[v_j(a,x)]$ is linear in $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$, it follows from Sion's (1958) min-max theorem that

$$\min_{\tilde{\gamma}_{-j}\in\tilde{\Gamma}_{-j}} \max_{\tilde{\gamma}_j\in\tilde{\Gamma}_j} \min_{\tilde{\lambda}\in\Lambda^*(\tilde{\gamma})} \mathbf{E}_{z(\tilde{\gamma},\tilde{\lambda})}[v_j(a,x)] = \max_{\tilde{\gamma}_j\in\tilde{\Gamma}_j} \min_{\tilde{\gamma}_{-j}\in\tilde{\Gamma}_{-j}} \min_{\tilde{\lambda}\in\Lambda^*(\tilde{\gamma})} \mathbf{E}_{z(\tilde{\gamma},\tilde{\lambda})}[v_j(a,x)],$$

and then to apply Ghosh and Han (2018, Lemma 1). This expresses the fact that, if the payoffs (1)–(2)—which are specified only in terms of the primitives of the model—coincide, no harsher punishment can be forced upon any principal by letting the other principal post an indirect mechanism.

4.3 Failure of the Folk Theorem

In contrast with Lemma 1, Lemma 2 suggests that our second example should a priori be a more promising setting for the min-max logic to apply and thus for deriving a folk theorem. Indeed, a mechanical application of Yamashita's (2010) and Peters and Troncoso-Valverde's (2013) results would yield that any incentive-feasible allocation in which each principal obtains a payoff above 0 can be supported in an SPNE-PSP of a game G^M with sufficiently rich message spaces; for instance, the message spaces may be large enough as to incorporate all direct mechanisms as in Yamashita (2010), allowing principals to use recommendation mechanisms. Yet this approach would be mislead, as we shall now see.

To first illustrate this point in an intuitive manner, notice that the only way P1 and P2 can reach payoffs $(1, 1)$ is by deterministically choosing the actions (x_{11}, x_{21}). Suppose then, by way of contradiction, that there exists an SPNE-PSP (γ^*, λ^*) of some game G^M in which P1 and P2 choose these actions, and only these actions, on the equilibrium path. For this to be the case, there must exist a message $m_1^1 \in M_1^1$ such that P1 chooses x_{11} upon receiving it from A1—recall that, by assumption, A1 strictly prefers to participate with P1, while A2 and A3 strictly prefer to participate with P2. But P2 could then deviate to a mechanism γ_2 committing her to choose x_{22} regardless of the messages she receives from A2 and A3. Because of (4), A1 would then strictly prefer that P1 chooses x_{11}. In the subgame (γ_1^*, γ_2), he can ensure this by sending P1 the message m_1^1. This yields P2 a payoff of $2 > 1$, and hence γ_2 is a profitable deviation for P2. Intuitively, to prevent such a deviation one would need to implement x_{12} or x_{13} in P1's mechanism off the equilibrium path. This would indeed be possible if P1 were able to communicate with all three agents, committing to choose any of the above punishments when reported by a majority of them, as done by Yamashita (2010) with recommendation mechanisms. However, as long as participation and communication are tied, this possibility cannot be guaranteed. In the example, the implementation of the punishment conflicts with the incentives of A1, who is the only one communicating with P1. Thus there is no SPNE-PNP of G^M supporting the payoffs $(1, 1)$ for P1 and P2. The following result, for which we provide a proof in the Appendix, generalizes this insight to all payoffs above $(0, 0)$ for P1 and P2, even if they play mixed strategies in equilibrium.

Proposition 3 *The principals obtain payoffs* $(0, 0)$ *in any SPNE of any game* G^M.

It should be noted that Proposition 3 does not depend on the size of the message spaces. In particular, it holds true if these are large enough to incorporate the infinite-dimensional space of direct mechanisms, allowing principals to use recommendation mechanisms. This finding contrasts with the folk theorems of Yamashita (2010) and Peters and Troncoso-Valverde (2013), in which each agent participates and communicates with all the principals. Finally, it should be noted that our second example is generic in that payoffs in Table 6 can be modified as long as no additional threats are generated, which guarantees equilibrium uniqueness. Similarly, it is immediate to see that all preference orderings for A1 consistent with (4) lead to the same result. In particular, we can freely disturb A1's preference ordering over P1's actions when P2 plays x_{21} or x_{23}.

