

Keeping the Agents in the Dark: Private Disclosures in Competing Mechanisms*

Andrea Attar[†] Eloisa Campioni[‡]
Thomas Mariotti[§] Alessandro Pavan[¶]

October 31, 2021

Abstract

We study competing-mechanism games, in which several principals contract with several privately informed agents. We show that enabling principals to engage into private disclosures—whereby a principal sends to the agents contractible private signals about how her decision rule will respond to the agents’ messages—can significantly affect the predictions of such games. Our first result is that equilibrium outcomes and payoffs of games without private disclosures need no longer be supported once private disclosures are allowed for. This challenges the robustness of the folk theorem of Yamashita (2010). Our second result is that allowing for private disclosures may generate equilibrium outcomes that cannot be supported in any game without private disclosures, no matter how rich message spaces are. This challenges the canonicity of the universal mechanisms of Epstein and Peters (1999). These findings call for a novel approach to the analysis of competing-mechanism games.

Keywords: Incomplete Information, Competing Mechanisms, Private Disclosures, Folk Theorems, Universal Mechanisms.

JEL Classification: D82.

*We thank Laura Doval, Piotr Dworzak, Johannes Hörner, Michael Peters, and Takuro Yamashita for extremely valuable feedback. We also thank seminar audiences at Brown University, Higher School of Economics, New Economic School, Northwestern University, University of British Columbia, UCLA, and the Virtual Market Design Seminar, as well as conference participants at the 2020 Econometric Society World Congress and the 2021 ASSA Virtual Annual Meeting for many useful discussions. This research has benefited from financial support of the ANR (Programme d’Investissements d’Avenir ANR-17-EURE-0010), the MIUR (PRIN 2015), and the research foundation TSE-Partnership.

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

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

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

[¶]Northwestern University, Evanston, United States, and CEPR. Email: alepavan@northwestern.edu.

1 Introduction

Classical mechanism-design theory identifies the holding of private information by economic agents as a fundamental constraint on the allocations of resources (Hurwicz (1973)). How agents communicate their private information then becomes crucial for determining the set of allocations that can be implemented. In pure incomplete-information environments in which all payoff-relevant decisions are taken by a single uninformed principal, one can with no loss of generality restrict all private communication to be one-sided, from the agents to the principal (Myerson (1979)). Indeed, in that case, the principal need only post a mechanism selecting a (possibly random) decision for every profile of messages she may receive from the agents—what we hereafter refer to as a *standard mechanism*. Communication from the principal to the agents is then limited to the public announcement of such a mechanism; private communication from the principal to the agents is redundant, as it has no bearing on the set of allocations that the principal can implement.

In this paper, we argue that these basic insights from classical mechanism-design theory do not extend to competitive settings. To this end, we consider competing-mechanism games, in which the implementation of an allocation is no longer in the hands of a single principal, but of several principals who noncooperatively design mechanisms, each of which controls a specific dimension of the allocation, and in which all the agents participate. In this context, we show that allowing for private communication from the principals to the agents can significantly affect the set of allocations that can be supported in an equilibrium of such a game, and this, even if one focuses as we do on pure incomplete-information environments in which agents take no payoff-relevant decisions—arguably the least favorable scenario for this form of private communication to have bite. The general lesson from our results is that, in spite of their universal use in the theoretical and applied literature, the restriction to standard mechanisms is unwarranted in competitive settings, which calls for a drastic shift in the conceptualization and analysis of competing-mechanism games.

To put our findings in perspective, it is useful to keep in mind the two theoretical pillars of the competing-mechanism literature. First, following Epstein and Peters (1999), one can construct a space of (indirect) *universal mechanisms* whereby each agent can communicate to each principal his endogenous market information—that is, the profile of mechanisms posted by the other principals—in addition to his exogenous private information. An analogue of the revelation principle holds: any equilibrium outcome of any competing-mechanism game can be supported as an equilibrium outcome of the game in which the principals can only post universal mechanisms. This addresses the infinite-regress problem raised by the failure

of the standard revelation principle in competing-mechanism games (McAfee (1993), Peck (1997)). Second, following Yamashita (2010), one can obtain an explicit characterization of equilibrium outcomes and equilibrium payoffs for a large class of competing-mechanism games. An analogue of the folk theorem holds: any incentive-compatible allocation that yields each principal a payoff above a well-defined min-max-min bound can be supported in equilibrium by each principal posting a *recommendation mechanism*, whereby the agents, in addition to reporting their exogenous private information, are asked to vote on the direct mechanism the principal should use, which enables them to implement punishments following any deviation by some other principal.

From our perspective, the key point is that these two central results are established under the assumption that principals are restricted to post standard mechanisms, so that all private communication is from the agents to the principals. We challenge these results by considering a richer class of communication protocols in which the principals can privately disclose information to the agents before the agents have the opportunity to send messages to any of them. While so doing, we maintain two informational assumptions made in most of the literature. First, principals cannot directly communicate with each other. Second, principals' decisions and/or mechanisms cannot directly refer to each other; that is, principals cannot directly condition the decisions they take and/or the mechanisms they choose on the other principals' decisions and/or mechanisms.

We model private disclosures from the principals to the agents as contractible private signals sent by each principal to the agents about the decision rule she implements contingent on the messages she receives from them. The key informational assumption is that all these private signals are sent to the agents before any principal receives any messages from them. Each principal fully commits, as parts of the mechanism she posts, to a distribution of private signals and to a decision rule mapping the private signals she initially sends to the agents and the messages she ultimately receives from them into a (possibly random) decision. Notice that the extension of the game we thereby propose has no bite in the case of a single principal who can control all the dimensions of the allocation, as private disclosures do not affect the set of allocations that she can implement. In practice, such private disclosures may take the form of sellers privately informing buyers in an auction of personalized reservation prices before they submit their bids, manufacturers privately informing retailers of the amount of output supplied, or policy makers privately informing voters of the characteristics of a public good implemented as a result of the solicitation of the voters' preferences.

Our first result is that equilibria in standard mechanisms need not be robust to private

disclosures. To establish this result, we start from an example of a competing-mechanism game without public disclosures in which message spaces are sufficiently rich for the two principals in the game to post recommendation mechanisms. In this game, any feasible vector payoff for the principals can be supported in equilibrium.¹ However, we show that, by posting a mechanism that asymmetrically discloses information about her decisions to the agents, one of the principals can ensure that the agents no longer have the incentives to carry out the punishment in the other principal's mechanism that would be necessary to make the deviation unprofitable. Intuitively, this is because, by privately informing one of the agents of her decision while leaving the other agents in the dark, this principal is able to perfectly align her preferences in each state with those of that agent, making him no longer willing to participate to the required punishment. As a result, this principal can guarantee herself a payoff strictly above her minimum feasible payoff, regardless of the mechanism offered by the other principal and of the continuation equilibrium played by the agents. Many equilibrium payoffs that can be supported when principals compete in standard mechanisms can thus no longer be supported when, in addition to standard mechanisms, principals can also post mechanisms with private signals. The upshot of this example is that equilibrium outcomes of competing-mechanism games without private disclosures, and, in particular, those supported by recommendation mechanisms as in the folk theorem of Yamashita (2010), need not be robust to the possibility for principals to engage into private disclosures.

Our second result is that equilibrium outcomes of competing-mechanism games with private disclosures need not be equilibrium outcomes of any game in which principals are restricted to post standard mechanisms, no matter how rich message spaces are. The reason is that private disclosures may help principals correlate their decisions with the agents' exogenous private information in a way that cannot be replicated by principals responding to the agents' messages when these are solely based on the agents' common knowledge of the mechanisms and on their exogenous private information. To establish this result, we provide an example of a competing-mechanism game with public disclosures in which the equilibrium correlation between the principals' decisions and the agents' exogenous private information requires that (1) the agents receive information about one principal's decision and pass it on the other principal *before* the latter finalizes her own decision, and (2) such information not create common knowledge among the agents about the former principal's decision before they communicate with the latter. The example illustrates the possibility to achieve both (1) and (2) with private disclosures, and the necessity of both (1) and (2) when

¹This reflects that, in the example, the min-max-min bound characterized by Yamashita (2010) is equal to the minimum feasible payoff for each principal.

it comes to supporting certain outcomes, which implies the impossibility of supporting these outcomes with standard mechanisms, no matter how rich message spaces are. In equilibrium, the requirement (2) is satisfied by letting one principal send private signals that are, from any agent's perspective, completely uninformative of her decision, but that, once passed to the other principal in an incentive-compatible way, together perfectly reveal it; thus private disclosures take the form of encryption keys that leave all agents in the dark. The upshot of this example is that the universal mechanisms of Epstein and Peters (1999) are no longer canonical when principals can engage into private disclosures.

Taken together, the above results challenge the existing modeling approaches and suggest that private disclosures from the principals to the agents should be central to the theory of competing mechanisms.

Related Literature

This paper contributes to the theoretical foundations of competing-mechanism games. In a seminal paper, McAfee (1993) points out that the equilibria of such games may require that agents report *all* their private information to the principals; that is, their exogenous information about their private types and their market information about the mechanisms posted by the other principals. To overcome the resulting infinite-regress problem, Epstein and Peters (1999) construct a space of universal mechanisms that enables them to establish an analogue of the revelation principle for competing-mechanism games. Subsequent work provides explicit characterizations of the equilibrium outcomes of such games. Yamashita (2010) shows that, if there are at least three agents, every deterministic incentive-compatible allocation yielding each principal a payoff at least equal to a well-defined min-max-min bound can be supported in equilibrium. Peters and Troncoso-Valverde (2013), in a slightly different context, argue that any allocation that is incentive-compatible and individually rational in the sense of Myerson (1979) can be supported in equilibrium provided there are sufficiently many players. Our results indicate that such characterizations are sensitive to the assumption that principals are restricted to standard mechanisms and do not extend when principals can engage into private disclosures.

In classical mechanism-design theory (Myerson (1982)), private communication from a single principal to the agents is key when certain payoff-relevant actions can solely be taken by the agents, as in moral-hazard settings. Such communication, taking the form of action recommendations to the agents, has been shown to serve as a correlating device between the players' decisions in several problems of economic interest, such as Rahman and

Obara’s (2010) model of partnerships. Perhaps surprisingly, however, private disclosures have been neglected in competing-mechanism settings even when agents take payoff-relevant actions, as in Prat and Rustichini’s (2003) model of lobbying. To the best of our knowledge, the only exception is the recent work of Attar, Campioni, and Piaser (2019), who study complete-information games in which agents take observable actions. They construct an example in which equilibrium allocations supported by standard mechanisms fail to be robust against a deviation by a principal to a mechanism with private recommendations. In equilibrium, principals correlate their decisions with the agents’ actions in a way that cannot be achieved without private recommendations.

Whereas private signals from the principals to the agents can still be interpreted, in the above papers, in light of the traditional role they play in single-principal settings, the present paper uncovers two novel roles for such signals, making them very different from pure action recommendations. Private signals first alleviate the punishments that can be inflicted on a principal. In our first example, disclosing her decision to one agent while leaving the other agents in the dark enables one principal to guarantee herself a payoff strictly above her min-max-min bound by making it impossible for the agents to coordinate on the required punishment. Private signals also help overcoming the lack of a direct communication channel between the principals. In our second example, one principal can correlate her decisions with those of other principal and with the agents’ exogenous private information by sending encryption keys to the agents, to be passed to the other principal, in a way that cannot be replicated with standard mechanisms. Such a possibility appears particularly relevant for the applications of competing-mechanism games favored in the literature, where contracting is typically decentralized—so that principals may find it difficult to rely on a common mediator to coordinate their decisions—and direct communication among the principals—either in the form of principals exchanging cheap-talk messages, or in the form of principals using semi-private correlation devices whose realizations are observed by the other principals but not by the agents at the time they communicate with the principals—is unlikely to be feasible.

In our setting, the correlation in the principals’ decisions is generated by the agents communicating to each principal both their exogenous private information and the private signals they receive from other principals. The principals cannot directly condition on other principals’ decisions and/or mechanisms, or directly exchange information among themselves. Instead, Kalai, Kalai, Lehrer and Samet (2010), Peters and Szentes (2012), Peters (2015), and Szentes (2015) consider settings in which players can make commitments contingent

on other players' commitments and communication is unrestricted. The result in our first example extends to these settings: by deviating to a mechanism with private disclosures, a principal can guarantee herself a payoff strictly above her min-max-min bound, regardless of whether or not principals can make commitments contingent on other principals' decisions and/or mechanisms. The result in our second example, instead, extends to settings in which principals can condition their mechanism on other principals' mechanisms, as in Peters' (2015) model of reciprocal contracting, but not to settings in which principals can directly condition their decisions on other principals' decisions. This is because private disclosures in this example take the form of encryption keys passed from one principal to the other through the agents in an incentive-compatible way, something which can be achieved in a much more direct way if communication is unrestricted.

The role of private signals we highlight in this paper hinges on at least two principals contracting with at least two agents. When there is a single principal, the revelation principle obviously holds and private disclosures have no bearing on the set of equilibrium outcomes. Similarly, when multiple principals contract with a single agent, the menu theorems of Peters (2001), Martimort and Stole (2002), and Pavan and Calzolari (2009, 2010) guarantee that any equilibrium allocation of a game in which the principals compete by offering arbitrary message-contingent decision rules can be reproduced in a game in which the principals offer subsets (menus) of their decisions to the agent and delegate to him the choice of the final allocation. Thus private disclosures play no role in such settings either.

The paper is organized as follows. Section 2 introduces a general model of competing mechanisms under incomplete information. Sections 3 and 4 present the results. Section 5 discusses the different roles of private disclosures. Section 6 concludes.

