

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

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

December 4, 2022

Abstract

We study the design of market information in games in which several principals contract with several privately informed agents. We investigate a new dimension of these games, namely, the possibility for the principals to asymmetrically inform the agents about how their mechanisms respond to their messages. We document two effects of such private disclosures. First, they raise the principals' individual payoff guarantees, protecting them against their competitors' threats. Second, by enlarging the set of incentive-compatible correlation patterns between the principals' decisions and the agents' types, they can be used to support equilibrium outcomes and payoffs that cannot be supported in their absence, no matter how rich the message spaces are allowed to be. These results challenge the folk theorems à la Yamashita (2010) and the canonicity of the universal mechanisms of Epstein and Peters (1999), calling for a novel approach to the analysis of competition in mechanism design.

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

JEL Classification: D82.

*We thank Laura Doval, Piotr Dworzak, Johannes Hörner, Michael Peters, Martin Pollrich, Balazs Szentes, and Takuro Yamashita for extremely valuable feedback. We also thank seminar audiences at Brown University, EIEF, Higher School of Economics, LSE, LUISS, New Economic School, Northwestern University, Seminars in Economic Theory, UCLA, UC Riverside, Universität Bonn, Université de Montréal, University of British Columbia, University of Chicago, and the Virtual Market Design Seminar, as well as conference participants at the 2020 Econometric Society World Congress, the 2021 ASSA Virtual Annual Meeting, the 2021 CEPR Workshop on Contracts, Incentives, and Information, the 2021 Università degli Studi di Roma "Tor Vergata" Workshop on Markets, Information, and the Macroeconomy, the 2022 Conference of the Society for the Advancement of Economic Theory, the 2022 Conference on Mechanism and Institution Design, the 2022 European Summer Symposium in Economic Theory, and the 2022 Stony Brook International Conference on Game Theory for many useful discussions. This research has benefited from financial support of the ANR (Programme d'Investissements d'Avenir ANR-17-EURE-0010), the MIUR (PRIN 2015), and the research foundation TSE-Partnership.

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

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

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

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

1 Introduction

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

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

Our only departure from classical mechanism-design theory consists in letting every principal privately disclose information to the agents about her effective decision rule, that is, the mapping describing how the decision she implements responds to the messages she in fine receives from the agents. We model such private disclosures as contractible private signals, one for each agent, each summarizing what the corresponding agent knows about the principal's effective decision rule. The key point is that different agents may now have different beliefs about how the principal's decision responds to their messages. In the resulting competing-mechanism game with private disclosures, each principal fully commits, as parts of the mechanism she posts, to a distribution of private signals and to

an extended decision rule mapping the private signals she initially sends to the agents and the messages she ultimately receives from them into a (possibly random) decision. It follows from the Revelation Principle (Myerson (1979, 1989)) that the extension of the game we thereby propose has no bite in the case of a single principal who controls all the dimensions of the allocation, as private disclosures do not affect the set of allocations that can be implemented. In practice, such private disclosures may be used by sellers privately informing buyers in an auction of personalized reservation prices before they submit their bids, by manufacturers privately informing retailers of the amount of output they supply, or by policymakers privately informing voters of the characteristics of a public good implemented as a result of the solicitation of the voters' preferences.

We identify two new channels through which private disclosures modify equilibrium behavior in competing-mechanism games.

Our first main result is that private disclosures may enable principals to guarantee higher equilibrium payoffs relative to what they could do with standard mechanisms. Hence equilibria in standard mechanisms need not be robust to private disclosures. To establish this result, we provide an example of a competing-mechanism game in which message spaces are sufficiently rich for the two principals to post *recommendation mechanisms*, whereby, in line with Yamashita (2010), each agent can recommend a direct mechanism and make a report about his type to each principal. In this game, an extreme version of Yamashita's (2010) folk theorem holds when private disclosures are not feasible: namely, any feasible payoff vector for the principals can be supported in equilibrium using standard mechanisms.¹ However, we show that, by posting a mechanism that asymmetrically discloses information about her decision to the agents, one of the two principals can ensure that the agents no longer have the incentives to carry out the punishment in the other principal's mechanism that would be necessary to make the deviation unprofitable. Intuitively, this is because, by privately informing one of the agents of her decision while keeping the others in the dark, this principal perfectly aligns that agent's preferences with hers, making the latter no longer willing to participate in the required punishment. As a result, this principal can guarantee herself a payoff strictly above her minimum feasible payoff, regardless of the mechanism posted by the other principal and of the continuation equilibrium played by the agents. Many equilibrium payoffs that can be supported when principals compete in standard mechanisms can thus no longer be supported when, in addition to standard mechanisms, principals can also post mechanisms with private disclosures. The upshot

