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

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

October 22, 2021

Abstract

We study games in which several principals contract with several privately-informed agents. We show that enabling the principals to engage in contractible private disclosures – by sending private signals to the agents about how the mechanisms will respond to the agents' messages – can significantly affect the predictions of such games. Our first result shows that equilibrium outcomes (and payoffs) of games without private disclosures need not be sustainable when private disclosures are allowed. The result thus challenges the robustness of the "folk theorems" of Yamashita (2010) and Peters and Troncoso-Valverde (2013). Our second result shows that private disclosures may generate equilibrium outcomes that cannot be supported in any game without private disclosures, no matter the richness of the message spaces and the availability of public randomizing devices. The result thus 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, Signals, Universal Mechanisms, Folk Theorems.
JEL Classification: D82.

*For comments and useful suggestions, we thank participants at various conferences and seminars where the paper was presented. This research has benefited from financial support of the ANR (Programme d'Investissement 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, and Università degli Studi di Roma "Tor Vergata," Roma, Italy. Email: andrea.attar@tse-fr.eu.

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

[§]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 information in a mechanism then becomes crucial for the decisions that are implemented. In incomplete-information environments in which all payoff-relevant decisions are taken by the principal, private communication can be restricted to be from the agents to the principal (Myerson (1979)). In that case, the principal may, without loss of generality, post a public mechanism that specifies a (possibly random) decision for every profile of messages she may receive from the agents; hereafter, we refer to such communication protocols as *standard mechanisms*. Communication from the principal to the agents takes the form of the announcement of such public mechanism. Any private communication from the principal to the agents is redundant, in that it has no impact on the set of allocations and payoffs that can be implemented.

In this paper, we show that these standard insights from classical mechanism design theory do not extend to settings with competing principals. Specifically, we show that, when several agents contract with several principals, private disclosures from the principals to the agents can significantly affect the set of equilibrium outcomes. We introduce such private disclosures by allowing the principals to send to the agents contractible private signals about the implications of the agents' messages for the selected allocations. Importantly, such private signals are sent before receiving messages from the agents. We derive our results in pure incomplete-information environments in which only the principals take payoff-relevant decisions. In such environments, private disclosures play no role when a single principal contracts with the agents.

The two theoretical pillars of this literature can be described as follows. First, one can construct a space of universal mechanisms whereby every agent can communicate to every principal his endogenous market information – that is, the profile of mechanisms posted by the other principals – in addition to his exogenous type (Epstein and Peters (1999)). An analogue of the revelation principle then 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. Second, one can obtain an explicit characterization of the equilibrium outcomes and of the equilibrium payoffs for a large class of games (Yamashita (2010), Peters and Troncoso-Valverde (2013)). An analogue of the folk theorem then holds: any incentive-feasible allocation that yields every principal a payoff above a well-defined min-max-min bound can be supported in equilibrium.

From our perspective, the key point is that these two central results are derived in a framework in which only the agents privately communicate with the principals, who post standard mechanisms. We challenge these findings by considering a richer class of communication protocols in which the principals are allowed to disclose some information about their mechanisms privately to the agents. While so doing, we maintain two main informational assumptions of this literature. First, principals do not directly communicate among them. Second, principals' mechanisms cannot directly refer to each other, that is, a principal's mechanism cannot condition its allocations directly on the allocations selected in other principals' mechanisms.

Our first result establishes the non-robustness of equilibria in standard mechanisms and the non-validity of the folk-theorems established for such games. We provide an example of equilibrium outcomes of a game in which principals compete in standard mechanisms (with rich message spaces) that cannot be sustained in a game where, in addition to the original mechanisms, the principals can also offer mechanisms with private disclosures– i.e. with signals sent by the principals to the agents before the agents send messages to the principals. In the original game, standard mechanisms are sufficiently rich to include *recommendation* ones, as in Yamashita (2010). In a recommendation mechanism, agents, in addition to reporting their exogenous payoff-relevant information, are asked to "vote" on the direct mechanism the principal should use. Any such game is known to be amenable to folk-theorem-type of results whereby any allocation that is incentive-compatible for the agents (in the usual sense), and yields each principal a payoff above an appropriately-defined min-max-min threshold can be sustained in equilibrium. We show that, by deviating to a mechanism that discloses information about her decisions privately and asymmetrically to the agents, a principal can ensure that the agents no longer have the incentives to carry out the punishments with the non-deviating principals necessary to make the deviation unprofitable. Furthermore, the equilibrium allocations that are not robust to the principals deviating to mechanisms with private disclosure are incentive-compatible and individually rational in the sense of Myerson (1979). The result thus implies that the equilibria characterized by Yamashita (2010) and by Peters and Troncoso-Valverde (2013), as well as the folk theorems established in these papers and in the related literature on contracts-on-contracts (Kalai et al. (2010), Peters and Szentes (2012), Peters (2015), and Szentes (2015)) need not be robust to deviations to mechanisms with private disclosures: certain payoffs above the min-max-min threshold cannot be sustained in games with private disclosures.

Our second result is that equilibrium outcomes of competing-mechanism games in which principals are allowed to post mechanisms with private disclosures need not be equilibrium outcomes in any game in which principals are restricted to standard mechanisms, no matter the richness of the message spaces. The reason is that private disclosures may help the principals correlate their decisions with the information privately held by the agents in a way that cannot be replicated by the principals responding to the agents' messages when the latter are based solely on the agents' common knowledge of the mechanisms and of the agents' exogenous private information. We establish the result by means of an example in which the equilibrium correlation between the principals' decisions and the agents' exogenous private information requires that (a) the agents receive information about a principal's decision and pass it on to another principal *before* the latter principal finalizes her own decision, and (b) such information not create common knowledge among the agents about the former principal's decision before they communicate with the latter principal. The example shows the necessity of both (a) and (b) when it comes to sustaining certain outcomes, the possibility to accomplish both (a) and (b) with private disclosures, and the impossibility to accomplish (a) and (b) with standard mechanisms, regardless of the richness of the message spaces and the availability of public randomizing devices. The result thus also implies that the universal mechanisms of Epstein and Peters (1999) are not canonical when principals can engage in private disclosures.