2 The Model

We consider a pure incomplete-information setting in which several principals, indexed by $j = 1, \dots, J$, contract with several agents, indexed by $i = 1, \dots, I$. As anticipated in the introduction, our results hinge on $I \geq 2$ and $J \geq 2$.

Information Every agent i (he) possesses some exogenous private information summarized by his type ω^i , which belongs to some finite set Ω^i . Thus the set of exogenous states of the world $\omega \equiv (\omega^1, \dots, \omega^I)$ is $\Omega \equiv \Omega^1 \times \dots \times \Omega^I$. Principals and agents commonly believe that the state ω is drawn from Ω according to the distribution \mathbf{P} .

Decisions and Payoffs Every principal j (she) takes a decision x_j in some finite set X_j . We let $v_j : X \times \Omega \rightarrow \mathbb{R}$ and $u^i : X \times \Omega \rightarrow \mathbb{R}$ be the payoff functions of principal j and of agent i , respectively, where $X \equiv X_1 \times \dots \times X_J$ is the set of possible profiles of decisions for the principals. Agents take no payoff-relevant decisions. An *allocation* is a function $z : \Omega \rightarrow \Delta(X)$ assigning a lottery over the set X to every state of the world. The *outcome* induced by an allocation z is the restriction of z to the set of states occurring with positive probability under \mathbf{P} .²

Mechanisms with Signals A principal can engage into private disclosures by posting a *mechanism with signals*. Such a mechanism consists, first, of a probability distribution over the signals that the principal privately sends to the agents, and, second, of a decision rule that assigns a lottery over her decisions to every profile of signals she sends to the agents and every profile of messages she receives from them. Formally, a mechanism with signals for principal j is a pair $\gamma_j \equiv (\sigma_j, \phi_j)$ such that

1. $\sigma_j \in \Delta(S_j)$ is a Borel probability measure over the profiles of signals $s_j \equiv (s_j^1, \dots, s_j^I)$ that principal j sends to the agents, where $S_j \equiv S_j^1 \times \dots \times S_j^I$ for some collection of Polish spaces S_j^i of signals from principal j to every agent i .
2. $\phi_j : S_j \times M_j \rightarrow \Delta(X_j)$ is a Borel-measurable function assigning a lottery over principal j 's decisions to every profile of signals $s_j \in S_j$ she sends to the agents and every profile of messages $m_j \equiv (m_j^1, \dots, m_j^I) \in M_j$ she receives from them, where $M_j \equiv M_j^1 \times \dots \times M_j^I$ for some collection of Polish spaces M_j^i of messages from every agent i to principal j .

We assume that $\Omega^i \subset M_j^i$ for all i and j , so that the language through which agent i communicates with principal j is rich enough for him to report his type to her. A (potentially indirect) *standard mechanism* for principal j is a special case of a mechanism with signals in which S_j^i is a singleton for all i ; hereafter, we will often simplify the notation by omitting σ_j and representing a standard mechanism solely by a Borel-measurable function $\phi_j : M_j \rightarrow \Delta(X_j)$ describing the principal's response to the messages she receives from the agents. The requirement that signal and message spaces be Polish entails no loss of generality; in particular, the universal standard mechanisms of Epstein and Peters (1999) involve uncountable Polish message spaces.

²The distinction between an allocation and an outcome is relevant when the agents' types are correlated.

Admissibility A general requirement for defining expected payoffs in the game to be described below is that, for each j , the evaluation mapping $(\phi_j, s_j, m_j) \mapsto \phi_j(s_j, m_j)$ be measurable. To do so, we must define a measurable structure over the space of admissible functions ϕ_j . If S_j and M_j are countable, we can take this space to be $\Delta(X_j)^{S_j \times M_j}$, endowed with the product Borel σ -field. If S_j or M_j are uncountable, however, there is no measurable structure over the space of all Borel-measurable functions $\phi_j : S_j \times M_j \rightarrow \Delta(X_j)$ such that the evaluation mapping for principal j is measurable (Aumann (1961)); in that case, there is no choice but to restrict the space of admissible functions ϕ_j . Admissibility can be shown to coincide with the requirement that this space be of bounded Borel class (Aumann (1961), Rao (1971)), which still allows for a rich class of mechanisms for our analysis. We hereafter fix an admissible space Φ_j , endowed with a σ -field \mathcal{F}_j , so that $\Gamma_j \equiv \Delta(S_j) \times \Phi_j$ is the space of admissible mechanisms for principal j , endowed with the product σ -field \mathcal{G}_j generated by the Borel subsets of $\Delta(S_j)$ and the elements of \mathcal{F}_j . When attention is restricted to standard mechanisms, the set of admissible mechanisms is simply denoted by Φ_j , with the understanding that signal spaces are singletons.

Timing and Strategies The competing-mechanism game G^{SM} with private disclosures unfolds in four stages:

1. the principals simultaneously post mechanisms, observed by all agents;
2. the principals' mechanisms send private signals to the agents;
3. the agents simultaneously send messages to the principals;
4. the principals' decisions are implemented and all payoffs accrue.

A mixed strategy for principal j is a probability measure $\mu_j \in \Delta(\Gamma_j)$ over \mathcal{G}_j . A strategy for agent i is a measurable function $\lambda^i : \Gamma \times S^i \times \Omega^i \rightarrow \Delta(M^i)$ that assigns to every profile of mechanisms $\gamma \equiv (\gamma_1, \dots, \gamma_J) \in \Gamma \equiv \Gamma_1 \times \dots \times \Gamma_J$ that the principals may post, every profile of signals $s^i \equiv (s_1^i, \dots, s_J^i) \in S^i \equiv S_1^i \times \dots \times S_J^i$ that agent i may receive, and every type $\omega^i \in \Omega^i$ of agent i a Borel probability measure over the profiles of messages $m^i \equiv (m_1^i, \dots, m_J^i) \in M^i \equiv M_1^i \times \dots \times M_J^i$ sent by agent i , where $\Gamma \times S^i \times M^i$ is endowed with the appropriate product σ -field. The allocation $z_{\mu, \lambda} : \Omega \rightarrow \Delta(X)$ induced by the strategies $(\mu, \lambda) \equiv (\mu_1, \dots, \mu_J, \lambda^1, \dots, \lambda^I)$ is then defined by

$$z_{\mu, \lambda}(x|\omega) \equiv \int_{\Gamma} \int_S \int_M \prod_{j=1}^J \phi_j(s_j, m_j)(x_j) \bigotimes_{i=1}^I \lambda^i(dm^i | \gamma, s^i, \omega^i) \bigotimes_{j=1}^J \sigma_j(ds_j) \bigotimes_{j=1}^J \mu_j(d\gamma_j)$$

for all $(\omega, x) \in \Omega \times X$, where $S \equiv S_1 \times \dots \times S_J$ and $M \equiv M_1 \times \dots \times M_J$. For every profile of mechanisms γ , a behavior strategy for agent i in the subgame γ played by the agents is a Borel-measurable function $\beta^i : S^i \times \Omega^i \rightarrow \Delta(M^i)$ assigning a Borel probability measure over the profile of messages $m^i \in M^i$ she sends to the principals to every profile of signals $s^i \in S^i$ she may receive and to every realization $\omega^i \in \Omega^i$ of her type. We let $z_{\gamma, \beta}$ be the allocation induced by the profile of behavior strategies $\beta \equiv (\beta^1, \dots, \beta^I)$ in the subgame γ ; the latter is defined in the same way as $z_{\mu, \lambda}$, except that γ is fixed and $\lambda^i(\cdot | \gamma, s^i, \omega^i)$ is replaced by $\beta^i(\cdot | s^i, \omega^i)$ for all i .

A degenerate case of the game G^{SM} arises when S_j^i is a singleton for all i and j , so that principals cannot engage into private disclosures and can only post standard mechanisms. To distinguish this situation, we denote by G^M the corresponding competing-mechanism game without private disclosures; the games studied by Epstein and Peters (1999) and Yamashita (2010) are prominent examples.

Equilibrium In line with the standard practice of the common-agency literature (Peters (2001), Martimort and Stole (2002)) and the competing-mechanism literature (Epstein and Peters (1999), Yamashita (2010), Peters and Troncoso-Valverde (2013), Peters (2014)), we will assume that the agents treat the mechanisms posted by the principals at face value. This means that we can identify any subgame $\gamma \equiv (\gamma_1, \dots, \gamma_J) \in \Gamma$ of G^{SM} with a Bayesian game played by the agents, with type space $S^i \times \Omega^i$ and action space M^i for every agent i , and in which the agents' beliefs are pinned down by the prior \mathbf{P} and the signal distributions $(\sigma_1, \dots, \sigma_J)$ to which the principals are committed through the mechanisms they post in γ , whether or not γ is reached on the equilibrium path. The strategy profile (μ, λ) is a perfect Bayesian equilibrium (PBE) of G^{SM} whenever

1. for each $\gamma \in \Gamma$, $(\lambda^1(\gamma), \dots, \lambda^I(\gamma))$ is a Bayes–Nash equilibrium (BNE) of the subgame γ played by the agents;
2. given the continuation equilibrium strategies λ , μ is a Nash equilibrium of the game played by the principals.

An allocation z is *incentive-compatible* if, for all i and $\omega^i \in \Omega^i$,

$$\omega^i \in \arg \max_{\hat{\omega}^i \in \Omega^i} \sum_{\omega^{-i} \in \Omega^{-i}} \sum_{x \in X} \mathbf{P}[\omega^{-i} | \omega^i] z(x | \hat{\omega}^i, \omega^{-i}) u^i(x, \omega^i, \omega^{-i}).$$

It follows from the definition of a BNE in any subgame played by the agents that any allocation $z_{\mu, \lambda}$ supported by a PBE (μ, λ) of G^{SM} is incentive-compatible; otherwise, some

type ω^i of some agent i would be strictly better off mimicking the strategy $\lambda^i(\cdot \mid \cdot, \cdot, \hat{\omega}^i)$ of some other type $\hat{\omega}^i$ —this is an instance of the revelation principle (Myerson (1982)). This observation implies that, when there is a single principal, any allocation that can be implemented by a mechanism with signals can also be implemented via a direct revelation mechanism; as agents take no payoff-relevant actions, such direct revelation mechanisms involve no private disclosures from the principal to the agents. As we show below, the situation is markedly different when several principals contract with several agents.

3 A Challenge to Folk Theorems

In this section, we address the question of whether equilibrium outcomes of competing-mechanism games without private disclosures, in which principals are restricted to posting standard mechanisms, yet with potentially rich message spaces, are robust to the possibility for principals to post mechanisms with signals. This question is especially relevant in light of the fact that, as shown by Yamashita (2010) and Peters and Troncoso-Valverde (2013), such games typically lend themselves to folk-theorem-types of results. Notice for future reference that similar results are also pervasive in the literature on contractible contracts and reciprocal contracting; see, for instance, Kalai, Kalai, Lehrer, and Samet (2010), Peters and Szentes (2012), Peters (2015), and Szentes (2015).

The construction of Yamashita (2010), which we exploit in Section 3.1 below, is based on the idea that, given rich enough message spaces, each principal’s equilibrium mechanism can be made sufficiently flexible to punish other principals’ potential deviations. This can be achieved by enabling the agents to recommend to every principal j a (deterministic) direct mechanism $d_j : \Omega \rightarrow X_j$ selecting a decision for any profile of reported types she may receive from them. Specifically, let us consider a competing-mechanism game without private disclosures in which every message space M_j^i is sufficiently rich to enable agent i to recommend any direct mechanism to principal j and to make a report about his type; that is, letting D_j be the finite set of all such mechanisms, $D_j \times \Omega^i \subset M_j^i$ for all i and j . Accordingly, a *recommendation mechanism* ϕ_j^r for principal j stipulates that, if every agent i sends a message $m_j^i \equiv (d_j^i, \omega^i) \in D_j \times \Omega^i$ to principal j , then

$$\phi_j^r(m_j^1, \dots, m_j^I) \equiv \begin{cases} d_j(\omega^1, \dots, \omega^I) & \text{if } \text{card} \{i : d_j^i = d_j\} \geq I - 1 \\ \bar{x}_j & \text{otherwise} \end{cases}, \quad (1)$$

where \bar{x}_j is some fixed decision in X_j ; if, instead, some agent i sends a message $m_j^i \notin D_j \times \Omega^i$ to principal j , then ϕ_j^r treats this message as if it coincided with some fixed element $(\bar{d}_j, \bar{\omega}_j^i)$ of $D_j \times \Omega^i$, once again applying rule (1). Intuitively, recommendation mechanisms provide

a flexible system of punishments against other principals' potential deviations that can be used to support many equilibrium allocations. Indeed, Yamashita (2010) establishes the following folk theorem: if $I \geq 3$, then every deterministic incentive-compatible allocation yielding each principal a payoff at least equal to a well-defined min-max-min payoff bound can be supported in equilibrium.³

We now provide an example showing that the possibility for principals to use private disclosures undermines this characterization result. In this example, a folk theorem holds for any competing-mechanism game without private disclosures but with rich enough message spaces; however, a continuum of equilibrium payoff vectors of any such game can no longer be supported when principals can post mechanisms with signals. This shows that equilibrium outcomes supported by standard mechanisms may not be robust to private disclosures.

Example 1 Let $J \equiv 2$ and $I \equiv 3$. We denote the principals by P1 and P2, and the agents by A1, A2, and A3. The decision sets are $X_1 \equiv \{x_{11}, x_{12}\}$ for P1 and $X_2 \equiv \{x_{21}, x_{22}\}$ for P2. A1 and A2 can each be of two types, with $\Omega^1 = \Omega^2 \equiv \{\omega_L, \omega_H\}$, whereas A3 can only be of a single type, which we omit from the notation for the sake of clarity. A1's and A2's types are perfectly correlated: only the states (ω_L, ω_L) and (ω_H, ω_H) can occur with positive probability under \mathbf{P} .