¹This reflects that, in the example, the min-max-min payoff of each principal—computed with respect to her opponent's mechanism, her own mechanism, and the agents' continuation equilibrium—is equal to her minimum feasible payoff.

of this example is that equilibrium outcomes and payoffs of competing-mechanism games without private disclosures—in particular, those supported by recommendation mechanisms as in Yamashita’s (2010) folk theorem—need no longer be supported once principals can engage into private disclosures.

Our second main result is that private disclosures may enable principals to achieve new equilibrium payoffs relative to what they could do with standard mechanisms. Specifically, we show that equilibrium outcomes and payoffs of competing-mechanism games with private disclosures need not be supported in any game in which principals are restricted to posting standard mechanisms, no matter how rich the message spaces are. The reason is that private disclosures help the principals correlate their decisions with the agents’ types in a way that cannot be replicated by principals responding to the agents’ messages when these are solely based on their types and on their common knowledge of the mechanisms. To establish this result, we provide an example of a competing-mechanism game with private disclosures in which the equilibrium correlation between the principals’ decisions and the agents’ types requires that (1) the agents receive information about one principal’s decision and pass it on to the other principal *before* the latter finalizes her own decision, and (2) such information does not create common knowledge among the agents about the former principal’s decision before they communicate with the latter. The example illustrates the possibility to achieve both (1) and (2) with private disclosures, and the necessity of both (1) and (2) when it comes to supporting certain outcomes and payoffs, which implies that it is impossible to support the latter with standard mechanisms, no matter how rich the message spaces are. In equilibrium, the requirement (2) is satisfied by letting one principal send private signals, each of which is completely uninformative of her decision from any agent’s perspective, but which together, once they are passed on by the agents to the other principal in an incentive-compatible way, perfectly reveal the principal’s decision; thus private disclosures in this example play the role of an encrypted message that one principal passes on to the other through the agents while keeping the latter in the dark. The upshot of this example is that standard mechanisms, and in particular the *universal mechanisms* of Epstein and Peters (1999), are no longer canonical once principals can engage into private disclosures.

Taken together, the above results imply that the sets of equilibrium outcomes and payoffs of competing-mechanism games with and without private disclosures are not nested. These results highlight new strategic effects in competing-mechanism games, which hinge on the principals designing the market information to be disclosed to the agents, and which cannot be captured by the existing modeling approaches. This suggests that private disclosures from the principals to the agents should be central for equilibrium characterization in the

theory of competing mechanisms.

These results are not mere theoretical curiosities: they have implications for how firms compete in markets. Our first result suggests that auctioneers may do better by disclosing their reserve prices to some bidders while keeping them secret to others, a practice that some auctioneers have started following in recent years but whose merits have not yet been explored theoretically.² Our second result, in turn, suggests that manufacturers may collude more effectively by asymmetrically informing common retailers of how their supplies respond to the latter’s private information about market conditions. This is because private disclosures allow the manufacturers to relax the retailers’ incentive-compatibility constraints, thus facilitating coordination without resorting to illegal explicit collusive agreements.

Related Literature

This paper contributes to the theoretical foundations of competing-mechanism games. In a seminal paper, McAfee (1993) points out that characterizing the equilibria of such games may require to let agents report *all* their private information to the principals; that is, their exogenous information about their types and their endogenous market information about the mechanisms posted by the other principals. To overcome the resulting infinite-regress problem, Epstein and Peters (1999) construct a space of universal mechanisms and establish an analog of the revelation principle for competing-mechanism games: 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. Subsequent work has focused on providing explicit characterizations of the equilibrium outcomes of such games. Yamashita (2010) shows that, if there are at least three agents, every deterministic incentive-compatible allocation yielding each principal a payoff at least equal to a well-defined min-max-min bound can be supported in equilibrium. This folk-theorem-type of result is established by letting each principal commit to a recommendation mechanism, whereby the agents, in addition to reporting their types, are asked to vote on the direct mechanism the principal should use. This, in turn, enables the agents to implement punishments adapted to any deviation by the other principals. From our perspective, the key point is that these two central results are established under the assumption that principals are restricted to posting standard mechanisms, so that all private communication is from the agents to the principals—in particular, both universal and recommendation mechanisms are instances of standard mechanisms. We show that these results do not extend to games in which

²In a similar vein, ad buyers are sometimes left in the dark about the precise auction format they are bidding in (see, for instance, <https://digiday.com/marketing/ad-buyers-programmatic-auction>).

the principals can privately disclose information to the agents, and, by so doing, generate asymmetric market information among the latter.