Taken together, the above results challenge the existing modeling approach to competing mechanisms 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 which principals fully commit to mechanisms in the presence of privately-informed agents. McAfee (1993) is the first to point out that equilibrium outcomes in such games may rely on agents reporting *all* their private information to the principals, i.e. their exogenous payoff-relevant types and the market information contained in the other principals' posted mechanisms. Epstein and Peters (1999) construct a space of universal mechanisms that permits one to establish the analog of the revelation principle for competing-mechanism games. Subsequent work by Yamashita (2010) and Peters and Troncoso-Valverde (2013) provides an explicit characterization of the equilibrium allocations and payoffs of such games and show that the latter coincide with those that are incentive-compatible and individually rational in the sense of Myerson (1979). Our results indicate that such a characterization is sensitive to the assumption that principals are restricted to standard mechanisms and does not extend to settings where principals can engage in private but contractible disclosures.

As it is well known from classical mechanism design theory (Myerson (1982)), private communication from a single principal to the agents is key when certain payoff-relevant actions can be taken solely by the agents, as in moral-hazard settings. Such a communication, which takes the form of action recommendations for the agents, has been shown to serve as a correlating device between the players' decisions in several problems of economic interest, such as the partnership model of Rahman and Obara (2010). Perhaps surprisingly, however, private disclosures have been neglected in competing-mechanism settings *even when agents take private actions*, as in the lobbying model of Prat and Rustichini (2003). To the best of our knowledge, the only exception is the recent work of Attar, Campioni and Piaser (2019), which considers a complete-information game in which agents' actions are observable. They construct an example in which equilibrium allocations sustained by standard mechanisms fail to be robust against unilateral deviations by the principals to mechanisms with private recommendations. In equilibrium, a principal implements a correlation between her decisions and the agents' actions that cannot be sustained without sending private recommendations.

While the insights in the above papers can be interpreted in light of the traditional role that private signals (from the principals to the agents) play in single-principal settings (Myerson (1982, 1986), Forges (1986)), the present paper uncovers a novel role for such signals, which is different from action recommendations. Private disclosures to the agents about the decisions implemented in response to the agents' messages permit the principals to guarantee themselves a payoff strictly above the min-max-min threshold by making it difficult to the agents and the other principals to punish a deviation. They play a role also in settings where the principals may use a mediator to coordinate their offers to the agents and/or the agents may use a mediator to coordinate their responses to the principals.[1]

Private signals may also help overcoming the lack of a direct communication channel among the principals. By disclosing information asymmetrically to the agents, the principals may correlate their decisions with those of other principals and with the agents' exogenous private information in a way that is not replicable with standard mechanisms without

[1]The use of a mediator in competing-mechanism games is somewhat controversial because it is in contrast with the spirit of these games where contracting is decentralized. However, irrespective of whether or not a mediator is available, our results show that private disclosures may protect a principal against possible punishments by the agents and the other principals in a way that is not replicable in their absence, affecting the validity of existing folk theorems. In fact, as our Example 1 shows, the only way a mediator could succeed in inflicting the min-max-min payoff to a principal deviating to a mechanism with private disclosures is by forcing the agents to send messages to the non-deviating principals before receiving any information from the deviating principal. This appears both unrealistic and artificial, as it requires confidence on an ad-hoc timing of communication between the agents and the principals.

private disclosures (Example 2). Such a possibility appears particularly relevant for the applications of competing-mechanism games favored in the literature, where contracting is typically decentralized (that is, the principals find it difficult to rely on a common mediator to coordinate their decisions) and direct communication among the principals (either in the form of the principals exchanging cheap-talk messages, or in the form of the 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 unlike to be feasible.

In our setting, the correlation in the principals' decisions is generated by the agents communicating to the principals both their exogenous types and the information privately received from other principals. The principals cannot directly condition on other principals' mechanisms, or exchange information directly among them. Instead, Kalai et al. (2010), Peters and Szentes (2012), Peters (2015), and Szentes (2015), consider settings in which players make commitments that are contingent on the commitments of others and communication is unrestricted. The result in our Example 1 extends to these settings: by deviating to a mechanism with private disclosures a principal can guarantee herself a payoff strictly above the min-max-min threshold, irrespective of whether or not principals can make commitments contingent on other principals' mechanisms and decisions. The result in Example 2, instead, extends to settings in which principals can condition their mechanism on other principals' mechanisms (the case considered in Peters (2015)) but not to settings in which principals can condition their final payoff-relevant decision directly on other principals' payoff-relevant decisions. Intuitively, in such settings, disclosing information privately to the agents as a way of "encrypting" it and passing it on to opponent principals while respecting the agents' incentives is not useful given that the principals can communicate directly among themselves.

The role of signals we document hinges on the presence of at least two principals and two agents. With a single principal, the Revelation Principle obviously holds and private disclosures have no bearing on the set of equilibrium outcomes and payoffs. Similarly, when multiple principals contract with a single agent, the menu theorems of Martimort and Stole (2002), Peters (2001), and Pavan and Calzolari (2009, 2010) guarantee that any equilibrium allocation of a game in which the principals compete by offering arbitrary message-contingent allocations can be reproduced in the game in which the principals offer subset (menus) of their decisions to the agent and delegate to the latter the choice of the final allocations. In such settings, private disclosures play no role.

The rest of 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 $j = 1, \ldots, J$, contract with several agents, indexed by $i = 1, \ldots, I$. As anticipated above, our results hinge on $J, I \geq 2$.