The players' payoffs are represented in Tables 1 and 2 below, in which the first payoff is that of P2 and the last two payoffs are those of A1 and A2, respectively. P1's and A3's payoffs are constant over $X \times \Omega$ and hence play no role in the analysis.

	x_{21}	x_{22}
x_{11}	5, 8, 8	5, 1, 1
x_{12}	6, 4.5, 4.5	6, 4.5, 4.5

Table 1: Payoffs in state (ω_L, ω_L) .

	x_{21}	x_{22}
x_{11}	6, 4.5, 4.5	6, 4.5, 4.5
x_{12}	5, 1, 1	5, 8, 8

Table 2: Payoffs in state (ω_H, ω_H) .

³As pointed out by Peters (2014), however, these bounds typically depend on the message spaces M_j^i . The requirement that there be at least three agents reflects that, according to (1), near unanimity unequivocally pins down a direct mechanism for each principal posting a recommendation mechanism. Relatedly, Attar, Campioni, Mariotti, and Piaser (2021) show that this and related folk theorems crucially hinge on each agent participating and communicating with each principal, regardless of the profile of posted mechanisms.

3.1 A Folk Theorem in Standard Mechanisms

In the context of this example, let us first consider, as in Yamashita (2010), a general competing-mechanism game G_1^M without private disclosures, and with message spaces such that $D_j \times \Omega^i \subset M_j^i$ for all i and j , so that principals can post recommendation mechanisms. To guarantee the existence of a BNE in every subgame $\phi \equiv (\phi_1, \phi_2)$ of G_1^M , we assume that all the message spaces M_j^i are finite. Our first result characterizes the set of equilibrium payoffs for P2 in G_1^M .

Lemma 1 *Any payoff for P2 in $[5, 6]$ can be supported in a PBE of G_1^M .*

Proof. The proof consists of two steps. Step 1 proves that G_1^M admits a PBE in which P2 obtains his minimum feasible payoff of 5, which also shows that 5 is P2's min-max-min payoff in G_1^M . Step 2 then leverages on the construction in Step 1 to prove that any payoff for P2 in $(5, 6]$ can be supported in a PBE of G_1^M , which completes the proof.

Step 1 We first show that the outcome

$$z(\omega_L, \omega_L) \equiv \delta_{(x_{11}, x_{21})}, \quad z(\omega_H, \omega_H) \equiv \delta_{(x_{12}, x_{22})}, \quad (2)$$

in which P2 obtains her minimum feasible payoff of 5, can be supported in a PBE of G_1^M .⁴ To establish this result, we first show that, if P1 and P2 post recommendation mechanisms, then there exists a continuation BNE supporting the outcome (2). We next show that, in every subgame in which P1 posts her equilibrium recommendation mechanism, there exists a continuation BNE in which P2 obtains a payoff of 5. The result then follows from these two properties along with the fact that P1 has no profitable deviation as her payoff is constant over $X \times \Omega$.

On Path Suppose that both P1 and P2 post recommendation mechanisms ϕ_1^r and ϕ_2^r . We assume that, for each j , $\bar{\omega}_j^1 = \bar{\omega}_j^2 = \omega_L$, so that, if some agent $i = 1, 2$ sends a message $m_j^i \notin D_j \times \Omega^i$ to principal j , ϕ_j^r treats this message as if agent i reported to principal j to be of type ω_L . We claim that, in the subgame (ϕ_1^r, ϕ_2^r) , it is a BNE for the three agents to recommend the direct mechanisms (d_1^*, d_2^*) defined by

$$d_1^*(\omega) \equiv \begin{cases} x_{11} & \text{if } \omega = (\omega_L, \omega_L) \\ x_{12} & \text{otherwise} \end{cases} \quad \text{and} \quad d_2^*(\omega) \equiv \begin{cases} x_{21} & \text{if } \omega = (\omega_L, \omega_L) \\ x_{22} & \text{otherwise} \end{cases} \quad (3)$$

for all $\omega \equiv (\omega^1, \omega^2) \in \Omega^1 \times \Omega^2$, and for A1 and A2 to report their types truthfully to P1 and P2. To see this, we only need to observe that these strategies implement the outcome

⁴For any finite set A and for each $a \in A$, δ_a is the Dirac measure over A assigning probability 1 to a .

(2), which yields A1 and A2 their maximum feasible payoff of 8 in each state; because A3's payoff is constant over $X \times \Omega$, these strategies thus form a BNE of the subgame (ϕ_1^r, ϕ_2^r) . The claim follows.

Off Path Because P1's payoff is constant over $X \times \Omega$, she has no profitable deviation. Suppose then that P2 deviates to some arbitrary standard mechanism $\phi_2 : M_2 \rightarrow \Delta(X_2)$, and let $p(m_2)$ be the probability that the lottery $\phi_2(m_2)$ assigns to decision x_{21} when the agents send the messages $m_2 \equiv (m_2^1, m_2^2, m_2^3) \in M_2$ to P2. Now, let

$$\bar{p} \equiv \max_{m_2 \in M_2} p(m_2) \quad (4)$$

and select a message profile $\bar{m}_2 \equiv (\bar{m}_2^1, \bar{m}_2^2, \bar{m}_2^3) \in M_2$ that achieves the maximum in (4); similarly, let

$$\underline{p} \equiv \min_{(m_2^1, m_2^2) \in M_2^1 \times M_2^2} p(m_2^1, m_2^2, \bar{m}_2^3) \quad (5)$$

and select a message profile $(\underline{m}_2^1, \underline{m}_2^2) \in M_2^1 \times M_2^2$ for A1 and A2 that, given \bar{m}_2^3 , achieves the minimum in (5). That \bar{p} , \bar{m}_2 , \underline{p} , and $(\underline{m}_2^1, \underline{m}_2^2)$ are well-defined for any given ϕ_2 follows from the fact that M_2 is finite. We now prove that there exist BNE strategies for the agents in the subgame (ϕ_1^r, ϕ_2) such that P2 obtains a payoff of 5, so that the deviation is not profitable. We consider two cases in turn.

Case 1: $\bar{p} \geq \frac{1}{2}$ Suppose first that ϕ_2 is such that $\bar{p} \geq \frac{1}{2}$. We claim that the subgame (ϕ_1^r, ϕ_2) admits a BNE that satisfies the following properties: (i) all agents recommend the direct mechanism d_1^* to P1, as if P2 did not deviate; (ii) A1 and A2 truthfully report their types to P1; (iii) A3 sends message \bar{m}_2^3 to P2; (iv) P2 obtains a payoff of 5. As for (i), the argument is that unilaterally sending a different recommendation to P1 is of no avail as no agent is pivotal. As for (iii), sending \bar{m}_2^3 to P2 is optimal for A3 given that his payoff is constant over $X \times \Omega$. Consider then (ii). Suppose first that the state is (ω_L, ω_L) . Because $\bar{p} \geq \frac{1}{2}$, $8\bar{p} + (1 - \bar{p}) \geq 4.5$. From Table 1, and by definition of d_1^* and \bar{m}_2 , it thus follows that, if A2 reports ω_L to P1 and sends \bar{m}_2^2 to P2, and if A3 sends \bar{m}_2^3 to P2, then A1 best responds by reporting ω_L to P1 and sending \bar{m}_2^1 to P2; notice, in particular, that, because $\bar{\omega}_1^1 = \omega_L$, if A1 sends a message $m_1^1 \notin D_1 \times \Omega^1$ to P1, then P1 takes the same decision as if A1 truthfully reported his type to her. The argument for A2 is identical. Suppose next that the state is (ω_H, ω_H) . If either A1 or A2 truthfully reports his type to P1, then, by definition of d_1^* , the other informed agent A2 or A1 cannot induce P1 to take a decision other than x_{12} . These properties, along with the finiteness of M_2 , imply that the subgame (ϕ_1^r, ϕ_2) admits a BNE

satisfying (i)–(iii). In this BNE, P1 takes decision x_{11} in state (ω_L, ω_L) and decision x_{12} in state (ω_H, ω_H) , yielding a payoff of 5 to P2, as required by (iv). The claim follows.

Case 2: $\bar{p} < \frac{1}{2}$ Suppose next that ϕ_2 is such that $\bar{p} < \frac{1}{2}$. We claim that the subgame (ϕ_1^r, ϕ_2) admits a BNE that satisfies the following properties: (i) all agents recommend the direct mechanism

$$d_1(\omega) \equiv \begin{cases} x_{12} & \text{if } \omega = (\omega_H, \omega_H) \\ x_{11} & \text{otherwise} \end{cases} \quad (6)$$

to P1; (ii) A1 and A2 truthfully report their types to P1; (iii) A3 sends message \bar{m}_2^3 to P2; (iv) P2 obtains a payoff of 5. The arguments for (i) and (iii) are the same as in Case 1. Consider then (ii). Suppose first that the state is (ω_L, ω_L) . If either A1 or A2 truthfully reports his type to P1, then, by definition of d_1 , the other informed agent A2 or A1 cannot induce P1 to take a decision other than x_{11} . Suppose next that the state is (ω_H, ω_H) . Because $\underline{p} \leq \bar{p} < \frac{1}{2}$, $\underline{p} + 8(1 - \underline{p}) > 4.5$. From Table 2, and by definition of d_1 and $(\underline{m}_2^1, \underline{m}_2^2)$, it thus follows that, if A2 reports ω_H to P1 and sends \underline{m}_2^2 to P2, and if A3 sends \bar{m}_2^3 to P2, then A1 best responds by reporting ω_H to P1 and sending \underline{m}_2^1 to P2; notice, in particular, that, because $\bar{\omega}_1^1 = \omega_L$, if A1 sends a message $m_1^1 \notin D_1 \times \Omega^1$ to P1, then P1 takes the same decision as if A1 misreported his type. The argument for A2 is identical. These properties, along with the finiteness of M_2 , imply that the subgame (ϕ_1^r, ϕ_2) admits a BNE satisfying (i)–(iii). The argument for (iv) is then the same as in Case 1. The claim follows.

This completes Step 1 of the proof. We refer to the Appendix for Step 2 of the proof. The result follows. ■

The arguments in Step 1 of the proof rely on the same intuition as in Yamashita (2010, Theorem 1). The possibility for the agents to recommend a different direct mechanism to P1 for every mechanism posted by P2 allows them to implement punishments contingent on P2's deviations. In particular, the argument in Case 2 shows that any deviation by P2 to a mechanism that implements x_{21} with a probability strictly less than $\frac{1}{2}$ is blocked by recommending to P1 the direct mechanism d_1 , which is different from the equilibrium mechanism d_1^* . Observe that, although we allow principals to post stochastic mechanisms, unlike in Yamashita (2010), the threat of agents choosing a deterministic direct mechanism is sufficient to yield P2 her minimum feasible payoff of 5 in equilibrium.

Stochastic mechanisms, however, can be used to support random allocations; see, for instance, Xiong (2013). Step 2 of the proof in turn shows that this permits one to support any payoff for P2 in the feasible set $[5, 6]$ in equilibrium. The proof, which is provided in

the Appendix, only requires adjusting the principals' behavior on path—off path, letting the agents coordinate on the mechanisms d_1^* or d_1 used in Step 1 suffices to deter P2's deviations. To this end, we modify Yamashita's (2010) definition of a recommendation mechanism to allow principals to randomize over their decisions on path, while maintaining the assumption that the agents' message spaces are finite.

In related work, Peters and Troncoso-Valverde (2013) establish a folk theorem in a generalized version of Yamashita (2010). In the game they study, any outcome corresponding to an allocation that is incentive-compatible and individually rational in the sense of Myerson (1979) can be supported in equilibrium provided there are at least seven players. It is straightforward to check that the outcome (2) satisfies these conditions, which guarantees that it can also be supported in equilibrium in their framework.⁵ Notice finally that, whereas, in general, a principal's min-max-min payoff may be sensitive to the richness of the available message spaces, in our example P1 posting a recommendation mechanism is sufficient to inflict P2 her minimum feasible payoff of 5, leaving no role for additional messages beyond those contained in $D_1 \times \Omega^i$ for all i . In other words, that P2's relevant min-max-min payoff is equal to 5 is fairly uncontroversial.

3.2 Nonrobustness to Private Disclosures

We now show that many of the equilibrium outcomes characterized in Lemma 1 are not robust to private disclosures. Specifically, we show that, in any enlarged game in which principals can post mechanisms with signals, P2 can guarantee herself a payoff higher than her min-max-min payoff of 5. To this end, we consider a general competing-mechanism game G_1^{SM} with private disclosures; this notably includes the case where $D_j \times \Omega^i \subset M_j^i$ for all i and j , as in the game G_1^M studied in Section 3.1. To guarantee that the result is not driven by the possible nonexistence of equilibria, we assume that all the signal spaces S_j^i and the message spaces M_j^i are finite.⁶ The following result then holds.

Lemma 2 G_1^{SM} admits a PBE. Moreover, if $\text{card } S_2^1 \geq 2$, then P2's payoff in any PBE of G_1^{SM} is at least equal to $5 + \frac{\mathbf{P}[(\omega_L, \omega_L)]\mathbf{P}[(\omega_H, \omega_H)]}{2 - \mathbf{P}[(\omega_L, \omega_L)]}$.

Proof. We first show that a PBE exists. We next establish the desired bound on P2's equilibrium payoff.

⁵The requirement on the number of players can be met by adding additional agents identical to A3.

⁶As the arguments below reveal, the second part of Lemma 2 does not hinge on this simplifying assumption, and extends to any infinite game G_1^{SM} for which a PBE exists.