5 Communication without Participation

Our examples have lead us to question the logic of folk theorems for competing-mechanism games. We now revisit the reasons why this logic fails to hold, and identify an alternative communication protocol under which a folk theorem can be established.

Consider first our second example, which illustrates a situation in which no payoff for the principals above Yamashita's (2010) min-max-min bound can be supported in equilibrium even when agents can send messages from arbitrarily rich message spaces. As discussed in Section 4, the assumption driving this result is that an agent can communicate with a principal only if he chooses to participate with her. It is thus natural to ask under which additional conditions a folk theorem would obtain in our general setting. In this respect, it should be noted that folk theorems in the spirit of Yamashita (2010) and Peters and Troncoso-Valverde (2013) require that, following any deviation by a principal, at least three agents communicate with each nondeviating principal, which guarantees that the deviator can be punished by implementing actions in her competitors' mechanisms that minimize her payoff. This turns out to be very demanding when communication is tied to participation. Indeed, it is problematic to identify non ad hoc conditions on primitives guaranteeing that agents' equilibrium strategies have this feature. Besides, under exclusive competition, this requirement imposes that there be at least $3J$ agents interacting with J principals, a condition that may fail to be satisfied in some applications.

A more direct way to retrieve a folk theorem in competing-mechanism games may hence consist in suppressing any restriction on communication, as in Yamashita (2010) and Peters and Troncoso-Valverde (2013). Even then, however, a distinguishing feature of our setting remains that agents take payoff-relevant participation decisions. We now show, in the context of our first example, that this enables us to sustain equilibria yielding a principal a payoff below her min-max-min payoff *despite communication being unrestricted.* That is, Proposition 2 is robust to the specific communication protocol considered. To see this, let us now allow A1 to send the message m to P1 even when he does not participate with her. This modifies the subgame played by the agents following P2's deviation to $\tilde{\gamma}_2$ only in that an additional line must be included in Table 4:

	YN	NY
$(t,Y)N$	$(10-5\delta_{(t,Y),Y}, \frac{9}{2}+8\delta_{(t,Y),Y})$	$(10-5\delta_{(t,Y),N}, 8+4\sigma_{N,Y})$
$(m,Y)N$	$(10-5\delta_{(m,Y),Y}, \frac{9}{2}+8\delta_{(m,Y),Y})$	$(10-5\delta_{(m,Y),N}, 8+4\sigma_{N,Y})$
$(t,N)Y$	$(9-\sigma_{Y,N}, 8+4\delta_{(t,N),Y})$	$(4+36\sigma_{Y,Y}, 13-6\sigma_{Y,Y})$
$(m,N)Y$	$(9-\sigma_{Y,N}, 8+4\delta_{(m,N),Y})$	$(4+36\sigma_{Y,Y}, 13-6\sigma_{Y,Y})$

Table 7: The subgame $(\gamma_1^*, \tilde{\gamma}_2)$ of G^M under unrestricted communication.

Letting P1 post the mechanism γ_1^* as in Proposition 2 with the additional feature that $\delta^*_{(t,N),Y} = \delta^*_{(m,N),Y} = 1$, we can then mimic the proof of Proposition 2 to reach the same conclusion. Thus, making agents' communication opportunities independent of their participation decisions is not sufficient to reformulate a folk theorem in terms of the principals' min-max-min payoffs.

6 Concluding Remarks

Our results question the logic of folk theorems for competing-mechanism games in two ways.

Our first example illustrates a situation in which principals' payoffs below the min-max-min bounds identified by Peters and Troncoso-Valverde (2013), Peters (2014), and Ghosh and Han (2018) can be supported in equilibrium. Our second example illustrates a situation in which no payoff for the principals above the min-max-min bound identified by Yamashita (2010) can be supported in equilibrium even if agents can choose messages from arbitrarily rich message spaces. These examples highlight the two key reasons why the folk-theorem logic fails to hold. In the first example, agents' participation decisions are payoff-relevant. In the second example, an agent can communicate with a principal only if he chooses to participate with her. Both features are prominent in economic applications of competing-mechanism games.[7] We have argued that these insights are robust to slight perturbations of players' payoffs and to alternative extensive forms and communication protocols.