Peters and Troncoso-Valverde (2013), in a different context, argue that any allocation that is incentive-compatible and individually rational in the sense of Myerson (1979) can be supported in equilibrium provided there are sufficiently many players—there is no distinction between principals and agents in their setup. A distinctive feature of their modeling approach is that each player commits to a mechanism and to irreversibly sending an encrypted message about her type before observing the mechanisms posted by the other players and privately communicating with them. Each player, in particular, sends her encrypted type before knowing whether or not she will have to participate in punishing some other player, which allows for harsh punishments that are not incentive-compatible once the mechanisms are observed. By contrast, our approach fits more squarely into classical mechanism-design theory by maintaining the standard distinction between principals and agents, as well as the standard informational assumption that agents do not communicate among themselves and release no information before observing the mechanisms posted by the principals.

In classical mechanism-design theory (Myerson (1982)), private communication from a single principal to the agents is key when certain payoff-relevant actions can solely be taken by the agents, as in moral-hazard settings. Such communication, taking the form of action recommendations to the agents, has been shown to serve as a correlating device between the players' actions in several problems of economic interest, such as Rahman and Obara's (2010) model of partnerships. Perhaps surprisingly, however, private disclosures have been neglected even in competing-mechanism settings where agents take payoff-relevant actions, as in Prat and Rustichini's (2003) model of lobbying. To the best of our knowledge, the only exception is the recent work of Attar, Campioni, and Piaser (2019), who study complete-information games in which agents take observable actions. They construct an example in which equilibrium allocations supported by standard mechanisms fail to be robust against a deviation by a principal to a mechanism with private recommendations. In equilibrium, principals correlate their decisions with the agents' actions in a way that cannot be achieved without private recommendations. In a pure incomplete-information environment, Attar, Campioni, and Piaser (2013) explore the idea that principals may recommend to the agents which reports to make in their opponents' (direct) mechanisms. Such recommendations are shown to correlate final decisions; however, any correlation achieved in this way can also be obtained by letting the agents randomize over the messages they send to the principals. Hence private communication from the principals to the agents plays no essential role in that paper, in contrast with the present work.

In the above papers, private signals from the principals to the agents play a role similar to the one they play in single-principal settings. In contrast, the present paper uncovers two novel roles for private disclosures—raising the principals’ individual payoff guarantees and enabling principals to correlate their decisions with the agents’ types in ways that could not be achieved through standard mechanisms—that make them fundamentally different from action recommendations.

Private disclosures generate endogenous asymmetric information about the decision rule effectively followed by a principal. A similar role is played by the private information held by the agents when contracting is bilateral (see, for instance, Segal and Whinston (2003)). The literature, however, focuses on situations in which a single principal contracting with multiple agents cannot commit to a public mechanism specifying how she responds to the messages received from the agents. This inability has important welfare implications. In vertical contracting, for example, efficient production may not obtain because a retailer rejects profitable offers given his beliefs about the trades that a manufacturer conducts with other retailers (Rey and Tirole (1986), Hart and Tirole (1990), McAfee and Schwartz (1994), Segal (1999)).³ In our setting, instead, the principals can fully commit to their mechanisms. They strategically choose to disclose private information to the agents in order to discipline how the agents behave towards their opponents. Our results suggest that, in several markets of interest, private (or secret) contracting need not be the result of high transaction costs of processing information, or of the limits imposed to multilateral contracting by the relevant antitrust laws. A principal may find it optimal to asymmetrically inform the agents even if she is able to publicly commit to a mechanism.

The role of private disclosures uncovered by our analysis appears particularly relevant for the applications of competing-mechanism games favored in the literature, such as competing auctions (McAfee (1993), Peters (1997), Peters and Severinov (1997), Virág (2010)) or competitive search (Moen (1997), Eeckhout and Kircher (2010), Wright, Kircher, Julien, and Guerrieri (2021)). In these settings, contracting is typically decentralized—so that principals may find it difficult to rely on a common mediator to coordinate their decisions—and direct communication among the principals—either in the form of principals exchanging cheap-talk messages, or in the form of principals using semi-private correlation devices whose realizations are observed by the other principals but not by the agents at the time they communicate with the principals—is unlikely to be feasible.