Existence of a PBE Because, for each j , the sets S_j and M_j are finite, the space $\Gamma_j \equiv \Delta(S_j) \times \Delta(X_j)^{S_j \times M_j}$ of mechanisms for principal j in G_1^{SM} is compact, and every subgame $(\gamma_1, \gamma_2) \in \Gamma_1 \times \Gamma_2$ is finite; moreover, the agents' information structures and payoffs are continuous functions of (γ_1, γ_2) . Hence the BNE of the subgame (γ_1, γ_2) form a nonempty compact set $B^*(\gamma_1, \gamma_2)$, and the correspondence $B^* : \Gamma_1 \times \Gamma_2 \rightarrow \prod_{i=1}^3 \Delta(M^i)^{S^i \times \Omega^i}$ is upper hemicontinuous (Milgrom and Weber (1985, Theorem 2)) and, therefore, admits a Borel-measurable selection $b^* \equiv (b^{1*}, b^{2*}, b^{3*})$ by Kuratowski and Ryll-Nardzewski's selection theorem (Aliprantis and Border (2006, Theorem 18.13)); for each i , the corresponding strategy for agent i in G_1^{SM} is defined by $\lambda^{i*}(m^i | \gamma_1, \gamma_2, s^i, \omega^i) \equiv b^{i*}(\gamma_1, \gamma_2)(m^i | s^i, \omega^i)$. Now, suppose that P1 posts the mechanism γ_1^* that equiprobably randomizes between decisions x_{11} and x_{12} regardless of the signals P1 sends to the agents and the messages she receives from them. Then, from Tables 1–2, P2 obtains an expected payoff of 5.5 regardless of the mechanism she posts. Because P1's payoff is constant over $X \times \Omega$, it follows that, for each $\gamma_2^* \in \Gamma_2$, $(\gamma_1^*, \gamma_2^*, \lambda^{1*}, \lambda^{2*})$ is a PBE of G_1^{SM} .

A Tighter Payoff Bound for P2 For each $\sigma \in (\frac{1}{2}, 1)$, we first construct a mechanism $\gamma_2(\sigma) \in \Gamma_2$ that guarantees P2 a payoff of $5 + \frac{(1-\sigma)\mathbf{P}[(\omega_L, \omega_L)]\mathbf{P}[(\omega_H, \omega_H)]}{1-\sigma\mathbf{P}[(\omega_L, \omega_L)]}$ regardless of the mechanism posted by P1 and of the agents' continuation equilibrium strategies; that is,

$$\begin{aligned} \inf_{\gamma_1 \in \Gamma_1} \inf_{\beta \in B^*(\gamma_1, \gamma_2(\sigma))} \sum_{\omega \in \Omega} \sum_{x \in X} \mathbf{P}[\omega] z_{\gamma_1, \gamma_2(\sigma), \beta}(x | \omega) v_2(x, \omega) \\ \geq 5 + \frac{(1-\sigma)\mathbf{P}[(\omega_L, \omega_L)]\mathbf{P}[(\omega_H, \omega_H)]}{1-\sigma\mathbf{P}[(\omega_L, \omega_L)]}, \end{aligned} \quad (7)$$

where $z_{\gamma_1, \gamma_2(\sigma), \beta}(x | \omega)$ is the probability that the decision profile x is implemented when the agents' private information is ω , the principals' mechanisms are $(\gamma_1, \gamma_2(\sigma))$, and the agents play according to β . To see this, suppose without loss of generality that $\{1, 2\} \subset S_2^1$ and $\emptyset \in S_2^i$ for $i = 2, 3$. Fix then some $\sigma \in (\frac{1}{2}, 1)$, and let $\gamma_2(\sigma)$ be the mechanism with signals for P2 such that

- with probability $\sigma_2(1, \emptyset, \emptyset) \equiv \sigma$, P2 sends signal $s_2^1 = 1$ to A1 and signals $s_2^2 = s_2^3 = \emptyset$ to A2 and A3 and takes decision x_{21} regardless of the profile of messages she receives from the agents;
- with probability $\sigma_2(2, \emptyset, \emptyset) \equiv 1 - \sigma$, P2 sends signal $s_2^1 = 2$ to A1 and signals $s_2^2 = s_2^3 = \emptyset$ to A2 and A3 and takes decision x_{22} regardless of the profile of messages she receives from the agents.

Thus, given the private signals sent by P2, A1 knows exactly P2's decision, while A2 and A3 remain uninformed. That is, A2 and A3 believe that P2 takes decision x_{21} with probability

σ and decision x_{22} with probability $1 - \sigma$; yet they know that A1 knows P2's decision. We claim that $\gamma_2(\sigma)$ satisfies (7).

Indeed, suppose, by way of contradiction, that there exists $(\gamma_1, \beta) \in \Gamma_1 \times B^*(\gamma_1, \gamma_2(\sigma))$ such that, given (γ_1, β) , posting $\gamma_2(\sigma)$ yields P2 a payoff of $5 + \varepsilon$, where

$$0 \leq \varepsilon < \frac{(1 - \sigma)\mathbf{P}[(\omega_L, \omega_L)]\mathbf{P}[(\omega_H, \omega_H)]}{1 - \sigma\mathbf{P}[(\omega_L, \omega_L)]}. \quad (8)$$

Observe that the mechanism $\gamma_2(\sigma)$ implements decisions in X_2 that are independent of any messages P2 may receive from the agents and, hence, of any signals sent by γ_1 . Thus the only role that signals in γ_1 could play, given $\gamma_2(\sigma)$, would be to affect the distribution over P1's decisions induced by the agents; but it follows from standard arguments (Myerson (1982)) that messages are enough to this end, and thus that signals are redundant. We can thus assume that γ_1 is a standard mechanism ϕ_1 , involving no signals.

We first establish some useful accounting inequalities. Given $(\phi_1, \gamma_2(\sigma))$ and β , the probability that P1 takes decision x_{11} in state (ω_L, ω_L) can be written as

$$\pi_{11}(\omega_L, \omega_L) \equiv \sigma\pi_{11}(\omega_L, \omega_L, 1) + (1 - \sigma)\pi_{11}(\omega_L, \omega_L, 2) \quad (9)$$

where, for each $s_2^1 \in \{1, 2\}$,

$$\pi_{11}(\omega_L, \omega_L, s_2^1) \equiv \sum_{(m_1^1, m_1^2, m_1^3) \in M_1} \beta^1(m_1^1 | s_2^1, \omega_L) \beta^2(m_1^2 | \omega_L) \beta^3(m_1^3) \phi_1(x_{11} | m_1^1, m_1^2, m_1^3) \quad (10)$$

is the probability that P1 takes decision x_{11} in state (ω_L, ω_L) , conditional on P2 sending signal s_2^1 to A1. Similarly, the probability that P1 takes decision x_{12} in state (ω_H, ω_H) can be written as

$$\pi_{12}(\omega_H, \omega_H) \equiv \sigma\pi_{12}(\omega_H, \omega_H, 1) + (1 - \sigma)\pi_{12}(\omega_H, \omega_H, 2) \quad (11)$$

where, for each $s_2^1 \in \{1, 2\}$,

$$\pi_{12}(\omega_H, \omega_H, s_2^1) \equiv \sum_{(m_1^1, m_1^2, m_1^3) \in M_1} \beta^1(m_1^1 | s_2^1, \omega_H) \beta^2(m_1^2 | \omega_H) \beta^3(m_1^3) \phi_1(x_{12} | m_1^1, m_1^2, m_1^3) \quad (12)$$

is the probability that P1 takes decision x_{12} in state (ω_H, ω_H) , conditional on P2 sending signal s_2^1 to A1. By definition of ε , we have

$$\mathbf{P}[(\omega_L, \omega_L)][6 - \pi_{11}(\omega_L, \omega_L)] + \mathbf{P}[(\omega_H, \omega_H)][6 - \pi_{12}(\omega_H, \omega_H)] = 5 + \varepsilon,$$

or, equivalently,

$$\mathbf{P}[(\omega_L, \omega_L)]\pi_{11}(\omega_L, \omega_L) + \mathbf{P}[(\omega_H, \omega_H)]\pi_{12}(\omega_H, \omega_H) = 1 - \varepsilon,$$

which implies

$$\pi_{11}(\omega_L, \omega_L) \geq 1 - \frac{\varepsilon}{\mathbf{P}[(\omega_L, \omega_L)]} \quad \text{and} \quad \pi_{12}(\omega_H, \omega_H) \geq 1 - \frac{\varepsilon}{\mathbf{P}[(\omega_H, \omega_H)]} \quad (13)$$

as both $\pi_{11}(\omega_L, \omega_L)$ and $\pi_{12}(\omega_H, \omega_H)$ are at most equal to 1. Notice that (8) ensures that the right-hand side of each inequality in (14) is strictly positive, and thus can be interpreted as a probability as it is at most equal to 1. Similarly, it follows from (9) and from the first inequality in (13) that

$$\pi_{11}(\omega_L, \omega_L, 2) \geq 1 - \frac{\varepsilon}{(1 - \sigma)\mathbf{P}[(\omega_L, \omega_L)]}. \quad (14)$$

Again, (8) ensures that the right-hand side of (14) is strictly positive, and thus can be interpreted as a probability as it is at most equal to 1.

We now come to the bulk of the argument. From Table 1, in state (ω_L, ω_L) , and upon receiving signal $s_2^1 = 2$ from P2, A1 wants to minimize the probability that P1 takes decision x_{11} . It follows that, given the reporting strategies $\beta^2(\cdot | \omega_L)$ and β^3 of A2 and A3, any message that A1 sends with positive probability to P1 in state (ω_L, ω_L) upon receiving signal $s_2^1 = 2$ from P2 induces P1 to take decision x_{11} with probability $\pi_{11}(\omega_L, \omega_L, 2)$, and, by (10) and (14), that, for any message $m_1^1 \in M_1^1$,

$$\sum_{(m_1^2, m_1^3) \in M_1^2 \times M_1^3} \beta^2(m_1^2 | \omega_L) \beta^3(m_1^3) \phi_1(x_{11} | m_1^1, m_1^2, m_1^3) \geq 1 - \frac{\varepsilon}{(1 - \sigma)\mathbf{P}[(\omega_L, \omega_L)]}; \quad (15)$$

otherwise, by (14), A1 could induce P1 to take decision x_{11} with a probability lower than $\pi_{11}(\omega_L, \omega_L, 2)$, which would be a strict better response, a contradiction. Integrating (15) with respect to the measure $\sigma\beta^1(\cdot | \omega_H, 1) + (1 - \sigma)\beta^1(\cdot | \omega_H, 2)$ then yields

$$\begin{aligned} & \sum_{(m_1^1, m_1^2, m_1^3) \in M_1} [\sigma\beta^1(m_1^1 | \omega_H, 1) + (1 - \sigma)\beta^1(m_1^1 | \omega_H, 2)] \beta^2(m_1^2 | \omega_L) \beta^3(m_1^3) \phi_1(x_{11} | m_1^1, m_1^2, m_1^3) \\ & \geq 1 - \frac{\varepsilon}{(1 - \sigma)\mathbf{P}[(\omega_L, \omega_L)]}, \end{aligned}$$

which exactly states that, by deviating to $\beta^2(\cdot | \omega_L)$ in state (ω_H, ω_H) , A2 can ensure that P1 takes decision x_{11} with probability at least $1 - \frac{\varepsilon}{(1 - \sigma)\mathbf{P}[(\omega_L, \omega_L)]}$, hereby obtaining, because $4.5 > \sigma + 8(1 - \sigma)$ as $\sigma > \frac{1}{2}$, an expected payoff at least equal to

$$4.5 \left\{ 1 - \frac{\varepsilon}{(1 - \sigma)\mathbf{P}[(\omega_L, \omega_L)]} \right\} + [\sigma + 8(1 - \sigma)] \frac{\varepsilon}{(1 - \sigma)\mathbf{P}[(\omega_L, \omega_L)]}, \quad (16)$$

By contrast, if A2 plays $\beta^2(\cdot | \omega_H)$ in state (ω_H, ω_H) , as he must do in equilibrium, then, by the first inequality in (13), he obtains an expected payoff at most equal to

$$4.5 \frac{\varepsilon}{\mathbf{P}[(\omega_H, \omega_H)]} + [\sigma + 8(1 - \sigma)] \left\{ 1 - \frac{\varepsilon}{\mathbf{P}[(\omega_H, \omega_H)]} \right\}. \quad (17)$$

Comparing (16) and (17), and using again the fact that $4.5 > \sigma + 8(1 - \sigma)$, we obtain that this deviation is profitable for A2 for every ε satisfying (8), contradicting the assumption that $\beta \in B^*(\phi_1, \gamma_2(\sigma))$. Thus $\gamma_2(\sigma)$ satisfies (7), as claimed.

To conclude the proof, observe that, because P2 can, for any $\sigma \in (\frac{1}{2}, 1)$, guarantee herself a payoff of $5 + \frac{(1-\sigma)\mathbf{P}[(\omega_L, \omega_L)]\mathbf{P}[(\omega_H, \omega_H)]}{1-\sigma\mathbf{P}[(\omega_L, \omega_L)]}$ by posting the mechanism $\gamma_2(\sigma)$, her payoff in any PBE of G_1^{SM} must at least be equal to

$$\sup_{\sigma \in (\frac{1}{2}, 1)} 5 + \frac{(1 - \sigma)\mathbf{P}[(\omega_L, \omega_L)]\mathbf{P}[(\omega_H, \omega_H)]}{1 - \sigma\mathbf{P}[(\omega_L, \omega_L)]} = 5 + \frac{\mathbf{P}[(\omega_L, \omega_L)]\mathbf{P}[(\omega_H, \omega_H)]}{2 - \mathbf{P}[(\omega_L, \omega_L)]}.$$

The result follows. ■

The proof of Lemma 2 crucially exploits the fact that, by posting a mechanism with signals, P2 can asymmetrically inform the agents of her decision. Specifically, the mechanism $\gamma_2(\sigma)$ we construct is such that, when communicating with P1, A1 is perfectly informed of P2's decision, while A2 and A3 are kept in the dark. Such an asymmetry in the information transmitted by P2 to the agents, which is possible only under private disclosures, is precisely what enables P2 to guarantee herself a payoff strictly above her min-max-min payoff of 5 regardless of the mechanism offered by P1 and of the continuation equilibrium in the subgame played by the agents.