As pointed out in the Introduction, our model shares several building blocks with the competing-auction and competitive-search literatures: namely, principals post mechanisms, after which each agent can subsequently participate and communicate with at most one principal. Admittedly, in competing auctions, sellers compete to trade with buyers who are privately informed of their valuations. We do expect that our results, established under complete information, would carry over to such incomplete-information settings. A key step for such an extension would be to develop an adequate formulation for the bound on principals' payoffs, a task that may not be obvious under incomplete information (Peters (2014)). By contrast, competitive-search models typically postulate complete information—as in the standard case where workers apply for the vacancies and wages offered by competitive firms (Wright, Kircher, Julien, and Guerrieri (2019))—and thus seem closer to a literal interpretation of our model. The primary focus of this literature, however, is on symmetric equilibria in which agents play a mixed strategy, an approach often motivated by referring to the difficulty of coordinating agents in large anonymous markets. Whereas our general model allows us to consider similar situations, the examples we have developed crucially exploit heterogeneity among principals and/or agents. It thus remains an open question whether similar examples could be designed under the stricter conditions on environments and equilibria postulated in this literature.

A common feature of applied competing-auction and competitive-search models is the restriction to direct mechanisms. That is, each seller only requires the buyers who participate with her to submit their exogenous private information. Under complete information, this

[7] Clearly, our results do not contradict the possibility of multiple equilibrium allocations arising in a specific class of applications. In this respect, Han (2016) shows how indirect mechanisms can be used to sustain multiple symmetric equilibria in a simple competitive-search setting.

amounts to each seller posting take-it-or-leave-it offer for the buyers participating with her. The difficulties in establishing a general folk theorem that we have highlighted may lead one to ask whether, under exclusive competition, direct mechanisms achieve a full characterization of equilibria. Our first example provides a clear negative answer. Indeed, an immediate implication of Propositions 1–2 is the existence of many equilibrium outcomes that cannot be supported by direct mechanisms. Thus our first example documents a failure of the revelation principle in competing-mechanism games of exclusive competition: direct mechanisms are not flexible enough to reproduce all the threats that P1 can implement using the payoff-irrelevant message m.[8] This failure is dramatic: while the game $\tilde{G}$ has a unique equilibrium payoff vector, a continuum of Pareto-ranked equilibria can be sustained in the game G^M as soon as a single agent has the opportunity to send a single additional message to a single principal.

To interpret this failure of the revelation principle under exclusive competition, it is helpful to contrast our first example with related results in the literature. First, the example is cast in a complete-information framework; this contrasts with Martimort (1996), Peck (1997), and Attar, Campioni, and Piaser (2018), who crucially exploit the agents' private information. Second, the example does not rely on direct externalities between principals, but instead exploits the strategic role of agents' participation and communication decisions; hence it does not rely on the intuitions developed by Martimort (1996) in the context of competing hierarchies. Third, principals in the example play pure strategies in equilibrium; hence we do not rely on the limited power of direct mechanisms to extract the agents' information on the realization of principals' mixed strategies, in contrast with Peck (1997). Thus, in a sense, we have put ourselves in the worst possible scenario for communication to play a role: agents have no private information of their own and each principal chooses her equilibrium mechanism deterministically, so that there is no need to use the agents to reveal it to the other principal. Despite these drastic features, our example shows the potentially destabilizing role of communication even in complete-information environments. This again questions the prevalent use of direct mechanisms in the applied literature.