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 $\omega \equiv (\omega^1, ..., \omega^I)$ is $\Omega \equiv \Omega^1 \times \ldots \times \Omega^I$. 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 $v_j : X \times \Omega \to \mathbb{R}$ and $u^i : X \times \Omega \to \mathbb{R}$ be the payoff functions of principal j and of agent i, respectively, where $X \equiv X_1 \times \ldots \times X_J$ is the set of possible profiles of decisions for the principals. Agents take no payoff-relevant decisions. An *allocation* is a function $z : \Omega \to \Delta(X)$ 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 *mechanism with signals* for a principal 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 the principal's decisions to every profile of signals sent to the agents and every profile of messages received 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 $s_j \equiv (s_j^1, \ldots, s_j^I)$ that principal j sends to the agents, where $S_j \equiv S_j^1 \times \ldots \times S_j^I$ for some collection of Polish spaces S_j^i of signals from principal j to every agent i.

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

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

We assume that $\Omega^i \subset M^i_j$ 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^i_j 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 $\phi_j : M_j \to \Delta(X_j)$ 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.

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 on 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 $\phi_j : S_j \times M_j \to \Delta(X_j)$ 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} unfolds in four stages:

1. the principals simultaneously post mechanisms, observed by all agents;

2. the principals' mechanisms simultaneously and privately send 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 $\lambda^i : \Gamma \times S^i \times \Omega^i \to \Delta(M^i)$ that assigns to every profile of mechanisms $\gamma \equiv (\gamma_1, ..., \gamma_J) \in \Gamma \equiv \Gamma_1 \times \ldots \times \Gamma_J$ that the principals may post, every profile of signals $s^i \equiv (s^i_1, ..., s^i_J) \in S^i \equiv S^i_1 \times \ldots \times S^i_J$ 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 $m^i \equiv (m^i_1, ..., m^i_J) \in M^i \equiv M^i_1 \times \ldots \times M^i_J$ sent by agent i, where $\Gamma \times S^i \times M^i$ is endowed with the appropriate product σ-field. The allocation $z_{\mu,\lambda} : \Omega \to \Delta(X)$ induced by the strategies $(\mu, \lambda) \equiv (\mu_1, \ldots, \mu_J, \lambda^1, \ldots, \lambda^I)$ is then defined by

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

for all $(\omega, x) \in \Omega \times X$, where $S \equiv S_1 \times \ldots \times S_J$ and $M \equiv M_1 \times \ldots \times M_J$. For every profile of mechanisms γ, a behavioral strategy for agent i in the subgame γ played by the agents is a Borel-measurable function $\beta^i : S^i \times \Omega^i \to \Delta(M^i)$ 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 $\beta \equiv (\beta^1, \ldots, \beta^I)$ 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 special case of the game G^{SM} arises when S^i_j is a singleton for all i and j, so that the principals can only post standard mechanisms. To distinguish this situation, we denote by G^M the corresponding competing-mechanism game *without signals*; the games studied by Epstein and Peters (1999) and Yamashita (2010) are prominent examples.

Equilibrium The strategy profile (μ, λ) is a perfect Bayesian equilibrium (PBE) of G^{SM} whenever

1. for each $\gamma \in \Gamma$, $(\lambda^1(\gamma), \ldots, \lambda^I(\gamma))$ 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.

Notice that, in any subgame γ, the beliefs of the agents are pinned down by the prior $\mathbf{P}$ and the signal distributions $(\sigma_1, \ldots, \sigma_J)$ to which the principals are committed through the mechanisms γ. An allocation z is *incentive-compatible* if, for all i and $\omega^i \in \Omega^i$,

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

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 Subsection 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 $d_j : \Omega \to X_j$ 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^i_j 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, $D_j \times \Omega^i \subset M^i_j$ for all i and j. Accordingly, a *recommendation mechanism* ϕ^r_j for principal j stipulates that, if every agent i sends a message $m^i_j \equiv (d^i_j, \omega^i) \in D_j \times \Omega^i$ to principal j, then

$$\phi^r_j(m^1_j, \ldots, m^I_j) \equiv \begin{cases} d_j(\omega^1, \ldots, \omega^I) & \text{if card}\{i : d^i_j = d_j\} \geq I - 1 \\ \overline{x}_j & \text{otherwise} \end{cases}, \quad (1)$$

where $\overline{x}_j$ is some fixed decision in X_j; if, instead, some agent i sends a message $m^i_j \notin D_j \times \Omega^i$

to principal j, then ϕ_j^r treats this message as if it coincided with some fixed element $(\overline{d}_j, \overline{\omega}_j^i)$ 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 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).

[3] As 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.

	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)

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 min-max-min payoff that can be inflicted on P2 by using recommendation mechanisms, and shows how this payoff can be attained in an equilibrium of G_1^M.[4]

Claim 1 *The outcome*

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

in which P2 obtains her minimum feasible payoff of 5, can be supported in a PBE of G_1^M*.*

Proof. 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 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. We assume that $\overline{\omega}_1^1 = \overline{\omega}_1^2 = \omega_L$, so that, if some agent $i = 1, 2$ sends a message $m_1^i \notin D_1 \times \Omega^i$ to P1, P1 treats this message as if agent i reported to her to be of type ω_L. We claim that, in the corresponding 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

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

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

[4] For 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 $\phi_2 : M_2 \to \Delta(X_2)$, 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

$$\overline{p} \equiv \max_{m_2 \in M_2} p(m_2) \tag{4}$$

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

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

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 $\overline{m}_2^3$, achieves the minimum in (5). That $\overline{p}$, $\overline{m}_2$, $\underline{p}$, 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 Suppose first that ϕ_2 is such that $\overline{p} \geq \frac{1}{2}$. 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 $\overline{m}_2^3$ 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 $\overline{m}_2^3$ 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 $\overline{p} \geq \frac{1}{2}$, $8\overline{p} + (1 - \overline{p}) \geq 4.5$. From Table 1, and by definition of d_1^* and $\overline{m}_2$, it thus follows that, if A2 reports ω_L to P1 and sends $\overline{m}_2^2$ to P2, and if A3 sends $\overline{m}_2^3$ to P2, then A1 best responds by reporting ω_L to P1 and sending $\overline{m}_2^1$ to P2; notice, in particular, that, because $\overline{\omega}_1^1 = \omega_L$, 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 Suppose next that ϕ_2 is such that $\overline{p} < \frac{1}{2}$. 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