To see this, notice that the only way to keep P2's payoff down to 5 is for P1 to take decision x_{11} in state (ω_L, ω_L) and decision x_{12} in state (ω_H, ω_H) . However, by privately informing A1 of her decision, P2 can exploit the fact that, in state (ω_L, ω_L) , and upon learning that $x_2 = x_{22}$, A1's preferences over X_1 are perfectly aligned with hers; this guarantees that, if A1 could influence P1's decision in state (ω_L, ω_L) , she would induce P1 to take decision x_{12} with positive probability, bringing P2's payoff strictly above 5. Hence, given the other agents' messages, A1 must not be able to influence P1's decision in state (ω_L, ω_L) . A similar argument implies that, given the other agents' messages, A1 must not be able to influence P1's decision in state (ω_H, ω_H) either.⁷ Moreover, because A3 does not observe the state, his message to P1 must be the same in each state. As a result, A2 must de facto have full control over P1's decision. However, when P2 is expected to take decision x_{21} with probability $\sigma > \frac{1}{2}$, A2, without receiving further information from P2, strictly prefers to induce P1 to take decision x_{11} in both state. Hence, if she has the possibility to do so, which we just argued must be the case, she has no incentive to induce the distribution over X_1 that inflicts the min-max-min payoff of 5 on P2.

⁷Otherwise, in state (ω_H, ω_H) , and upon learning that $x_2 = x_{21}$, A1 would induce P1 to take decision x_{11} with positive probability, bringing P2's payoff again strictly above 5.

Lemma 2 generalizes this insight by constructing a lower bound for P2’s equilibrium payoff that is strictly higher than her min-max-min payoff. This lower bound is independent of the size of the signal spaces S_1^i and of the message spaces M_1^i used by P1 in G_1^{SM} . In particular, replacing all sums by the appropriate integrals in the proof of Lemma 2 reveals that this bound remains relevant even if some agent can send infinitely many messages to P1—provided, of course, an equilibrium still exists. It should be noted that, if the probability of either state goes to 1, the lower bound for P2’s equilibrium payoff converges to her minimum feasible payoff of 5. In the example, this reflects that P1, if informed of the state, can always inflict this payoff on P2. More generally, the role of private disclosures we emphasize does not extend to complete-information settings in which agents take no payoff-relevant actions. This is because, by posting a recommendation mechanism, a principal can then ensure that her decisions are only implemented through near-unanimous recommendations by the agents, independently of any signals they may receive from the other principals.

Because A1’s and A2’s preferences are perfectly aligned and A3’s payoff is constant over $X \times \Omega$, the reader may wonder why P2 would want to inform the agents in an asymmetric way. The reason is that, if the agents had the same information about P2’s decision, then they could discipline each other, which would enable them to implement incentive-compatible punishments for P2 as in Yamashita’s (2010) construction. For example, if all the agents are perfectly informed of P2’s decision, then there exists a mechanism for P1 and a continuation equilibrium in the subgame played by the agents that jointly implement the distribution over $X_1 \times \Omega$ that inflict 5 on P2. The possibility for P2 to asymmetrically inform the agents of her decision is precisely what allows her to prevent the agents from selecting a direct mechanism that punishes her in case she deviates.

Lemmas 1–2 together imply the following result.

Proposition 1 *PBE outcomes of competing-mechanism games without private disclosures need not be robust to the possibility for principals to post mechanisms with private signals. In particular, PBE payoff vectors of competing-mechanism games without private disclosures but with rich message spaces such that $D_j \times \Omega^i \subset M_j^i$ for all i and j need no longer be supported once principals can engage into private disclosures.*

Example 1 illustrates how private disclosures, taking the form of signals privately sent by the principals to the agents about the decisions they implement in response to the agents’ messages, can affect the set of equilibrium outcomes in competing-mechanism games. For

this to be the case, it is crucial that such signals be received by the agents before they send their messages to the principals. For instance, in Example 1, if P1 could make sure that agents send their messages to her after observing P2's mechanism but before receiving private signals from P2, then the outcome (2) could be supported in equilibrium, yielding P2 a payoff of 5. Intuitively, this could be achieved by letting P1 post a recommendation mechanism, as in Step 1 of the proof of Lemma 1, which would effectively neutralize P2's signals to the agents.

4 A Challenge to Universal Standard Mechanisms

In the previous section, we have shown that equilibrium outcomes of competing-mechanism games in which principals are restricted to standard mechanisms, yet with potentially rich message spaces, need not be robust to the possibility for principals to post mechanisms with signals. In this section, we address the dual question of whether competing-mechanism games in which principals can engage into private disclosures may admit equilibria whose outcomes cannot be supported when principals are restricted to standard mechanisms, no matter how rich message spaces are.

We now provide an example showing that this is indeed the case. In this example, a principal can use private signals to make the agents' messages to the other principal depend on information that correlates with her own decision. In turn, this allows the principals to correlate their decisions with the agents' exogenous private information in a way that cannot be achieved, when private disclosures are not feasible, by letting the principals randomize over their choices of mechanisms or by letting the agents randomize over the messages they send to the principals.

Example 2 Let $I = J \equiv 2$. We denote the principals by P1 and P2, and the agents by A1 and A2. The decision sets are $X_1 \equiv \{x_{11}, x_{12}, x_{13}, x_{14}\}$ for P1 and $X_2 \equiv \{x_{21}, x_{22}\}$ for P2. A2 can be of two types, with $\Omega^2 \equiv \{\omega_L, \omega_H\}$, whereas A1 can only be of a single type, which we omit from the notation for the sake of clarity. The states ω_L and ω_H are commonly believed to occur with probabilities $\mathbf{P}[\omega_L] = \frac{1}{4}$ and $\mathbf{P}[\omega_H] = \frac{3}{4}$, respectively.

The players' payoffs are represented in Tables 3 and 4 below, in which the first payoff is that of P2 and the last two payoffs are those of A1 and A2, respectively; $\zeta < 0$ is an arbitrary loss for P2. P1's payoff is constant over $X \times \Omega$ and hence plays no role in the analysis.

	x_{21}	x_{22}
x_{11}	$\zeta, 4, 1$	$\zeta, 8, 3.5$
x_{12}	$\zeta, 2, 5$	$\zeta, 9, 8$
x_{13}	$10, 3, 3$	$\zeta, 5.5, 3.5$
x_{14}	$\zeta, 1, 3.5$	$10, 7.5, 7.5$

Table 3: Payoffs in state ω_L .

	x_{21}	x_{22}
x_{11}	$\zeta, 1, 6$	$10, 7.5, 5$
x_{12}	$10, 3, 9$	$\zeta, 5.5, 6$
x_{13}	$\zeta, 8, 7$	$\zeta, 4.5, 7$
x_{14}	$\zeta, 9, 6$	$\zeta, 3, 9$

Table 4: Payoffs in state ω_H .

4.1 An Equilibrium with Private Disclosures

To illustrate the key ideas in the simplest possible manner, we consider a specific competing-mechanism game G_2^{SM} with private disclosures in which only P2 can send signals to the agents, and these signals are binary; that is, we let $S_1^1 = S_1^2 \equiv \{\emptyset\}$ and $S_2^1 = S_2^2 \equiv \{1, 2\}$. Furthermore, we consider the simplest possible message spaces that allow the agents to report their private information to the principals; that is, we let $M_1^i \equiv \Omega^i \times S_2^i$ and $M_2^i \equiv \Omega^i$ for all i .⁸ The following result then holds.

Lemma 3 For $\alpha = \frac{2}{3}$, the outcome

$$z(\omega_L) \equiv \alpha \delta_{(x_{13}, x_{21})} + (1 - \alpha) \delta_{(x_{14}, x_{22})}, \quad (18)$$

$$z(\omega_H) \equiv \alpha \delta_{(x_{12}, x_{21})} + (1 - \alpha) \delta_{(x_{11}, x_{22})}, \quad (19)$$

in which P2 obtains her maximum feasible payoff of 10, can be supported in a PBE of G_2^{SM} .

Proof. Let P2 post the mechanism $\gamma_2^* \equiv (\sigma_2^*, \phi_2^*)$ such that

$$\sigma_2^*(s_2) \equiv \begin{cases} \frac{\alpha}{2} & \text{if } s_2 = (1, 1) \\ \frac{\alpha}{2} & \text{if } s_2 = (2, 2) \\ \frac{1-\alpha}{2} & \text{if } s_2 = (1, 2) \\ \frac{1-\alpha}{2} & \text{if } s_2 = (2, 1) \end{cases}$$

and, for each $(s_2, m_2) \in S_2 \times M_2$,

$$\phi_2^*(s_2, m_2) \equiv \begin{cases} \delta_{x_{21}} & \text{if } s_2 \in \{(1, 1), (2, 2)\} \\ \delta_{x_{22}} & \text{if } s_2 \in \{(1, 2), (2, 1)\} \end{cases} \quad (20)$$

⁸As the arguments below reveal, Lemma 3 does not hinge on these simplifying assumptions, and extends to games with richer signal and message spaces as long as $\Omega^i \times S_2^i \subset M_1^i$ and $\Omega^i \subset M_2^i$ for all i .

irrespective of the messages $m_2 \in M_2$ received from the agents. A key feature of this mechanism is that, regardless of the signal he receives from P2, every agent's posterior distribution about P2's decision coincides with his prior distribution; that is, each agent believes that P2 takes decision x_{21} with probability α and decision x_{22} with probability $1 - \alpha$. For the same reason, each agent believes that the other agent received the same signal as his with probability α and a different signal with probability $1 - \alpha$. Thus γ_2^* completely leaves both agents in the dark.

As for P1, let her post the deterministic mechanism $\gamma_1^* \equiv (\delta_{(\emptyset, \emptyset)}, \phi_1^*)$ such that, for each $(m_1^1, m_1^2) \in M_1$,

$$\phi_1^*(\emptyset, \emptyset, m_1) \equiv \begin{cases} \delta_{x_{13}} & \text{if } m_1 \in \{(1, \omega_L, 1), (2, \omega_L, 2)\}, \\ \delta_{x_{14}} & \text{if } m_1 \in \{(1, \omega_L, 2), (2, \omega_L, 1)\}, \\ \delta_{x_{12}} & \text{if } m_1 \in \{(1, \omega_H, 1), (2, \omega_H, 2)\}, \\ \delta_{x_{11}} & \text{if } m_1 \in \{(1, \omega_H, 2), (2, \omega_H, 1)\}, \end{cases} \quad (21)$$

in which, for instance, $(1, \omega_L, 1)$ stands for $m_1^1 = 1$ and $m_1^2 = (\omega_L, 1)$; that is, A1 reports to P1 that she received signal $s_2^1 = 1$ from P2, whereas A2 reports that her type is ω_L and that she received signal $s_2^2 = 1$ from P2. Observe from (20)–(21) that the outcome (18)–(19) is implemented in the subgame (γ_1^*, γ_2^*) if every agent reports truthfully to P1 his type and the signal he receives from P2. We now show that truthful reporting is consistent with a BNE of the subgame (γ_1^*, γ_2^*) . The proof consists of two steps.

Step 1 Consider first A1's incentives, under the belief that A2 is truthful to P1. Because A1 has only one type, we only need to check A1's incentives to truthfully report to P1 the signal he receives from P2.

If A1 truthfully reports his signal to P1, then, regardless of the signal he receives from P2, his expected payoff is

$$\begin{aligned} & \frac{1}{4} [\alpha u^1(x_{13}, x_{21}, \omega_L) + (1 - \alpha)u^1(x_{14}, x_{22}, \omega_L)] \\ & + \frac{3}{4} [\alpha u^1(x_{12}, x_{21}, \omega_H) + (1 - \alpha)u^1(x_{11}, x_{22}, \omega_H)] = 3\alpha + 7.5(1 - \alpha). \end{aligned} \quad (22)$$

If, instead, A1 misreports his signal to P1, then, regardless of the signal he receives from P2, his expected payoff is

$$\begin{aligned} & \frac{1}{4} [\alpha u^1(x_{14}, x_{21}, \omega_L) + (1 - \alpha)u^1(x_{13}, x_{22}, \omega_L)] \\ & + \frac{3}{4} [\alpha u^1(x_{11}, x_{21}, \omega_H) + (1 - \alpha)u^1(x_{12}, x_{22}, \omega_H)] = \alpha + 5.5(1 - \alpha), \end{aligned}$$

which is strictly less than (22) for all $\alpha \in [0, 1]$.

Step 2 Consider next A2's incentives, under the belief that A1 is truthful to P1. We need to check A2's incentives to truthfully report to P1 both his type and the signal he receives from P2.

Case 1: $\omega^2 = \omega_L$ We first consider the behavior of A2 when he is of type ω_L . If A2 truthfully reports both his type and his signal to P1, then, regardless of the signal he receives from P2, his expected payoff is

$$\alpha u^2(x_{13}, x_{21}, \omega_L) + (1 - \alpha)u^2(x_{14}, x_{22}, \omega_L) = 3\alpha + 7.5(1 - \alpha). \quad (23)$$

If, instead, A2 truthfully reports his type but misreports his signal to P1, then, regardless of the signal he receives from P2, his expected payoff is

$$\alpha u^2(x_{14}, x_{21}, \omega_L) + (1 - \alpha)u^2(x_{13}, x_{22}, \omega_L) = 3.5,$$

which is at most equal to (23) if $\alpha \leq \frac{8}{9}$.