Taken together, our examples emphasize that disregarding agents' participation decisions in competing-mechanism games entails a severe loss of generality. Observe, in this respect, that our results do not depend on the specific assumption of exclusive competition. In

[8] Thus Propositions 1–2 extend to exclusive-competition environments the insights developed by Peters (2001) and Martimort and Stole (2002) in single-agent, nonexclusive-competition environments. The example suggests that, to implement all relevant threats, there is no need for each agent to communicate with all principals; this ex-post vindicates our assumption that communication is tied to participation.

particular, the second example can be reformulated in a nonexclusive scenario where each agent can participate with more than one principal at a time, provided that communication is tied to participation. The explicit consideration of agents' participation decisions introduces a fundamental constraint on principals' design of mechanisms. When designing a mechanism, a principal anticipates that her payoff will be affected by agents' participation decisions, which in turn depend on the entire profile of posted mechanisms. Dealing with this additional moral-hazard dimension may require a more sophisticated class of mechanisms, in which a principal can send private signals to agents. In recent work, Attar, Campioni, and Piaser (2019) provide an example of a complete-information competing-mechanism game in which principals use such private communication to correlate their actions with the agents' decisions in equilibrium. The resulting set of equilibrium allocations and the set of equilibrium allocations that can be supported with standard recommendation mechanisms turn out to be disjoint. All these insights point towards the need to develop new ideas and devices to approach equilibrium characterization in competing-mechanism games.

Appendix

Proof of Lemma 1. This is in fact an instance of a general result in the spirit of Myerson's (1982) revelation principle. The intuition is that, for any direct mechanism $\tilde{\gamma}_2$ of P2 and any mechanism γ_1 of P1, and for any Nash equilibrium λ of the subgame $(\gamma_1, \tilde{\gamma}_2)$ of G^M, there exist a direct mechanism $\tilde{\gamma}_1$ of P1 and a Nash equilibrium $\tilde{\lambda}$ of the subgame $\tilde{\gamma} \equiv (\tilde{\gamma}_1, \tilde{\gamma}_2)$ of $\tilde{G}$ such that the two resulting allocations coincide, $z(\tilde{\gamma}, \tilde{\lambda}) = z(\gamma_1, \tilde{\gamma}_2, \lambda)$. Indeed, P1 can reproduce the randomizations over messages performed by A1 in $\lambda^1(\gamma_1, \tilde{\gamma}_2)$ in case he decides to participate with P1 by offering a direct mechanism $\tilde{\gamma}_1$ such that

$$\tilde{\gamma}_1(Y, a_1^2) \equiv \frac{\lambda^1(\gamma_1, \tilde{\gamma}_2)((t, Y)N)\gamma_1((t, Y), a_1^2) + \lambda^1(\gamma_1, \tilde{\gamma}_2)((m, Y)N)\gamma_1((m, Y), a_1^2)}{\lambda^1(\gamma_1, \tilde{\gamma}_2)((t, Y)N) + \lambda^1(\gamma_1, \tilde{\gamma}_2)((m, Y)N)}$$

for all $a_1^2 \in A_1^2$. That is, if A1 chooses to participate with P1, then P1 first draws a lottery with outcomes (t, Y) and (m, Y), with probabilities $\frac{\lambda^1(\gamma_1,\tilde{\gamma}_2)((t,Y)N)}{\lambda^1(\gamma_1,\tilde{\gamma}_2)((t,Y)N)+\lambda^1(\gamma_1,\tilde{\gamma}_2)((m,Y)N)}$ and $\frac{\lambda^1(\tilde{\gamma}_1,\tilde{\gamma}_2)((m,Y)N)}{\lambda^1(\tilde{\gamma}_1,\tilde{\gamma}_2)((t,Y)N)+\lambda^1(\tilde{\gamma}_1,\tilde{\gamma}_2)((m,Y)N)}$, respectively, and then, depending on the outcome of this lottery, chooses x_{11} with probability $\gamma_1((t, Y), a_1^2)$ or $\gamma_1((m, Y), a_1^2)$. If A1 chooses not to participate with P1, then we set

$$\tilde{\gamma}_1(N, a_1^2) \equiv \gamma_1((\emptyset, N), a_1^2)$$

for all $a_1^2 \in A_1^2$. Turning to the agent's strategies, we define

$$\begin{aligned}
\tilde{\lambda}^1(\tilde{\gamma}_1, \tilde{\gamma}_2)(YN) &\equiv \lambda^1(\gamma_1, \tilde{\gamma}_2)((t, Y)N) + \lambda^1(\gamma_1, \tilde{\gamma}_2)((m, Y)N), \\
\tilde{\lambda}^1(\tilde{\gamma}_1, \tilde{\gamma}_2)(NY) &\equiv \lambda^1(\gamma_1, \tilde{\gamma}_2)((\emptyset, N)Y), \\
\tilde{\lambda}^2(\tilde{\gamma}_1, \tilde{\gamma}_2)(YN) &\equiv \lambda^2(\gamma_1, \tilde{\gamma}_2)(YN).
\end{aligned}$$