In our setting, the correlation between the principals’ decisions is generated by the agents

³These papers focus on complete-information settings. However, similar results obtain under incomplete information about the retailers’ characteristics (Dequiedt and Martimort (2015)). See also Akbarpour and Li (2020) for recent work on credible mechanisms when principals can only commit to bilateral contracts.

communicating to each principal both their exogenous types and the endogenous private signals they receive from the other principals. However, principals cannot directly condition on each other's decisions and/or mechanisms, nor directly exchange information among themselves. By contrast, Kalai, Kalai, Lehrer and Samet (2010), Peters and Szentes (2012), Peters (2015), and Szentes (2015) suppose that players can make commitments contingent on each other's commitments, and that communication is unrestricted. The conclusions of our first example remain valid in such settings: by deviating to a mechanism with private disclosures, a principal can guarantee herself a payoff strictly above her min-max-min bound, regardless of whether or not principals can make commitments contingent on each other's decisions and/or mechanisms. The conclusions of our second example still extend to the case where principals can condition their mechanisms on their competitors' mechanisms, as in Peters' (2015) model of reciprocal contracting, but not to the case where principals can directly condition their decisions on their competitors' decisions. This is because private disclosures in this example take the form of an encrypted message passed on from one principal to the other through the agents in an incentive-compatible way, something which can be achieved in a more direct way if principals can communicate directly among themselves after receiving the agents' messages.

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

In a collusion setting à la Laffont and Martimort (1997), von Negenborn and Pollrich (2020) show that a principal can prevent collusion between an agent and a supervisor by informing these two players asymmetrically of the decision she will take in response to their reports. In their model, the benefits of private disclosures disappear when the agent and the supervisor can write contracts conditioning their side payments on the principal's final decision. Instead, we focus on settings with competing principals, and show that the benefits of private disclosures remain even when principals can condition their decisions on the other principals' decisions. Most importantly, we establish that private disclosures

undermine folk-theorem results and the construction of universal mechanisms established in the competing-mechanism literature.

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 and the robustness of the results to alternative contracting assumptions and to the availability of public correlating devices. Section 6 concludes.

2 The Model

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

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

Decisions and Payoffs Every principal j (she) takes a decision x_j in some finite set X_j . We let $v_j : X \times \Omega \rightarrow \mathbb{R}$ and $u^i : X \times \Omega \rightarrow \mathbb{R}$ be the payoff functions of principal j and of agent i , respectively, where $X \equiv X_1 \times \dots \times X_J$ is the set of possible profiles of decisions for the principals. Agents take no payoff-relevant actions, so that our setting is one of pure incomplete information. An *allocation* is thus a function $z : \Omega \rightarrow \Delta(X)$ assigning a lottery over the set X to every state of the world. The *outcome* induced by an allocation z is the restriction of z to the set of states occurring with positive probability under \mathbf{P} .⁴

Standard Mechanisms In pure incomplete-information environments such as those we focus on this paper, the standard notion of a mechanism is that of a decision rule, announced to all the agents, selecting a (possibly random) decision for every profile of messages the principal may receive from the agents (Myerson (1979)). Formally, a *standard mechanism* for principal j is a decision rule $\phi_j : M_j \rightarrow \Delta(X_j)$ assigning a lottery over principal j 's decisions to every profile of messages $m_j \equiv (m_j^1, \dots, m_j^I) \in M_j$ she receives from the agents, where $M_j \equiv M_j^1 \times \dots \times M_j^I$ for some collection of nonempty sets M_j^i of messages from every agent i to principal j . We will assume that $\text{card} \Omega^i \leq \text{card} M_j^i$ for all i and j , so that the

⁴The distinction between an allocation and an outcome is relevant when the agents' types are correlated.

language through which agent i communicates with principal j is rich enough for him to report his type to her. Unless otherwise stated, we will also assume for convenience that the sets M_j^i are finite for all i and j . We will comment in due time on how the results extend to infinite message and signal spaces.