Next, if A2 misreports his type but truthfully reports his signal to P1, then, regardless of the signal he receives from P2, his expected payoff is

$$\alpha u^2(x_{12}, x_{21}, \omega_L) + (1 - \alpha)u^2(x_{11}, x_{22}, \omega_L) = 5\alpha + 3.5(1 - \alpha),$$

which is at most equal to (23) if $\alpha \leq \frac{2}{3}$.

Finally, if A2 misreports both his type and his signal to P1, then, regardless of the signal he receives from P2, his expected payoff is

$$\alpha u^2(x_{11}, x_{21}, \omega_L) + (1 - \alpha)u^2(x_{12}, x_{22}, \omega_L) = \alpha + 8(1 - \alpha),$$

which is at most equal to (23) if $\alpha \geq \frac{1}{5}$.

Case 2: $\omega^2 = \omega_H$ We next consider the behavior of A2 when he is of type ω_H . If A2 truthfully reports both his type and his signal to P1, then, regardless of the signal he receives from P2, his expected payoff is

$$\alpha u^2(x_{12}, x_{21}, \omega_H) + (1 - \alpha)u^2(x_{11}, x_{22}, \omega_H) = 9\alpha + 5(1 - \alpha). \quad (24)$$

If, instead, A2 truthfully reports his type but misreports his signal to P1, then, regardless of the signal he receives from P2, his expected payoff is

$$\alpha u^2(x_{11}, x_{21}, \omega_H) + (1 - \alpha)u^2(x_{12}, x_{22}, \omega_H) = 6,$$

which is at most equal to (24) if $\alpha \geq \frac{1}{4}$.

Next, if A2 misreports his type but truthfully reports his signal to P1, then, regardless of the signal he receives from P2, his expected payoff is

$$\alpha u^2(x_{13}, x_{21}, \omega_H) + (1 - \alpha)u^2(x_{14}, x_{22}, \omega_H) = 7\alpha + 9(1 - \alpha),$$

which is at most equal to (24) if $\alpha \geq \frac{2}{3}$.

Finally, if A2 misreports both his type and his signal to P1, then, regardless of the signal he receives from P2, his expected payoff is

$$\alpha u^2(x_{14}, x_{21}, \omega_H) + (1 - \alpha)u^2(x_{13}, x_{22}, \omega_H) = 6\alpha + 7(1 - \alpha),$$

which is at most equal to (24) if $\alpha \geq \frac{2}{5}$.

The above analysis implies that it is a BNE for A1 and A2 to truthfully report their private information to P1 in the subgame (γ_1^*, γ_2^*) if and only if $\alpha = \frac{2}{3}$. In this continuation equilibrium, P2 obtains her maximal feasible payoff of 10. Because P1's payoff is constant over $X \times \Omega$, there exists a PBE of G_2^{SM} in which P1 and P2 post the mechanisms γ_1^* and γ_2^* , and A1 and A2 play any BNE in any subgame following a deviation by P1 or P2—the existence of such an equilibrium being guaranteed by the fact that all these subgames are finite. The result follows. ■

Observe for future reference that, in this equilibrium, A1 obtains an expected payoff of 4.5, while A2 obtains an expected payoff of 4.5 if he is of type ω_L and an expected payoff of $\frac{23}{3}$ if he is of type ω_H .

Our equilibrium construction relies on the fact that, although the mechanism with signals γ_2^* is publicly disclosed to both A1 and A2, A1 and A2 receive different signals from P2. Specifically, private disclosures by P2 take the form of *encryption keys*: taken in isolation, each signal sent by P2 is completely uninformative of her decision, whereas, taken together, the two signals sent by P2 are perfectly informative of her decision; we will return to this point in Section 5. Notice in that respect that, if P2 were to inform the agents of her decision, then, after learning that P2 takes decision x_{21} , A2, when of type ω_L , would no longer be willing to induce P1 to take decision x_{13} .

Our construction also reveals that, for P2 to obtain her maximum feasible payoff of 10 while maintaining the agents' incentives, it is essential that both principals randomize over their decisions but do so in a perfectly correlated manner. Whereas it is technically feasible to achieve the equilibrium correlation between the principals' decisions by letting the agents randomize over the messages they send to the principals, while letting the principals

respond deterministically to the messages they receive from the agents, such a delegation is not incentive-compatible. It is thus essential that the randomization be carried out by the principals themselves. The correlation between the principals' decisions then requires that some information be passed on from one principal to the other, which, in the absence of a direct communication channel between the principals, is possible only through private disclosures. The analysis in Section 4.2 and the discussion in Section 5 confirm this intuition by showing that private signals are indispensable, no matter how rich message spaces are.

4.2 Indispensability of Private Signals

We now show that the outcome (18)–(19) for $\alpha = \frac{2}{3}$ cannot be supported in any PBE of any game in which principals are restricted to posting standard mechanisms, irrespective of the size of message spaces. That is, private disclosures are indispensable to support this outcome. To this end, we consider a general competing-mechanism game G_2^M without private disclosures, and with arbitrary message spaces M_i^j . This general formulation notably allows us to capture the case where every principal j 's message spaces are large enough—namely, uncountable Polish spaces—to encode the agents' information about the mechanism posted by her opponent, as in Epstein and Peters (1999).

Lemma 4 *There exists no PBE of G_2^M that supports the outcome (18)–(19) for $\alpha = \frac{2}{3}$.*

Proof. The arguments below more generally show that there is no joint probability measure $\mu \in \Delta(\Phi_1 \times \Phi_2)$ over $\mathcal{F}_1 \otimes \mathcal{F}_2$ and no equilibrium strategies $\lambda \equiv (\lambda^1, \lambda^2)$ for the agents that support the outcome (18)–(19) for any value of $\alpha \in [0, 1]$. In particular, we do not require that μ be a product measure. In other words, we allow the principals to coordinate their choice of a mechanism through arbitrary correlation devices. The proof is by contradiction, and consists of four steps.

Step 1 Observe first that, with probability 1, μ must select a pair of mechanisms $\phi \equiv (\phi_1, \phi_2)$ such that, in the subgame ϕ , the equilibrium behavior strategies $(\lambda^1(\phi), \lambda^2(\phi))$ support an outcome of the form

$$\begin{aligned} z^\phi(\omega_L) &\equiv \alpha_L^\phi \delta_{(x_{13}, x_{21})} + (1 - \alpha_L^\phi) \delta_{(x_{14}, x_{22})}, \\ z^\phi(\omega_H) &\equiv \alpha_H^\phi \delta_{(x_{12}, x_{21})} + (1 - \alpha_H^\phi) \delta_{(x_{11}, x_{22})}, \end{aligned}$$

for some $(\alpha_L^\phi, \alpha_H^\phi) \in [0, 1] \times [0, 1]$. Otherwise, with positive probability, P2 would incur a loss ζ , and his overall payoff would be strictly less than 10, a contradiction. Thus, for μ -almost

every ϕ and for $(\lambda^1(\phi), \lambda^2(\phi))$ -almost every message profile (m^1, m^2) sent by the agents under the equilibrium behavior strategies $(\lambda^1(\phi), \lambda^2(\phi))$, the lotteries $(\phi_1(m_1), \phi_2(m_2))$ over the principals' decisions must be degenerate.

Step 2 Consider then a subgame ϕ as in Step 1. We first claim that $\alpha_L^\phi \leq \alpha_H^\phi$. To see this, notice that, as A1 does not know which state prevails, it must be that, given A1's state-independent behavior strategy $\lambda^1(\phi)$, the state-dependent outcomes $z^\phi(\omega_L)$ and $z^\phi(\omega_H)$ are induced by A2's state-dependent behavior strategies $\lambda^2(\phi)(\cdot|\omega_L)$ and $\lambda^2(\phi)(\cdot|\omega_H)$. Then, for type ω_L of A2 to induce $z^\phi(\omega_L)$ instead of $z^\phi(\omega_H)$, it must be that

$$3\alpha_L^\phi + 7.5(1 - \alpha_L^\phi) \geq 5\alpha_H^\phi + 3.5(1 - \alpha_H^\phi). \quad (25)$$

Similarly, for type ω_H of A2 to induce $z^\phi(\omega_H)$ instead of $z^\phi(\omega_L)$, it must be that

$$9\alpha_H^\phi + 5(1 - \alpha_H^\phi) \geq 7\alpha_L^\phi + 9(1 - \alpha_L^\phi). \quad (26)$$

The two inequalities are satisfied only if $\alpha_L^\phi \leq \alpha_H^\phi$, as claimed. Because

$$\int \alpha_L^\phi \mu(d\phi) = \alpha = \int \alpha_H^\phi \mu(d\phi), \quad (27)$$

we obtain that, with μ -probability 1, $\alpha_L^\phi = \alpha_H^\phi$. Substituting $\alpha_L^\phi = \alpha_H^\phi$ into (25)–(26), it follows that, with μ -probability 1, $\alpha_L^\phi = \alpha_H^\phi = \frac{2}{3}$ and, hence, by (27), that $\alpha = \frac{2}{3}$. We thus conclude that, with probability 1, μ must select a pair of mechanisms ϕ such that the agents' equilibrium behavior strategies $(\lambda^1(\phi), \lambda^2(\phi))$ implement the outcome (18)–(19) for $\alpha = \frac{2}{3}$, yielding types ω_L and ω_H of A2 expected payoffs of 4.5 and $\frac{23}{3}$, respectively.

Step 3 Now, observe that, for $\lambda^2(\phi|\omega_H) \otimes \lambda^2(\phi|\omega_H)$ -almost every pair of message profiles (m^2, \hat{m}^2) sent by A2 in state ω_H , if A2 were to send the messages m_1^2 to P1 and the message \hat{m}_2^2 to P2, then, given A1's behavior strategy $\lambda^1(\phi)$, the mechanisms (ϕ_1, ϕ_2) would induce the following marginal distributions over $(x_{11}, x_{12}, x_{21}, x_{22})$:

$$\begin{aligned} \Pr(x_{11}, x_{21}) + \Pr(x_{11}, x_{22}) &= \frac{1}{3}, \\ \Pr(x_{12}, x_{21}) + \Pr(x_{12}, x_{22}) &= \frac{2}{3}, \\ \Pr(x_{11}, x_{21}) + \Pr(x_{12}, x_{21}) &= \frac{2}{3}, \\ \Pr(x_{11}, x_{22}) + \Pr(x_{12}, x_{22}) &= \frac{1}{3}. \end{aligned}$$

It is easy to check that this system has not full rank, and admits a continuum of solutions indexed by $p \equiv \Pr(x_{11}, x_{21}) = \Pr(x_{12}, x_{22})$, which allows us to write $\Pr(x_{11}, x_{22}) = \frac{1}{3} - p$

and $\Pr(x_{12}, x_{21}) = \frac{2}{3} - p$. Now, if type ω_L of A2 were to follow the same behavior, thus sending the messages m_1^2 and \hat{m}_2^2 according to the strategy described above, he would obtain an expected payoff of

$$p + 8p + 5\left(\frac{2}{3} - p\right) + 3.5\left(\frac{1}{3} - p\right) = 4.5 + 0.5p.$$

Because this payoff must at most be equal to his equilibrium payoff of 4.5, it follows that $p = 0$. This implies that, for $\lambda^1(\phi)$ -almost every m^1 , we have

$$(\phi_1(m_1^1, m_1^2), \phi_2(m_2^1, \hat{m}_2^2)) \in \{\delta_{(x_{11}, x_{22})}, \delta_{(x_{12}, x_{21})}\}.$$

But, for $\lambda^1(\phi)$ -almost every m_1 , we have

$$\begin{aligned} (\phi_1(m_1^1, m_1^2), \phi_2(m_2^1, m_2^2)) &\in \{\delta_{(x_{11}, x_{22})}, \delta_{(x_{12}, x_{21})}\}, \\ (\phi_1(m_1^1, \hat{m}_1^2), \phi_2(m_2^1, \hat{m}_2^2)) &\in \{\delta_{(x_{11}, x_{22})}, \delta_{(x_{12}, x_{21})}\}, \end{aligned}$$

and thus, as decisions are perfectly correlated,

$$(\phi_1(m_1^1, m_1^2), \phi_2(m_2^1, m_2^2)) = (\phi_1(m_1^1, \hat{m}_1^2), \phi_2(m_2^1, \hat{m}_2^2)). \quad (28)$$

Because this property is satisfied for $\lambda^2(\phi|\omega_H) \otimes \lambda^2(\phi|\omega_H)$ -almost every (m^2, \hat{m}^2) , we can conclude from Fubini's theorem (Aliprantis and Border (2006, Theorem 11.27)) that (28) holds for $\lambda^1(\phi) \otimes \lambda^2(\phi|\omega_H) \otimes \lambda^2(\phi|\omega_H)$ -almost every (m^1, m^2, \hat{m}^2) . From Fubini's theorem again, it follows that for $\lambda^1(\phi)$ -almost every m_1 , (28) holds for $\lambda^2(\phi|\omega_H) \otimes \lambda^2(\phi|\omega_H)$ -almost every (m_2, \hat{m}_2) , so that the mapping $(m_2^2, \hat{m}_2^2) \mapsto (\phi_1(m_1^1, m_1^2), \phi_2(m_2^1, m_2^2))$ is constant over a set of $\lambda^2(\phi|\omega_H)$ -measure 1.