By construction, $z(\tilde{\gamma}, \tilde{\lambda}) = z(\gamma_1, \tilde{\gamma}_2, \lambda)$. Moreover, because P1 reproduces the randomizations of A1 in case he decides to participate with her, the incentives of the agents are unchanged. Hence $\tilde{\lambda}$ is a Nash equilibrium of the subgame $\tilde{\gamma}$, as required. The result follows. ■

Proof of Proposition 3. Suppose, by way of contradiction, that there exists an SPNE of some game G^M in which at least one of the principals obtains a positive payoff. The proof consists of two steps.

Step 1 We first claim that, on the candidate SPNE equilibrium path, P1 chooses x_{11} with positive probability. Suppose, indeed, that P1 never chooses x_{11} in the SPNE under consideration. Then P2 must obtain a 0 payoff, which she can guarantee by committing to choose x_{23} with probability 1, regardless of the agents' communication and participation decisions. By assumption, P1 must, therefore, obtain a positive payoff. If P1 never chooses x_{11}, this can occur if and only if P1 and P2 choose (x_{12}, x_{21}) with positive probability. But P2 would then obtain a negative payoff, a contradiction. The claim follows.

Step 2 We next claim that P2's unique best response in the candidate SPNE consists in committing to play x_{22} with probability 1, regardless of the agents' communication and participation decisions. We distinguish two types of subgames, depending on the mechanism posted by P1 according to her—possibly mixed—equilibrium strategy $\mu_1^* \in \Delta(\Gamma_1^{M_1})$.

Case 1 Consider first the mechanisms γ_1 in the support of μ_1^* such that there is no message $m_1^1 \in M_1^1$ such that P1 chooses x_{11} with positive probability following the consistent profile $(m_1^1, \emptyset, \emptyset, Y, N, N)$; call Γ_1 the corresponding set of mechanisms. Committing herself to choose x_{22} with probability 1, regardless of the agents' communication and participation decisions, ensures P2 to obtain her maximal payoff of 0 in any subgame in which P1 posts a mechanism in Γ_1.

Case 2 Consider next the mechanisms γ_1 in the support of μ_1^* such that there is a message $m_1^1 \in M_1^1$ such that P1 chooses x_{11} with positive probability following the consistent profile $(m_1^1, \emptyset, \emptyset, Y, N, N)$; that is, $\gamma_1 \in \operatorname{supp} \mu_1^* \setminus \Gamma_1$. It follows from (4) that A1 will, among these messages, choose one that maximizes the probability $\gamma_1(m_1^1, \emptyset, \emptyset, Y, N, N)(x_{11})$ of P1 choosing x_{11}. Committing herself to choose x_{22} with probability 1, regardless of the agents' communication and participation decisions, ensures P2 to obtain her maximal payoff of $2 \max_{m_1^1 \in M_1^1} \{\gamma_1(m_1^1, \emptyset, \emptyset, Y, N, N)(x_{11})\}$ in any subgame in which P1 posts a mechanism $\gamma_1 \in \operatorname{supp} \mu_1^* \setminus \Gamma_1$.

According to Step 1, $\mu_1^*(\operatorname{supp} \mu_1^* \setminus \Gamma_1) > 0$, so that Case 2 arises with positive probability. The claim follows.

Step 3 According to Steps 1 and 2, in the candidate SPNE, P1 chooses x_{11} with positive probability while P2 chooses x_{22} with probability 1. Hence P1 must earn a negative payoff. This, however, is a contradiction, as she can guarantee herself a 0 payoff by committing to choose x_{13} with probability 1. Hence the result. ■

References

[1] Aliprantis, C.D., and K.C. Border (2006): *Infinite Dimensional Analysis*: *A Hitchhiker's Guide*, Berlin, Heidelberg, New York: Springer.