Mechanisms with Private Disclosures A *mechanism with private disclosures* is one in which a principal informs the agents privately of how her decision responds to their messages. The decision rule is now indexed by a family of parameters, one for each agent, where each parameter summarizes what the corresponding agent knows about the decision rule effectively followed by the principal. These parameters are drawn from a joint distribution that is part of the description of the mechanism. These draws in turn determine the principal's effective decision rule, after which the mechanism reveals to each agent his dedicated parameter in the form of a private signal. In fine, different agents may have different information about how the principal's decision responds to their messages. For instance, one agent may exactly know the decision rule the principal effectively follows, while another agent may be completely kept in the dark.

Formally, a mechanism with private disclosures consists of a joint probability distribution over the signals that the principal privately sends to the agents, and of an extended decision rule that assigns a lottery over her decisions to every profile of signals she sends to the agents and every profile of messages she receives from them. A mechanism with private disclosures for principal j is thus a pair $\gamma_j \equiv (\sigma_j, \phi_j)$ such that

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

Again, unless otherwise stated, we will assume that $\text{card} \Omega^i \leq \text{card} M_j^i$ and that the sets S_j^i and M_j^i are finite for all i and j . The space of mechanisms with private disclosures for every principal j is then $\Gamma_j \equiv \Delta(S_j) \times \Delta(X_j)^{S_j \times M_j}$, which is a compact and convex set in a Euclidean space.

For every draw $s_j \in S_j$ from σ_j , principal j 's effective decision rule is given by $\phi_j(s_j, \cdot) : M_j \rightarrow \Delta(X_j)$. The private signal s_j^i agent i receives from principal j is thus a private

disclosure from principal j to agent i about principal j 's effective decision rule $\phi_j(s_j, \cdot)$. Notice that a standard mechanism for principal j is a special case of a mechanism with private disclosures in which S_j^i is a singleton for all i . We will assume that $\text{card } \Omega^i \leq \text{card } M_j^i$ for all i and j , so that the language through which agent i communicates with principal j is rich enough for him to report his type to her.

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

1. the principals simultaneously post mechanisms with private disclosures and accordingly send to the agents private signals about their effective decision rules;
2. after observing their types, the principals' mechanisms, and their private signals, the agents simultaneously send messages to the principals;
3. the principals' decisions are implemented and all payoffs accrue.

A degenerate case of the game G^{SM} arises when S_j^i is a singleton for all i and j , so that principals cannot engage into private disclosures and can only post standard mechanisms. To distinguish this situation, we denote by G^M the corresponding competing-mechanism game without private disclosures; the games studied by Epstein and Peters (1999) and Yamashita (2010) are prominent examples.⁵ It should be noted that the only difference between the games G^{SM} and G^M is that, in the former, different agents may have different information about the principals' effective decision rules.

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

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

for all $(\omega, x) \in \Omega \times X$, where $S \equiv S_1 \times \dots \times S_J$ and $M \equiv M_1 \times \dots \times M_J$. For every profile of mechanisms γ , a behavior strategy for agent i in the subgame γ played by the

⁵Notice, however, that the message spaces for universal mechanisms in Epstein and Peters (1999) are by construction uncountably infinite, as they allow each agent to report to each principal his market information in addition to his type. We refer to Online Appendix A for a discussion of the measure-theoretic intricacies that have to be dealt with in this case.

agents is a function $\beta^i : S^i \times \Omega^i \rightarrow \Delta(M^i)$ assigning a lottery over the profile of messages $m^i \in M^i$ she sends to the principals to every profile of signals $s^i \in S^i$ she may receive and to every realization $\omega^i \in \Omega^i$ of her type. We let $z_{\gamma, \beta}$ be the allocation induced by the profile of behavior strategies $\beta \equiv (\beta^1, \dots, \beta^J)$ 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 . We denote by $\lambda^i(\gamma)$ the behavior strategy induced by the strategy λ^i in the subgame γ .