Step 4 We are now ready to complete the proof. The upshot from Step 3 is that A1 can force the decision when the state is ω_H . This implies that M^1 should include a message profile allowing A1 to implement $\delta_{(x_{11}, x_{22})}$ regardless of the message sent in equilibrium by A2. By sending this message, A1 can achieve a payoff of 7.5 when the state is ω_H . Thus, he can guarantee himself an expected payoff of at least $\frac{3}{4} \times 7.5$, which is higher than his equilibrium payoff of 4.5, a contradiction. The result follows. \blacksquare

The reason why the outcome (18)–(19) cannot be supported with standard mechanisms is the following. First, because the principals' decisions are perfectly correlated in both states, the mechanisms posted by the principals must respond deterministically to the messages sent by the agents (Step 1 in the proof of Lemma 4). Second, because only A2 observes the state, and because the distribution over the principals' decisions in (18)–(19) is state-dependent,

A2 must weakly prefer the distribution of messages he is supposed to carry out in each state to the one he is supposed to carry out in the other state. This constrains the joint distributions over principals' decisions that can be induced in the two states (Step 2 in the proof of Lemma 4). Third, for A2 to weakly prefer the distribution over the principals' decisions he is supposed to induce in state ω_L to the one that he could induce by mismatching the principals' decisions while preserving their marginal distributions in state ω_H , it must be that, given the mechanisms posted in equilibrium, the messages that A2 sends in state ω_H have no influence on the principals' decisions when combined with those sent with positive probability by A1 (Step 3 in the proof of Lemma 4). But then A1 has a profitable deviation (Step 4 in the proof of Lemma 4).

It should be noted that the result in Lemma 4 holds no matter how rich message spaces are. Hence, it also applies to the Epstein and Peters (1999) class of universal mechanisms, which, while they allow the agents to communicate all their market information to the principals, nonetheless remain standard mechanisms.

Lemmas 3–4 together imply the following result.

Proposition 2 *There exist allocations that can be supported in a PBE of a competing-mechanism game with private disclosures but that cannot be supported in any PBE of any competing-mechanism game without private disclosures—including, in particular, the game in which principals can post universal mechanisms—and this, even if the principals or the agents play mixed strategies in equilibrium.*⁹

This result shows that the universal mechanisms of Epstein and Peters (1999) fail to be canonical when principals can engage into private disclosures, that is, when they can send private signals to the agents about their decisions as a way of correlating their decisions with those of the other principals and with the agents' exogenous private information. Incidentally, it also shows that the folk theorem of Yamashita (2010) does not extend to stochastic allocations. Indeed, the allocation $(z(\omega_L), z(\omega_H))$ defined by (18)–(19) is certainly incentive-compatible; moreover, it yields P2 her maximum feasible payoff of 10, which is certainly at least equal to her min-max-min payoff, as defined by Yamashita (2010) over recommendation mechanisms. Yet Lemma 4 implies that this allocation cannot be supported in an equilibrium of G_2^M , even when $D_j \times \Omega^i \subset M_j^i$ for all i and j , so that recommendation mechanisms are feasible.

The proof of Lemma 4 does not make use of the property that the principals choose their mechanisms independently. The result in Lemma 4 thus carries over to the case where

⁹Epstein and Peters (1999) restrict attention to PBEs in which principals play pure strategies.

G_2^M is augmented by arbitrarily rich public randomizing devices that the principals may use to correlate their choices of mechanisms. On the other hand, the result does hinge on the principals not having access to private randomizing devices whose realizations are not known to the agents at the time they send their messages to the principals. Because the only role of the private disclosures in the example is to pass information from one principal to the other without changing the agents' beliefs, such private disclosures can be dispensed with if the principals can communicate directly with each other or have access to private correlation devices—that is, to devices whose realization is determined after the agents have sent their messages but before the principals have selected their payoff-relevant decisions.

5 Discussion

5.1 Informative versus Uninformative Signals

In this section, we discuss the different roles played by private disclosures in our examples; specifically, each example illustrates a key aspect of private signals.

In Example 2, private signals are used by P2 *on path* to correlate her decisions with P1's decisions and the agents' private information in a way that cannot be achieved in equilibrium with standard mechanisms. As we show in the proof of Lemma 3, the private signals that P2 may use to this end need not modify the agents' beliefs. Instead, they can work as pure encryption keys: in isolation, each key is completely uninformative of P2's decision but, taken together, the keys perfectly reveal it.

By contrast, in Example 1, the main thrust of private signals is the destabilizing role they play *off path* in undermining the robustness of equilibria in standard mechanisms. As we show in the proof of Lemma 2, P2 can guarantee herself a payoff strictly above her minimum feasible payoff by asymmetrically informing the agents about her decision, changing A1's beliefs before A1 has the opportunity to communicate with P1 while leaving A2 and A3 completely in the dark. In the discussion of Lemma 2, we argued that, if P2 were to perfectly inform all agents of her decision, or, more generally, of the decisions taken in response to the agents' messages—as in a standard mechanism—then it would be possible for P1 to post a mechanism that would inflict on P2 her minimum feasible payoff. We now show that the same conclusion is true of any signal structure that leaves all agents in the dark. Formally, we show that the analogue of Lemma 2 is false if P2 is restricted to posting mechanisms in which private signals take the form of uninformative encryption keys as those used in the context of Example 2.

To see this, consider the game G_1^{SM} studied in Claim 2; moreover, as in Lemma 1, assume that $D_j \times \Omega^i \subset M_j^i$ for all i and j , so that recommendation mechanisms are feasible, and assume that all the message spaces M_j^i are finite. We say that a mechanism $\gamma_2 \equiv (\sigma_2, \phi_2)$ of P2 has *uninformative signals* if

$$\sum_{s_2^{-i} \in S_2^{-i}} \sigma_2(s_2^{-i} | s_2^i) \phi_2(x_2 | s_2^i, s_2^{-i}, m_2) = \sum_{s_2 \in S_2} \sigma_2(s_2) \phi_2(x_2 | s_2, m_2) \quad (29)$$

for all i , $s_2^i \in S_2^i$, $m_2 \in M_2$, and $x_2 \in X_2$.¹⁰ That is, the signals s_2^i sent by P2 to any given agent i do not reveal to him anything about P2's effective decision rule $\phi_2(\cdot | s_2, \cdot)$. The following result then holds.

Lemma 5 *In G_1^{SM} , if P2 posts a recommendation mechanism ϕ_1^r , then, for every mechanism γ_2 of P2 that has uninformative signals, there exists a BNE of the subgame (ϕ_1^r, γ_2) in which P2 obtains her minimum feasible payoff of 5.*

Proof. Because A3's payoff is constant over $X \times \Omega$ and A1's and A2's payoff functions are identical, we can focus on A1's incentives. Suppose that, in the subgame (ϕ_1^r, γ_2) , A2 and A3 play behavior strategies β^2 and β^3 that prescribe the same play for any signals s_2^2 and s_2^3 they may receive from P2, respectively; that is, for each $\omega^2 \in \Omega^2$, $\beta^2(\cdot | s_2^2, \omega^2)$ is independent of s_2^2 , and similarly $\beta^3(\cdot | s_2^3)$ is independent of s_2^3 . Then, because every signal A1 receives from P2 is uninformative, A1 may as well best respond by playing a behavior strategy β^1 that prescribes the same play for any signal s_2^1 he may receive from P2; that is, for each $\omega^1 \in \Omega^1$, $\beta^1(\cdot | s_2^1, \omega^1)$ is independent of s_2^1 . Because all the message spaces M_j^i are finite, this implies that the subgame (ϕ_1^r, γ_2) admits a BNE in which all agents play behavior strategies that prescribe the same play for any signals they may receive from P2. According to (29), any such BNE of the subgame (ϕ_1^r, γ_2) can be straightforwardly turned into a BNE of the subgame $(\phi_1^r, \hat{\phi}_2)$ in which P1 posts the recommendation mechanism ϕ_1^r and P2 posts the standard mechanism $\hat{\phi}_2$ defined by

$$\hat{\phi}_2(x_2 | m_2) \equiv \sum_{s_2 \in S_2} \sigma_2(s_2) \phi_2(x_2 | s_2, m_2)$$

for all $m_2 \in M_2$ and $x_2 \in X_2$. Notice that, by construction, the same outcome is implemented in either case. Conversely, any BNE of the subgame $(\phi_1^r, \hat{\phi}_2)$ can be straightforwardly turned into a BNE of the subgame (ϕ_1^r, γ_2) in which all agents play behavior strategies that prescribe the same play for any signals they may receive from P2, and which implements the same

¹⁰That S_2 is a finite set plays no role for this discussion. In particular, Lemma 5 below remains valid if S_2 is an arbitrary Polish space.

outcome. To conclude, observe that, as $\hat{\phi}_2$ is a standard mechanism, we know from the proof of Lemma 1 that the subgame $(\phi_1^r, \hat{\phi}_2)$ admits a BNE in which P2 obtains a payoff of 5. The result follows. ■

This last result reflects that, if P1 posts a recommendation mechanism ϕ_1^r and P2 posts a mechanism γ_2 with uninformative signals, then there exists a one-to-one correspondence between the babbling equilibria of the subgame (ϕ_1^r, γ_2) , in which the agents ignore the signals they receive from P2, and the equilibria of the subgame $(\phi_1^r, \hat{\phi}_2)$ in which P2 posts the standard mechanism $\hat{\phi}_2$ obtained by averaging γ_2 over the profiles of signals s_2 . But then, because, by Lemma 1, P2's payoff can be kept down to 5 in the latter case, this must also be the true in the former case. Notice that there is no tension between this result and Lemma 3, which illustrates the power of mechanisms with uninformative signals in the context of Example 2; indeed, the key step in the proof of Lemma 3 precisely consists in constructing a non-babbling equilibrium of the agents' subgame in which they truthfully report to P1 the uninformative signals they receive from P2.

However, Lemma 5 points at a potential drawback of mechanisms with uninformative signals, namely, that they naturally lend themselves to babbling equilibria: if all agents but one ignore their uninformative signals, then the remaining agent may as well do the same. This contrasts with the mechanism constructed in the proof of Lemma 2, which, by disclosing P2's decision asymmetrically to the agents, allows P2 to guarantee herself more than her min-max-min payoff regardless of the mechanism posted by P1 and of the agents' continuation equilibrium. The general lesson that Lemma 2 thus illustrates is that, if a principal deviates and makes an informative private disclosure about her decision rule to an agent, then the agent cannot simply ignore it; this reasoning, of course, fully exploits the logic of sequential rationality and the standard assumption that the agents treat the mechanisms posted by the principals at face value.

5.2 Contractible Contracts and Reciprocal Mechanisms

The nonrobustness result in Proposition 1 is established under the assumption that no principal can directly condition the decisions she takes and/or the mechanism she chooses on the other principals' decisions and/or mechanisms. However, the result extends to settings in which such conditioning is feasible, as in the literature on contractible contracts (Kalai, Kalai, Lehrer, Samet (2010), Peters and Szentes (2012), Szentes (2015)) and reciprocal mechanisms (Peters (2015)). To see this, observe that, in Example 1, the only way to inflict on P2 her minimum feasible payoff of 5 is for P1 to take decision x_{11} in state (ω_L, ω_L) and decision x_{12}

in state (ω_H, ω_H) with probability 1. However, because the state is privately observed by A1 and A2, P1 must ultimately let them determine which decisions to implement in response to a deviation by P2, were P2 to post a mechanism with signals. Now suppose, for instance, that, as in the proof of Lemma 2, P2 posts a mechanism whereby she selects a decision at random and only informs A1 of her decision. Because this mechanism perfectly aligns P2's and A1's preferences in each state, P1's mechanism must not be responsive to A1's messages on pain of moving P2's payoff away from 5; notice that this remains true even if P1 can condition the decision she takes and/or the mechanism she chooses on P2's decision and/or mechanism. Thus P1 must entirely delegate to A2 the task of making her decision contingent on the state. Yet, by construction, A2 does not know which decision P2 is committed to; moreover, the additional possibility for P1 to, for instance, condition her decision on P2's is of little use if P2's payoff is to be kept down to 5, as this requirement uniquely pins down P1's decision in each state. It follows that P1 still has no way to reward A2 for truthfully reporting the state to her, and, along the same lines as in the proof of Lemma 2, that A2 has a profitable deviation. We conclude that, even if P1 can resort to contractible contracts or post a reciprocal mechanism, it is impossible for her to induce A1 and A2 to carry out the punishments necessary to block P2's deviation.

In this respect, the case of Proposition 2 is less clear cut. On the one hand, in Example 2, the proof of Lemma 4 goes through unaltered even if each principal can condition her choice of a mechanism on the other principal's mechanism—indeed, the argument is valid for any pair of mechanisms posted by the principals, no matter the way in which this pair is generated. This implies that Proposition 2 extends to settings in which principals can post reciprocal mechanisms (Peters (2015)). On the other hand, Lemma 4 does hinge on the impossibility for each principal to directly condition her decision on the other principal's decision. Indeed, if this were possible, or if the principals could use randomization devices whose realization is not known to the agents at the time they send their messages to the principals, then the outcome (18)–(19) could be supported even without private disclosures. The purpose of Example 2 is rather to show that, in the setting that has received most of the attention in the literature—in which direct communication between the principals is not feasible, private randomizing devices are not available, and principals cannot condition their decisions on other principals' decisions—private disclosures may have important implications for the set of equilibrium outcomes and for the validity of the universal revelation principle established for such a setting.

6 Concluding Remarks

Private disclosures, that is, signals about the decisions implemented in a mechanism sent by the principals to the agents before the latter send their messages back to the principals, have been ignored in previous work. They have profound effects on the equilibrium allocations of competing-mechanism games. They have a bearing for the canonicity of the universal mechanisms discussed in the literature and the validity of the folk theorems established for such games.