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

to P1; (ii) A1 and A2 truthfully report their types to P1; (iii) A3 sends message $\overline{m}_2^3$ 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 $\underline{p} \leq \overline{p} < \frac{1}{2}$, $\underline{p} + 8(1 - \underline{p}) > 4.5$. 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 $\overline{m}_2^3$ 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 $\overline{\omega}_1^1 = \omega_L$, 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. Hence the result. ■

The proof of Claim 1 relies 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 can be used to support random allocations, however; see, for instance, Xiong (2013). As the following folk theorem shows, this enables one to support many equilibrium payoffs for P2.

Claim 2 *Any payoff for P2 in* $[5, 6]$ *can be supported in a PBE of* G_1^M.

The proof of Claim 2 is provided in the Appendix. The proof only requires adjusting the principals' behavior on path—off path, letting the agents coordinate on the mechanism

d_1 used in the proof of Claim 1 is sufficient to deter P2's deviation. To this end, we modify Yamashita's (2010) definition of a recommendation mechanism to allow principals to randomize over their decisions on path.

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 minimal 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 the equilibrium outcomes characterized in Claims 1–2 are not robust to private disclosures. Specifically, we show that the outcome (2) cannot be supported in any equilibrium of any enlarged game in which principals can post mechanisms with signals, and, more generally, that, in any such game, 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.

Claim 3 *G_1^{SM} admits a PBE. Moreover, if $\operatorname{card} S_2^1 \geq 2$, 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.

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

[6]As the arguments below reveal, the second part of Claim 3 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 $B^* : \Gamma_1 \times \Gamma_2 \twoheadrightarrow \prod_{i=1}^3 \Delta(M^i)^{S^i \times \Omega^i}$ is upper hemicontinuous (Milgrom and Weber (1985, Theorem 2)) and thus admits a measurable selection β^* 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 $\lambda^{i*}(m^i \,|\, \gamma_1, \gamma_2, s^i, \omega^i) \equiv \beta^{i*}(\gamma_1, \gamma_2)(m^i \,|\, s^i, \omega^i)$. 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 a payoff of 5.5 regardless of the mechanism $\gamma_2^* \in \Gamma_2$ 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}. This shows the first part of Claim 3.

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,

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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 $m_1^1 \in M_1^1$,

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

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

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

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 $\sigma > \frac{1}{2}$, a payoff at least equal to

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

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 a payoff at most equal to

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

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

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

This shows the second part of Claim 3. Hence the result. ■

The proof of Claim 3 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 value 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 expected 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.

[7] Otherwise, 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.

Claim 3 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 Claim 3 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.

Claims 1–3 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 payoffs 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 in private disclosures.*

Example 1 above shows 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 guarantee 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 the proof of Claim 1, which would effectively neutralize P2's signals.

4 Non-universality of Standard Mechanisms

Example 1 in the previous section shows that equilibria in games in which principals are restricted to standard mechanisms (with a rich message space) need not be robust to the possibility that principals deviate and offer mechanisms with signals.

In this section, we address the dual question of whether competing-mechanism games in which principals can engage in private disclosures may admit equilibria whose outcomes are not sustainable in games in which principals are restricted to standard mechanisms, no matter the richness of the message spaces. We provide an example showing that this is indeed the case. The example shows that, with private signals, a principal can make the agents' messages to other principals 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 sustained by the principals randomizing over their mechanisms and/or by the agents randomizing over the messages they send to the principals.

Example 2 Let $I = J \equiv 2$. Again, 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 has two types, with $\Omega^2 \equiv \{\omega_L, \omega_H\}$, while A1's type space is a singleton and hence omitted to ease the notation. The states ω_L and ω_H are commonly believed to occur with probabilities $\mathbf{P}[\omega_L] = 1/4$ and $\mathbf{P}[\omega_H] = 3/4$, 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. The variable $\zeta < 0$ denotes an arbitrary loss for P2. P1's payoff is constant over X and hence omitted.

4.1 An Equilibrium with Private Disclosures

To illustrate the key ideas in the simplest possible manner, we consider a specific game with signals 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

	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.

possible message spaces that allow the agents to report their information to the principals. Specifically, we assume that $M_1^i = \Omega^i \times S_2^i$ and $M_2^i = \Omega^i$.[8] We refer to this game as G_2^{SM}. The following result then holds.

Claim 4 *For $\alpha = \frac{2}{3}$, the following is a PBE outcome of G_2^{SM}:*

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

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

P2's payoff in any equilibrium inducing the above outcome is equal to 10.

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

$$\sigma_2^*(s_2) \equiv \begin{cases} \frac{\alpha}{2} & \text{if } s_2 = (1,1), \\ \frac{\alpha}{2} & \text{if } s_2 = (2,2), \\ \frac{1-\alpha}{2} & \text{if } s_2 = (1,2), \\ \frac{1-\alpha}{2} & \text{if } s_2 = (2,1), \end{cases} \tag{20}$$

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

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

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 received from P2, every agent's posterior distribution

[8]As the arguments below reveal, result in Claim 4 does not hinge on these simplifying assumptions. In particular, the result also extends to games with richer signal and message provided that $M_1^i \supset \Omega^i \times S_2^i$ and $M_2^i \supset \Omega^i$.