Equilibrium In line with the standard practice of the common-agency literature (Peters (2001), Martimort and Stole (2002)) and the competing-mechanism literature (Epstein and Peters (1999), Yamashita (2010), Peters (2014), Szentes (2015)), we will assume that the agents treat the mechanisms posted by the principals as given. This means that we can identify any subgame $\gamma \equiv (\gamma_1, \dots, \gamma_J) \in \Gamma$ of G^{SM} with a Bayesian game played by the agents, with type space $S^i \times \Omega^i$ and action space M^i for every agent i , and in which the agents' beliefs are pinned down by the prior \mathbf{P} and the signal distributions $(\sigma_1, \dots, \sigma_J)$ to which the principals are committed through the mechanisms they post in γ , whether or not γ is reached on the equilibrium path. The strategy profile (μ, λ) is a perfect Bayesian equilibrium (PBE) of G^{SM} whenever

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

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

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

It follows from the definition of a BNE in any subgame played by the agents that any allocation $z_{\mu, \lambda}$ supported by a PBE (μ, λ) of G^{SM} is incentive-compatible; otherwise, some type ω^i of some agent i would be strictly better off mimicking the strategy $\lambda^i(\cdot | \cdot, \cdot, \hat{\omega}^i)$ of some other type $\hat{\omega}^i$ —this is an instance of the revelation principle (Myerson (1979)). This is the reason why, when there is a single principal, any allocation that can be implemented by a mechanism with private disclosures can also be implemented by 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.

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

3 Nonrobustness of Standard Mechanisms: How to Raise Payoff Guarantees with Private Disclosures

In this section, we show how principals may guarantee themselves higher payoffs by posting mechanisms with private disclosures—for example, in the case of competing auctions, by disclosing reservation prices only to a subset of bidders. This shows that equilibrium outcomes and payoffs of G^M games need not be robust to the introduction of private disclosures. The issue is especially relevant in light of the fact that, as shown by Yamashita (2010), G^M games in which principals can post standard mechanisms with rich message spaces typically lend themselves to folk-theorem types of results.⁶

Yamashita’s (2010) Theorem The construction of Yamashita (2010), which we exploit in Section 3.1 below, is based on the idea that, given rich enough message spaces, each principal’s equilibrium mechanism can be made sufficiently flexible to punish other principals’ potential deviations. This can be achieved by enabling the agents to recommend to every principal j a (deterministic) direct mechanism $d_j : \Omega \rightarrow X_j$ selecting a decision for any profile of reported types she may receive from them. Specifically, consider a G^M game in which every message space M_j^i is sufficiently rich to enable agent i to recommend any direct mechanism to principal j and to make a report about his type; that is, denoting by D_j the finite set of all such direct mechanisms, let us assume that $D_j \times \Omega^i \subset M_j^i$ for all i and j . Accordingly, a *recommendation mechanism* ϕ_j^r for principal j stipulates that, if every agent i sends a message $m_j^i \equiv (d_j^i, \omega^i) \in D_j \times \Omega^i$ to principal j , then

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

where \bar{x}_j is some fixed decision in X_j ; if, instead, some agent i sends a message $m_j^i \notin D_j \times \Omega^i$ to principal j , then ϕ_j^r treats this message as if it coincided with some fixed element $(\bar{d}_j, \bar{\omega}_j^i)$ of $D_j \times \Omega^i$, once again applying rule (1). Yamashita (2010) exploits recommendation mechanisms to establish the following folk theorem: if $I \geq 3$, then every deterministic incentive-compatible allocation yielding each principal a payoff at least equal to a well-defined min-max-min payoff bound can be supported in an equilibrium of G^M .⁷ The intuition

⁶Similar 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).

⁷As 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

for this result is that recommendation mechanisms allow the agents to select appropriate punishments following any unilateral deviation by some principal, taking the form of direct mechanisms to be implemented by the nondeviating principals.

Below we show that the possibility for principals to use private disclosures undermines this characterization result. To this end, we provide an example in which a principal, by asymmetrically disclosing to the agents how her decision responds to their messages, eschews the punishments needed to support payoffs close to her min-max-min payoff bound. The intuition is that well-chosen private disclosures prevent the agents from coordinating on the appropriate course of action, even if the other principals post recommendation mechanisms. In the example, one of two principals can achieve this goal by randomly drawing and privately disclosing her decision to one of three agents, while keeping the two other agents in the dark. Doing so enables the principal to align the selected agent’s preferences with her preferences. As a result, this agent can no longer be induced to participate in punishing the principal to the extent required to keep the latter’s payoff down to its min-max-min level. Actually, such a punishment would be possible only if the other principal’s mechanism did not respond at all to the selected agent’s messages. However, because, in the example, the remaining agents have neither the information nor the incentives to carry out the appropriate punishments themselves, this implies that this principal can guarantee herself strictly more than her min-max-min payoff, regardless of the mechanism posted by the other principal and the continuation equilibrium played by the agents. By design, the example abstracts from institutional details that may matter in applications. However, it illustrates in simple terms the power of private disclosures and how they can be used by principals to guarantee themselves higher payoffs.