An open question is what structure for the signal and message space is fully canonical, meaning that (1) it permits one to sustain all equilibrium allocations of competing-mechanism games with arbitrarily richer message and signal spaces, and (2) it guarantees that the equilibrium outcomes of the canonical game are robust to deviations to mechanisms with richer message and signal spaces. Identifying a canonical extensive form, a canonical class of mechanisms, and a set of fully robust equilibrium payoffs for such games are important next steps for this line of work.

Appendix

Proof of Lemma 1. We start with a definition. An *extended recommendation mechanism* $\tilde{\phi}_j^r : M_j \mapsto \Delta(X_j)$ for principal j implements the same decisions as the recommendation mechanism ϕ_j^r in (1), except if at least $I - 1$ agents i send messages $m_j^i \equiv (d_j^0, \omega^i) \in D_j \times \Omega^i$ to principal j , for some fixed direct mechanism $d_j^0 \in D_j$, in which case principal j disregards d_j^0 and implements a (possibly stochastic) direct mechanism $\tilde{d}_j : \Omega \rightarrow \Delta(X_j)$; again, if some agent i sends a message $m_j^i \notin D_j \times \Omega^i$ to principal j , then $\tilde{\phi}_j^r$ treats this message as if it coincided with some fixed element $(\bar{d}_j, \bar{\omega}_j^i)$ of $D_j \times \Omega^i$, for some $\bar{d}_j \neq d_j^0$.

Step 2 We now construct a family of PBEs of G_1^M , indexed by P2's payoff $v \in (5, 6]$, in which P2 posts the same recommendation mechanism ϕ_2^r as in Step 1 of the proof and P1 posts an extended recommendation mechanism $\tilde{\phi}_1^r$. Again, because P1's payoff is constant over $X \times \Omega$, she has no profitable deviation. If P2 deviates to some arbitrary standard mechanism $\phi_2 : M_2 \rightarrow \Delta(X_2)$, then we require that the agents' strategies implement the same punishments for P2 as in Step 1 of the proof. We suppose in particular that the direct mechanism d_1^0 differs from the direct mechanisms d_1^* and d_1 , defined by (3) and (6), which may be recommended by the agents to P1 following a deviation by P2; recall that these punishments inflict on P2 her minimum feasible payoff of 5. We consider two cases in turn.

Case 1: $v \in (5, 5.5]$ We specify $\tilde{\phi}_1^r$ as follows. First, we assume that $\bar{\omega}_1^1 = \bar{\omega}_1^2 = \omega_L$, so that, if some agent $i = 1, 2$ sends a message $m_1^i \notin D_1 \times \Omega^i$ to P1, then $\tilde{\phi}_1^r$ treats this message as if agent i reported to principal j to be of type ω_L ; recall from Step 1 of the proof that ϕ_2^r similarly satisfies $\bar{\omega}_2^1 = \bar{\omega}_2^2 = \omega_L$. Fixing some $\xi \in [\frac{1}{2}, 1)$, we then let

$$\tilde{d}_1(\omega) \equiv \begin{cases} \tilde{x}_1^\xi & \text{if } \omega = (\omega_L, \omega_L) \\ \tilde{x}_1^{1-\xi} & \text{otherwise} \end{cases},$$

where

$$\tilde{x}_1^\xi \equiv \xi \delta_{x_{11}} + (1 - \xi) \delta_{x_{12}} \quad \text{and} \quad \tilde{x}_1^{1-\xi} \equiv (1 - \xi) \delta_{x_{11}} + \xi \delta_{x_{12}}.$$

We now show that, for each $\xi \in [\frac{1}{2}, 1)$, the subgame $(\tilde{\phi}_1^r, \phi_2^r)$ admits a BNE in which: (i) each agent recommends to P1 the direct mechanism d_1^0 , and recommends to P2 the direct mechanism d_2^* defined by (3); (ii) A1 and A2 truthfully report their types to P1 and P2. The corresponding payoff for P2 in the subgame $(\tilde{\phi}_1^r, \phi_2^r)$ is $v = 6 - \xi \in (5, 5.5]$ as ξ varies in $[\frac{1}{2}, 1)$, as desired. Because A3's payoff is constant over $X \times \Omega$, we only need to focus on A1's and A2's incentives.

Consider first state (ω_L, ω_L) , and suppose that A2 and A3 recommend d_1^0 to P1 and d_2^* to P2, and that A2 truthfully reports his type to P1 and P2. Because A1 is not pivotal, recommending a different direct mechanism to either principal is of no avail to him; moreover, because $\bar{\omega}_1^1 = \bar{\omega}_2^1 = \omega_L$, sending a message $m_j^1 \notin D_j \times \Omega^1$ to any principal j amounts for A1 to truthfully reporting his type to her. We can thus with no loss of generality assume that A1 recommends d_1^0 to P1 and d_2^* to P2, and we only need to study A1's reporting decisions. (1) If A1 truthfully reports his type to P1 and P2, then P1 implements the lottery \tilde{x}_1^ξ , P2 takes decision x_{21} , and A1 obtains a payoff of $8\xi + 4.5(1 - \xi)$. (2) If A1 truthfully reports his type to P1 and misreports his type to P2, then P1 implements the lottery \tilde{x}_1^ξ , P2 takes decision x_{22} , and A1 obtains a payoff of $\xi + 4.5(1 - \xi) < 8\xi + 4.5(1 - \xi)$. (3) If A1 misreports his type to P1 and truthfully reports his type to P2, then P1 implements the lottery $\tilde{x}_1^{1-\xi}$, P2 takes decision x_{12} , and A1 obtains a payoff of $8(1 - \xi) + 4.5\xi \leq 8\xi + 4.5(1 - \xi)$ as $\xi \geq \frac{1}{2}$. (4) Finally, if A1 misreports his type to P1 and P2, then P1 implements the lottery $\tilde{x}_1^{1-\xi}$, P2 takes decision x_{22} , and A1 obtains a payoff of $1 - \xi + 4.5\xi < 8\xi + 4.5(1 - \xi)$. Thus A1 has no incentive to deviate from his candidate equilibrium strategy in state (ω_L, ω_L) , and neither has A2 by symmetry.

Consider next state (ω_H, ω_H) , and suppose that A2 and A3 recommend d_1^0 to P1 and d_2^* to P2, and that A2 truthfully reports his type to P1 and P2. Then P1 implements the lottery $\tilde{x}_1^{1-\xi}$ and P2 takes decision x_{22} regardless of the reports and/or messages of A1 to P1 and P2. Thus A1 has no incentive to deviate from his candidate equilibrium strategy in state (ω_H, ω_H) , and neither has A2 by symmetry. This concludes the discussion of Case 1.

Case 2: $v \in (5.5, 6]$ We specify $\tilde{\phi}_1^r$ as follows. First, we assume that $\bar{\omega}_1^1 = \bar{\omega}_1^2 = \omega_H$, so that, if some agent $i = 1, 2$ sends a message $m_1^i \notin D_1 \times \Omega^i$ to P1, then $\tilde{\phi}_1^r$ treats this message as if agent i reported to principal j to be of type ω_H ; the corresponding property for ϕ_2^r is irrelevant for the following arguments. Fixing some $\xi \in [\frac{1}{2}, 1]$, we then let

$$\tilde{d}_1(\omega) \equiv \begin{cases} \tilde{x}_1^\xi & \text{if } \omega = (\omega_H, \omega_H) \\ \tilde{x}_1^{1-\xi} & \text{otherwise} \end{cases},$$

where the lotteries \tilde{x}_1^ξ and $\tilde{x}_1^{1-\xi}$ are defined as in Case 1.

We now show that, for each $\xi \in (\frac{1}{2}, 1]$, the subgame $(\tilde{\phi}_1^r, \phi_2^r)$ admits a BNE in which: (i) each agent recommends to P1 the direct mechanism d_1^0 , and recommends to P2 the direct mechanism d_2^{**} that selects the decision x_{21} regardless of A1's and A2's reports; (b) A1 and A2 truthfully report their types to P1—because P2's decision is fixed, the messages they send to P2 are irrelevant. The corresponding payoff for P2 in the subgame $(\tilde{\phi}_1^r, \phi_2^r)$ is $v = 5 + \xi \in (5.5, 6]$ as ξ varies in $(\frac{1}{2}, 1]$, as desired. Because A3's payoff is constant over

$X \times \Omega$, we only need to focus on A1's and A2's incentives.

Consider first state (ω_H, ω_H) , and suppose that A2 and A3 recommend d_1^0 to P1 and d_2^{**} to P2, and that A2 truthfully reports his type to P1. Because A1 is not pivotal, recommending a different direct mechanism to either principal is of no avail to him; moreover, because $\bar{\omega}_1^1 = \omega_H$, sending a message $m_1^1 \notin D_1 \times \Omega^1$ to P1 amounts for A1 to truthfully reporting his type to her. We can thus with no loss of generality assume that A1 recommends d_1^0 to P1 and d_2^{**} to P2, and we only need to study A1's reporting decisions. (1) If A1 truthfully reports his type to P1, then P1 implements the lottery \tilde{x}_1^ξ and A1 obtains a payoff of $4.5\xi + 1 - \xi$. (2) If A1 misreports his type to P1, then P1 implements the lottery $\tilde{x}_1^{1-\xi}$ and A1 obtains a payoff of $4.5(1 - \xi) + \xi < 4.5\xi + 1 - \xi$ as $\xi > \frac{1}{2}$. Thus A1 has no incentive to deviate from his candidate equilibrium strategy in state (ω_H, ω_H) , and neither has A2 by symmetry.

Consider next state (ω_L, ω_L) , and suppose that A2 and A3 recommend d_1^0 to P1 and d_2^{**} to P2, and that A2 truthfully reports his type to P1. Then P1 implements the lottery $\tilde{x}_1^{1-\xi}$ regardless of A1's reports and/or messages to P1. Thus A1 has no incentive to deviate from his candidate equilibrium strategy in state (ω_L, ω_L) , and neither has A2 by symmetry. This concludes the discussion of Case 2.

To conclude the proof, observe that, because P1's payoff is constant over $X \times \Omega$, she has no profitable deviation, and that any deviation by P2 to some arbitrary standard mechanism $\phi_2 : M_2 \rightarrow \Delta(X_2)$ can be punished as in Step 1 of the proof, yielding her her minimum feasible payoff of 5, so that she has no profitable deviation either. The result follows. ■

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, T. Mariotti, and G. Piaser (2021): “Competing Mechanisms and Folk Theorems: Two Examples,” *Games and Economic Behavior*, 125, 79–93.
- [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] Epstein, L.G., and M. Peters (1999): “A Revelation Principle for Competing Mechanisms,” *Journal of Economic Theory*, 88(1), 119–160.
- [6] Hurwicz, L. (1973): “The Design of Mechanisms for Resource Allocation,” *American Economic Review*, 63(2), 1–30.
- [7] Kalai A., E. Kalai, E. Lehrer, and D. Samet (2010): “A Commitment Folk Theorem,” *Games and Economic Behavior*, 69(1), 127–137.
- [8] Martimort, D. and L. Stole (2002): “The Revelation and Delegation Principles in Common Agency Games,” *Econometrica*, 70(4), 1659–1673.
- [9] McAfee, R.P. (1993): “Mechanism Design by Competing Sellers,” *Econometrica*, 61(6), 1281–1312.
- [10] Milgrom, P.R., and R.J. Weber (1985): “Distributional Strategies for Games with Incomplete Information,” *Mathematics of Operations Research*, 10(4), 619–631.
- [11] Myerson, R.B. (1979): “Incentive Compatibility and the Bargaining Problem,” *Econometrica*, 47(1), 61–73.
- [12] Myerson, R.B. (1982): “Optimal Coordination Mechanisms in Generalized Principal-Agent Problems,” *Journal of Mathematical Economics*, 10(1), 67–81.
- [13] Pavan, A. and G. Calzolari (2009): “Sequential Contracting with Multiple Principals,” *Journal of Economic Theory*, 144(2), 503–531.

- [14] Pavan, A. and G. Calzolari (2010): “Truthful Revelation Mechanisms for Simultaneous Common Agency Games,” *American Economic Journal: Microeconomics*, 2(2), 132–190.
- [15] Peck, J. (1997): “A Note on Competing Mechanisms and the Revelation Principle,” Mimeo, Ohio State University.
- [16] Peters, M. (2001): “Common Agency and the Revelation Principle,” *Econometrica*, 69(5), 1349–1372.
- [17] Peters, M. (2014): “Competing Mechanisms,” *Canadian Journal of Economics*, 47(2), 373–397.
- [18] Peters, M. (2015): “Reciprocal Contracting,” *Journal of Economic Theory*, 158, 102–126.
- [19] Peters, M., and B. Szentes (2012): “Definable and Contractible Contracts,” *Econometrica*, 80(1), 363–411.
- [20] Peters, M., and C. Troncoso-Valverde (2013): “A Folk Theorem for Competing Mechanisms,” *Journal of Economic Theory*, 148(3), 953–973.
- [21] Prat, A., and A. Rustichini (2003): “Games Played Through Agents,” *Econometrica*, 71(4), 989–1026.
- [22] Rao, B.V. (1971): “Borel Structures for Function Spaces,” *Colloquium Mathematicum*, 23(1), 33–38.
- [23] Rahman, D., and I. Obara (2010): “Mediated Partnerships,” *Econometrica*, 78(1), 285–308.
- [24] Szentes, B. (2015): “Contractible Contracts in Common Agency Problems,” *Review of Economic Studies*, 82(1), 391–422.
- [25] Xiong, S. (2013): “A Folk Theorem for Contract Games with Multiple Principals and Agents,” Mimeo, Department of Economics, Rice University.
- [26] Yamashita, T. (2010): “Mechanism Games With Multiple Principals and Three or More Agents,” *Econometrica*, 78(2), 791–801.