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, every 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.[9]

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

$$\phi_1^*(\emptyset, m_1) \equiv \begin{cases} \delta_{x_{13}} & \text{if } m_1 \in \{(1,\omega_L,1),(2,\omega_L,2)\}, \\ \delta_{x_{14}} & \text{if } m_1 \in \{(1,\omega_L,2),(2,\omega_L,1)\}, \\ \delta_{x_{12}} & \text{if } m_1 \in \{(1,\omega_H,1),(2,\omega_H,2)\}, \\ \delta_{x_{11}} & \text{if } m_1 \in \{(1,\omega_H,2),(2,\omega_H,1)\}, \end{cases} \tag{22}$$

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 the signal $s_2^1=1$ from P2, whereas A2 reports that her type is ω_L and she received the signal $s_2^2=1$ from P2. Observe from (21)–(22) 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 received from P2. We now show that these behaviors constitute a BNE in 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, the only incentives we have to consider concern his report to P1 of 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

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

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

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

which is strictly less than his payoff from reporting the true signal received from P2, for all $\alpha \in [0,1]$.

[9] $\delta_\emptyset$ is the Dirac measure assigning probability 1 to $s_1=\emptyset$.

Step 2. Consider next A2's incentives, under the belief that A1 is truthful to P1. We first consider the behavior of A2 when he is of type ω_L. If he truthfully reports both his type and his signal to P1, then, regardless of the signal he receives from P2, his expected payoff is equal to

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

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 equal to

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

which is smaller or equal to his payoff from reporting both ω^2 and s_2^2 truthfully to P1 if $\alpha \leq \frac{8}{9}$.

If A2 misreports his type to P1 but reports the signal received from P2 truthfully, his expected payoff is equal to

$$\alpha u^2(x_{12}, x_{21}, \omega_L) + (1-\alpha)u^2(x_{11}, x_{22}, \omega_L) = 5\alpha + 3.5(1-\alpha), \tag{24}$$

which is no greater than his payoff from reporting both ω^2 and s_2^2 truthfully if $\alpha \leq \frac{2}{3}$.

Finally, if A2 misreports both his type and the signal received from P2, his expected payoff is equal to

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

which is smaller or equal to his payoff from reporting both ω^2 and s_2^2 truthfully if $\alpha \geq \frac{1}{5}$.

We next consider the behavior of A2 when he is of type ω_H. If he truthfully reports both his type and his signal to P1, then, regardless of the signal he receives from P2, his expected payoff is equal to

$$\alpha u^2(x_{12}, x_{21}, \omega_H) + (1-\alpha)u^2(x_{11}, x_{22}, \omega_H) = 9\alpha + 5(1-\alpha). \tag{25}$$

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 equal to

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

which is smaller or equal to his payoff from reporting both ω^2 and s_2^2 truthfully if $\alpha \geq \frac{1}{4}$.

If A2 misreports his type to P1 but reports the signal received from P2 truthfully, his expected payoff is equal to

$$\alpha u^2(x_{13}, x_{21}, \omega_H) + (1-\alpha) u^2(x_{14}, x_{22}, \omega_H) = 7\alpha + 9(1-\alpha), \tag{26}$$

which is smaller or equal to his payoff from reporting both ω^2 and s_2^2 truthfully if $\alpha \geq \frac{2}{3}$.

Finally, if A2 misreports both his type and the signal received from P2, his expected payoff is equal to

$$\alpha u^2(x_{14}, x_{21}, \omega_H) + (1-\alpha) u^2(x_{13}, x_{22}, \omega_H) = 6\alpha + 7(1-\alpha), \tag{27}$$

which is smaller or equal to his payoff from reporting both ω^2 and s_2^2 truthfully if $\alpha \geq \frac{2}{5}$.

The analysis above implies that it is a BNE for A1 and A2 to report truthfully 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 payoff of 10. Because P1's payoff is constant, there exists a PBE of G_2^{SM} in which P1 and P2 post the mechanisms γ_1^* and γ_2^* on the equilibrium path, 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 the subgames are finite. Hence the result. ■

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

The above construction relies on the fact that, although the mechanism with signals γ_2^* is publicly disclosed to both agents, A1 and A2 receive different signals from P2 (private disclosures). In the example, such signals are uninformative of P2's decision. If P2 were to inform the agents of her decision, then after learning that P2 takes decision x_{21}, A2 of type ω_L would not be willing to induce the decision x_{13} with P1.

The construction also reveals that, for P2 to obtain her maximal payoff of 10 while respecting 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 implement the equilibrium correlation between P1's and P2's decisions by letting the agents randomize over their messages to the principals and committing to respond deterministically to the received messages, such a delegation is not incentive compatible. It is thus essential that the randomization be carried out by the principals themselves. The correlation in 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 with private disclosures. The analysis in the next subsection (as well as the discussion in Section 5) confirms the above intuition by establishing the indispensability of the signals, no matter the richness of the message spaces.

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 without signals in which principals are constrained to post standard mechanisms, irrespective of the richness of message spaces. That is, private signals are indispensable to support the outcome in (18)–(19). To this end, we consider a general competing-mechanism game without signals in which every principal j can only post a standard mechanism $\phi_j : M_j \to \Delta(X_j)$. This general formulation notably allows us to capture the case where every principal j's message spaces M_j^1 and M_j^2 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). We generically refer to such a game as G_2^M.

Claim 5 *There exists no PBE of G_2^M that supports the outcome (18)–(19) for $\alpha = \frac{2}{3}$.*

Proof. The proof more generally shows 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 continuation equilibrium $\lambda = (\lambda^1, \lambda^2)$ that support the outcome (18)–(19), irrespective of the value of $\alpha \in [0, 1]$. Note that, in establishing the result, we do not impose that the joint distribution μ be equal to the product of its marginals. In other words, we allow for correlation devices that the principals may use to coordinate their choice of a mechanism. 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

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

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 expected 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 $\alpha^\phi_L \le \alpha^\phi_H$. To see this, notice that, as A1 does not know which state prevails, given A1's state-independent behavior, the state-dependent outcomes $z^\phi(\omega_L)$ and $z^\phi(\omega_H)$ must be the 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

$$3\alpha^\phi_L + 7.5(1-\alpha^\phi_L) \ge 5\alpha^\phi_H + 3.5(1-\alpha^\phi_H). \tag{28}$$