[2] Attar, A., E. Campioni, and G. Piaser (2018): "On Competing Mechanisms under Exclusive Competition," *Games and Economic Behavior*, 111, 1–15.

[3] Attar, A., E. Campioni, and G. Piaser (2019): "Private Communication in Competing Mechanism Games," *Journal of Economic Theory*, 183, 258–283.

[4] Aumann, R.J. (1961): "Borel Structures for Function Spaces," *Illinois Journal of Mathematics*, 5(4), 614–630.

[5] Eeckhout, J., and P. Kircher (2010): "Sorting vs Screening: Search Frictions and Competing Mechanisms," *Journal of Economic Theory*, 145(4), 1354–1385.

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

[7] Fershtman, C., K.L. Judd, and E. Kalai (1991): "Observable Contracts: Strategic Delegation and Cooperation," *International Economic Review*, 32(3), 551–559.

[8] Gal-Or, E. (1997): "Multiprincipal Agency Relationships as Implied by Product Market Competition," *Journal of Economics and Management Strategy*, 6(1), 235–256.

[9] Ghosh, S., and S. Han (2018): "Competing Mechanisms: One-Shot versus Repeated Games," Unpublished Manuscript, McMaster University.

[10] Han, S. (2007): "Strongly Robust Equilibrium and Competing-Mechanism Games," *Journal of Economic Theory*, 137(1), 610–626.

[11] Han, S. (2016): "Sellers' Implicit Collusion in Directed Search Markets," *The B.E. Journal of Theoretical Economics*, 16(2), 711–738.

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

[13] Martimort, D. (1996): "Exclusive Dealing, Common Agency, and Multiprincipal Incentives Theory," *RAND Journal of Economics*, 27(1), 1–31.

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

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

[16] Moen, E.R. (1997): "Competitive Search Equilibrium," *Journal of Political Economy*, 105(2), 385–411.

[17] Myerson, R.B. (1979): "Incentive Compatibility and the Bargaining Problem," *Econometrica*, 47(1), 61–73.

[18] Myerson, R.B. (1982): "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems," *Journal of Mathematical Economics*, 10(1), 67–81.

[19] Peck, J. (1997): "A Note on Competing Mechanisms and the Revelation Principle," Unpublished Manuscript, Ohio State University.

[20] Peters, M. (1997): "A Competitive Distribution of Auctions," *Review of Economic Studies*, 64(1), 97–123.

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

[22] Peters, M. (2014): "Competing Mechanisms," *Canadian Journal of Economics*, 47(2), 373–397.

[23] Peters, M. (2015): "Reciprocal Contracting," *Journal of Economic Theory*, 158(PA), 102–126.

[24] Peters, M., and S. Severinov (1997): "Competition among Sellers Who Offer Auctions Instead of Prices," *Journal of Economic Theory*, 75(1), 141–179.

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

[26] Peters, M., and C. Troncoso-Valverde (2013): "A Folk Theorem for Competing Mechanisms," *Journal of Economic Theory*, 148(3), 953–973.

[27] Rao, B.V. (1971): "Borel Structures for Function Spaces," *Colloquium Mathematicum*, 23(1), 33–38.

[28] Sion, M. (1958): "On General Minimax Theorems," *Pacific Journal of Mathematics*, 8(1), 171–176.

[29] Szentes, B. (2009): "A Note on 'Mechanism Games with Multiple Principals and Three or More Agents' by T. Yamashita," Unpublished Manuscript, University College London.

[30] Szentes, B. (2015): "Contractible Contracts in Common Agency Problems," *Review of Economic Studies*, 82(1), 391–422.

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

[32] Virag, G. (2010): "Competing Auctions: Finite Markets and Convergence," *Theoretical Economics*, 5(2), 241–274.

[33] Wright, R., P. Kircher, B. Julien, and V. Guerrieri (2019): "Directed Search and Competitive Search: A Guided Tour," *Journal of Economic Literature*, forthcoming.

[34] Xiong, S. (2013): "A Folk Theorem for Contract Games with Multiple Principals and Agents," Unpublished Manuscript, Rice University.

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