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

$$9\alpha^\phi_H + 5(1-\alpha^\phi_H) \ge 7\alpha^\phi_L + 9(1-\alpha^\phi_L). \tag{29}$$

The two inequalities are satisfied only if $\alpha^\phi_L \le \alpha^\phi_H$. The claim follows. Because

$$\int \alpha^\phi_L\,\mu(\mathrm{d}\phi) = \alpha = \int \alpha^\phi_H\,\mu(\mathrm{d}\phi), \tag{30}$$

we obtain that, with μ-probability 1, $\alpha^\phi_L = \alpha^\phi_H$. Substituting $\alpha^\phi_L = \alpha^\phi_H$ into (28)–(29), it follows that, with μ-probability 1, $\alpha^\phi_L = \alpha^\phi_H = \frac{2}{3}$ and, hence, by (30), that $\alpha = \frac{2}{3}$. We thsu 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 $(m^2_1, \hat{m}^2_2)$, then, given A1's behavior strategy $\lambda^1(\phi)$, the mechanisms (ϕ_1, ϕ_2) would induce the following marginal distributions over $(x_{11}, x_{12}, x_{21}, x_{22})$:

$$\Pr(x_{11}, x_{21}) + \Pr(x_{11}, x_{22}) = \frac{1}{3}, \tag{31}$$

$$\Pr(x_{12}, x_{21}) + \Pr(x_{12}, x_{22}) = \frac{2}{3}, \tag{32}$$

$$\Pr(x_{11}, x_{21}) + \Pr(x_{12}, x_{21}) = \frac{2}{3}, \tag{33}$$

$$\Pr(x_{11}, x_{22}) + \Pr(x_{12}, x_{22}) = \frac{1}{3}. \tag{34}$$

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 $(m^2_1, \hat{m}^2_2)$ 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 the latter must at most be his equilibrium payoff of 4.5, it follows that $p = 0$. This implies that, for $\lambda^1(\phi)$-almost every m^1, we have that

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

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

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

and thus, as decisions are perfectly correlated,

$$(\phi_1(m_1^1, m_1^2), \phi_2(m_2^1, m_2^2)) = (\phi_1(m_1^1, \hat{m}_1^2), \phi_2(m_2^1, \hat{m}_2^2)). \tag{35}$$

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 that (35) 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, (35) holds for $\lambda^2(\phi \mid \omega_H) \otimes \lambda^2(\phi \mid \omega_H)$-almost every $(m_2, \hat{m}_2)$, so that the mapping $(m_1^2, m_2^2) \mapsto (\phi_1(m_1^1, m_1^2), \phi_2(m_2^1, m_2^2))$ is constant over a set of $\lambda^2(\phi)$-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. Hence the result. ■

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 offered by the principals must respond deterministically to the messages sent by the agents (Step 1 in the proof of Claim 5). Second, because only A2 observes the state, and because the distribution over the principals' decisions varies with the state, A2 must weakly prefer the distribution over the messages he is supposed to send in each state to the one he is supposed to carry out in the other state. This constraints the joint distributions that can be sustained in the two states (Step 2 in the proof of Claim 5). Third, for A2 to prefer the distribution over the principals' decisions he is supposed to induce in state ω_L to the one that he can induce by "mis-matching" the principals' decisions while preserving the marginal distributions he induces in state ω_H, it must be that, given the mechanisms offered in equilibrium, the messages he sends in state ω_H are not influential when combined with those sent with positive probability by A1 (Step 3 in the proof of Claim 5). But then A1 has a profitable deviation (Step 4 in the proof of Claim 5).

It should be noted that the result in Claim 5 holds no matter the richness of the message spaces. Hence, it also applies to the Epstein and Peters (1999) class of universal mechanisms. Specifically, because these mechanisms are standard ones, the two claims above jointly imply the following result:

Proposition 2 *There exist allocations that can be supported in a PBE of a competing-mechanism game in which principals offer mechanisms with private disclosures but that cannot be supported in any PBE of any competing-mechanism game in which principals are restricted to offering standard mechanisms (including universal mechanisms)—and this, even if the principals or the agents mix in equilibrium.*[10]

The result suggests that the universal mechanisms of Epstein and Peters (1999) fail to be canonical when principals can engage in 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 other principals and with the agents' exogenous private information.

Remark. The proof of Claim 5 does not make use of the property that the principals' randomizations are independent. The result in Claim 5 thus carries over to the case where G_2^M is augmented by arbitrarily rich public randomizing devices that the principals may use to correlate their choice of the 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 they 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 one another 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

In this section, we discuss the different roles played by private disclosures in the two examples. We also comment on how the results due to private disclosures may extend to competing

[10]Epstein and Peters (1999, Theorem 4.1) restrict attention to PBEs in which principals play pure strategies.

mechanism games in which a principal can directly condition her mechanism on her opponents' ones.

First, consider Example 1. In that example, the possibility for P2 to guarantee herself a payoff strictly above her min-max-min threshold hinges on informing the agents asymmetrically about her decisions. Indeed, the main thrust of private signals in Example 1 is the destabilizing role they play *out of equilibrium.* As anticipated above, such a destabilizing role is relevant no matter whether the principals have access to private randomizing devices (or a mediator), and of whether principals can condition their decisions directly on other principals' decisions, as in the literature on contracts-on-contracts.

In Section 3, we argued that, if P2 were to perfectly inform *all* agents of her decisions, or, more generally, of the allocations selected in response to the agents' messages (as in standard mechanisms), then it would be possible for P1 to offer a mechanism that inflicts to P2 her min-max-min payoff. Below we show that the same conclusion is true under any signal structure that leaves the agents' in the dark. Formally, we show that the analogue of Claim 3 is false if P2 is restricted to offering mechanisms in which private signals take the form of uninformative encrypted keys, as in the proof of Claim 4 in Example 2. To see this, consider the game G_1^{SM} studied in Claim 3; moreover, as in Claim 1, assume that $M_j^i \supset D_j \times \Omega^i$ 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 (a) σ_2 is a product measure over S_2, and (b) for all i, $s_2^i \in S_2^i$, $m_2 \in M_2$, and $x_2 \in X_2$,

$$\sum_{s_2^{-i} \in S_2^{-i}} \sigma_2^{-i}(s_2^{-i})\phi_2(x_2 \,|\, s_2^i, s_2^{-i}, m_2) = \sum_{s_2 \in S_2} \sigma_2(s_2)\phi_2(x_2 \,|\, s_2, m_2), \tag{36}$$

where σ_2^{-i} is the marginal of σ_2 over S_2^{-i}.[11]

The first condition states that the private signal P2 sends to any agent is uninformative about the private signals she sends to the other agents. The second condition states that, given any profile of messages the agents may send to P2, the private signal P2 sends to any agent is uninformative about P2's decision. The following result then holds:

Claim 6 *In Example 1, if P1 posts a recommendation mechanism ϕ_1^r, then, for every mechanism γ_2 of P2 whose signals are uninformative, there exists a BNE of the subgame (ϕ_1^r, γ_2) that yields P2 a payoff of 5.*

[11] That S_2 is a finite set plays no role for the conclusion. In particular, Claim 6 remains valid if S_2 is an arbitrary Polish space.

Proof of Claim 6. Because A3's payoff is constant 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 $\omega^2 \in \Omega^2$, $\beta^2(\cdot \mid s_2^2, \omega^2)$ is independent of s_2^2, and $\beta^3(\cdot \,|\, s_2^3)$ 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 there exists 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. According to (36), 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 $\hat{\phi}_2$ such that, for all $m_2 \in M_2$ and $x_2 \in X_2$,

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

Notice that, by construction, the same outcome is implemented in each 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 outcome. To conclude, observe that, as $\hat{\phi}_2$ is a standard mechanism, we know from Claim 1 that there exists a BNE of the subgame $(\phi_1^r, \hat{\phi}_2)$ that yields P2 a payoff of 5. Hence the result. ■

This last result reflects that, if P1 posts a recommendation mechanism ϕ_1^r and P2 posts a mechanism γ_2 with uninformative signals, there exists a one-to-one correspondence between the babbling BNEs of the subgame (ϕ_1^r, γ_2) and the BNEs 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. Because 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 Claim 4, which showed the power of mechanisms with uninformative signals in the context of Example 2; indeed, the key step in the proof of Claim 4 was precisely to construct a non-babbling BNE of the agents' subgame in which they truthfully report to P1 the uninformative signals they receive from P2. Yet Claim 6 suggests that, from P2's perspective, a potential drawback of mechanisms with uninformative signals is that they accommodate for babbling equilibria that may keep her payoff down to 5. This contrasts with the mechanism for P2 constructed in Claim 3, in which disclosing her decision to the agents in an asymmetric way allows P2 to guarantee herself a payoff strictly above her min-max-min bound of 5, regardless of the mechanism posted by P1 and of the agents' continuation equilibrium.

The result in Proposition 1 is established by assuming that the principals cannot condition their mechanisms and decisions directly on other principals' mechanisms and decisions. However, the result extends to settings in which such conditioning is feasible, as in the literature on contracts-on-contracts and reciprocal mechanisms (e.g., Kalai et al. (2010), Peters and Szentes (2012), Peters (2015), and Szentes (2015)). In fact, in Example 1, the only way to inflict a payoff of 5 to P2 is by implementing x_{11} in state (ω_L, ω_L) and x_{12} in state (ω_H, ω_H) with certainty. Because the state is observed privately by the agents, P1 must let the agents determine which decisions to implement in response to a deviation by P2. But then arguments similar to those establishing Claim 3 imply that, when the agents are informed asymmetrically of P2's decisions, it is impossible to induce them to carry out the punishment necessary to block P2's deviations, irrespective of whether or not contracts-on-contracts are feasible.

Next, consider Example 2. In that example private signals are used *on the path* to induce a certain correlation between the principals' decisions and the agents' private information that cannot be sustained with standard mechanisms. In that example private signals do not modify the agents' beliefs. They work as pure "encryption keys" that, in isolation, are completely uninformative about a principal's decision, but, when combined with the keys given to other agents, perfectly reveal a principal's decision.

As discussed above, the impossibility to sustain the outcome in (18)–(19) in Example 2 extends to settings in which the principals and the agents have access to rich public randomizing devices (aka "sunspots") whose realizations are observed by all players prior to committing to their actions. In contrast, the signals of Example 2 could be dispensed with if the principals had access to correlation devices that are either realized after the agents have sent their messages, or that are observed by the principals but not by the agents.

The arguments in the proof of Claim 5 also make clear that its result is not driven by the impossibility for the principals to observe deviations by the other principals. In fact, the result extends to settings in which principals can post "reciprocal mechanisms" as defined in Peters (2015). On the other hand, the result does hinge on the impossibility for the principals to condition their final payoff-relevant decisions directly on other principals' decisions. Clearly, in the context of Example 2 above, if this was possible, the outcome (18)–(19) could be sustained even without private disclosures.

The value of the example is in showing that, absent private correlation devices and the possibility to write contracts-on-contracts, the principals can still correlate their decisions with the agents' private information in a rich way by disclosing information about their decisions privately to the agents and let them pass the information on to other principals.

Taken together, the two examples illustrate that signals in competing-mechanism games play a role fundamentally different from the one they play in single-principal settings. In the latter, signals are used to correlate the agents' behavior, when the agents have primitive actions that cannot be taken by the principals directly (as in moral-hazard settings). In competing-mechanism games, instead, signals are used to inform the agents asymmetrically of the decisions taken by the principals in response to the messages received from the agents. In other words, they inform the agents asymmetrically of the functioning of a mechanism.

6 Conclusions

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 it (a) permits one to sustain all equilibrium allocations of competing-mechanism games with arbitrarily rich message and signal spaces, and (b) 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.

7 Appendix

Proof of Claim 2. 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 $\tilde{d}_j : \Omega \to \Delta(X_j)$; 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 $(\overline{d}_j, \overline{\omega}_j^i)$ of $D_j \times \Omega^i$, for some $\overline{d}_j \neq d_j^0$.

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 the proof of Claim 1 and P1 posts an extended recommendation mechanism $\tilde{\phi}_1^r$. 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 (A.2), 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. Fix some $\xi \in [\frac{1}{2}, 1]$. If at least two agents recommend d_1^0 to P1, then:

$$\tilde{d}_1(\omega) \equiv \begin{cases} \tilde{x}_1^{\xi} \equiv \xi\delta_{x_{11}} + (1-\xi)\delta_{x_{12}} & \text{if at least one agent reports } \omega_L \\ \tilde{x}_1^{1-\xi} \equiv (1-\xi)\delta_{x_{11}} + \xi\delta_{x_{12}} & \text{if both agents report } \omega_H \end{cases} \tag{37}$$

Finally, we assume that $\overline{\omega}_1^1 = \overline{\omega}_1^2 = \omega_L$, 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 1 to be of type ω_L; recall from the proof of Claim 1 that ϕ_2^r similarly satisfies $\overline{\omega}_2^1 = \overline{\omega}_2^2 = \omega_L$.

We now show that, for each $\xi \in [\frac{1}{2}, 1]$, the subgame $(\tilde{\phi}_1^r, \phi_2^r)$ has a BNE in which: (a) each agent recommends to P1 the direct mechanism d_1^0, and recommends to P2 the direct mechanism d_2^* defined by (3); (b) 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 $\overline{\omega}_1^1 = \overline{\omega}_2^1 = \omega_L$, 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 $\tilde{x}_1^{\xi}$, 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 $\tilde{x}_1^{\xi}$, P2 takes decision x_{22}, and A1 obtains a payoff of $\xi + 4.5(1-\xi) < 8\xi + 4.5(1-\xi)$. If A1 misreports his type to P1, irrespective of the report to P2, P1's mechanism implements the lottery $\tilde{x}_1^{\xi}$ which leads to no profitable deviation.

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. Because A1 is not pivotal, recommending a different direct mechanism to either principal is of no avail to him, moreover, because $\overline{\omega}_1^1 = \omega_L$, sending a message $m_1^1 \notin D_1 \times \Omega^1$ to P1 amounts for A1 to lie on 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 $\tilde{x}_1^{1-\xi}$, P2 takes decision x_{22}, and A1 obtains a payoff of $4.5(1-\xi) + 8\xi$. (2) If A1 misreports his type to P1, irrespective of the report to P2, then P1 implements the lottery $\tilde{x}_1^{\xi}$ and P2 takes decision x_{22}, and A1 obtains a payoff of $4.5\xi + 8(1-\xi) < 8\xi + 4.5(1-\xi)$ as $\xi \in [\frac{1}{2}, 1]$.

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. Fix some $\xi \in [\frac{1}{2}, 1]$. If at least two agents recommend d_1^0 to P1, then:

$$\tilde{d}_1(\omega) \equiv \begin{cases} \tilde{x}_1^{\xi} \equiv \xi\delta_{x_{11}} + (1-\xi)\delta_{x_{12}} & \text{if both agents report } \omega_H \\ \tilde{x}_1^{1-\xi} \equiv (1-\xi)\delta_{x_{11}} + \xi\delta_{x_{12}} & \text{if at least one agent reports } \omega_L \end{cases} \tag{38}$$

Finally, we assume that $\overline{\omega}_1^1 = \overline{\omega}_1^2 = \omega_L$, 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 1 to be of type ω_H; the corresponding property for ϕ_2^r is irrelevant for the following arguments.

We now show that, for each $\xi \in (\frac{1}{2}, 1]$, the subgame $(\tilde{\phi}_1^r, \phi_2^r)$ has a BNE in which: (a) 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_{12} 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 $\overline{\omega}_1^1 = \omega_L$, sending a message $m_1^1 \notin D_1 \times \Omega^1$ to P1 amounts for A1 to lie on 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 $\tilde{x}_1^{\xi}$ 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 $\phi_2 : M_2 \to \Delta(X_2)$ can be punished as in the proof of Claim 1, yielding her the minimum feasible payoff of 5, so that she has no profitable deviation either. Hence the result. ■

References

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

[2] Attar, A., E. Campioni, 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] Forges, F. (1986): "An Approach to Communication Equilibria," *Econometrica*, 54(6), 1375–1385.

[7] Hurwicz, L. (1973): "The Design of Mechanisms for Resource Allocation," *American Economic Review*, 63(2), 1–30.

[8] Kalai, A., Kalai, T., Lehrer, E., and D. Samet (2010): "A commitment folk theorem," *Games and Economic Behavior.*, 69(1), 127–137.

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

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

[11] Milgrom, P.R., and R.J. Weber (1985): "Distributional Strategies for Games with Incomplete Information," *Mathematics of Operations Research*, 10(4), 619–631.

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

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

[14] Myerson, R.B. (1986): "Multistage Games with Communication," *Econometrica*, 54(2), 323–358.

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

[16] Pavan, A. and G. Calzolari (2009): "Sequential Contracting with Multiple Principals," *Journal of Economic Theory*, 144(2), 503–531.

[17] Pavan, A. and G. Calzolari (2010): "Truthful Revelation Mechanisms for Simultaneous Common Agency Games," *American Economic Journal: Microeconomics*, 2(2), 132–190.

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

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

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

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

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

[23] Prat, A., and A. Rustichini (2003): "Games Played Through Agents," *Econometrica*, 71(4), 989-1026.

[24] Rahman, D., and I. Obara (2010): "Mediated Partnerships," *Econometrica*, 78(1), 285–308.

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

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

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