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

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}	5, 8, 8	5, 1, 1
x_{12}	6, 4.5, 4.5	6, 4.5, 4.5

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

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

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

payoffs are constant over $X \times \Omega$ and hence play no role in the analysis.⁸

Roadmap Our argument in the remainder of this section is twofold. We first argue in Section 3.1 that, in Example 1, a folk theorem holds for any G^M game with rich enough message spaces. We next argue in Section 3.2 that a continuum of equilibrium payoff vectors of any such game can no longer be supported when principals can post mechanisms with private disclosures. Taken together, these two results imply that equilibrium outcomes and payoffs supported by standard mechanisms need not be robust.

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. Notice that the existence of a BNE in every subgame $\phi \equiv (\phi_1, \phi_2)$ of G_1^M is guaranteed by the fact that all the type spaces Ω^i and all the message spaces M_j^i are finite. Our first result characterizes the set of equilibrium payoffs for P2 in G_1^M .

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

The arguments in the proof rely on the same intuition as in Yamashita (2010, Theorem 1). The possibility for the agents to recommend a different direct mechanism to P1 for every mechanism posted by P2 enables them to implement punishments contingent on P2's deviations. Observe that, unlike Yamashita (2010), we allow principals to post stochastic mechanisms; yet the threat of agents choosing a deterministic direct mechanism is sufficient to yield P2 her minimum feasible payoff of 5 in equilibrium.⁹ We show that this in turn permits one to support any payoff for P2 in the feasible set $[5, 6]$ in equilibrium.¹⁰

⁸That A1's and A2's types are perfectly correlated, as well as that P1's and A3's preferences are constant in the game, simplifies some of the derivations, but is not essential for the result.

⁹Stochastic mechanisms, however, can be used to support random allocations; see, for instance, Xiong (2013).

¹⁰The proof only requires to modify Yamashita's (2010) definition of a recommendation mechanism to allow principals to randomize over their decisions on path, while maintaining the assumption that the agents' message spaces are finite.

In related work, Peters and Troncoso-Valverde (2013) establish a folk theorem in a generalized version of Yamashita (2010). In the game they study, any outcome corresponding to an allocation that is incentive-compatible and individually rational in the sense of Myerson (1979) can be supported in equilibrium provided there are at least seven players. The outcomes we construct in the proof of Lemma 1 obviously satisfy these conditions, which implies that they can also be supported in equilibrium in their framework.¹¹ 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 the example, posting a recommendation mechanism is sufficient for P1 to inflict on P2 her minimum feasible payoff of 5, leaving no role for additional messages beyond those contained in $D_1 \times \Omega^i$ for all i . In other words, that P2's relevant min-max-min payoff is equal to 5 is fairly uncontroversial.

3.2 Using Private Disclosures to Eschew Punishments

We now show that many of the equilibrium payoffs characterized in Lemma 1 cannot be supported when private disclosures are allowed for. Specifically, Lemma 2 below shows that, in any enlarged game in which principals can post mechanisms with private disclosures, P2 can guarantee herself a payoff strictly 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. It should be noted that the assumption that all the signal spaces S_j^i and the message spaces M_j^i are finite guarantees that the result is not driven by the possible nonexistence of equilibrium.¹²

The proof of Lemma 2 exploits the fact that, by posting a mechanism with private disclosures, P2 can asymmetrically inform the agents of her decision. Specifically, we construct a mechanism for P2 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 when private disclosures are allowed for, is precisely what enables P2 to guarantee herself a payoff strictly above her min-max-min payoff of 5 regardless of the mechanism posted by P1 and of the agents' continuation equilibrium strategies.

To see this, notice that the only way to keep P2's payoff down to 5 is for P1 to take decision x_{11} in state (ω_L, ω_L) and decision x_{12} in state (ω_H, ω_H) . However, by privately informing

¹¹The requirement on the number of players can be met by adding additional agents identical to A3.

¹²As the arguments below reveal, the second part of Lemma 2, which provides a lower bound for P2's payoff in G_1^{SM} regardless of the mechanism posted by P1, does not hinge on this simplifying assumption and extends to any infinite game G_1^{SM} that admits an equilibrium.

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 P2's; 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.¹³ 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 strictly higher than $\frac{1}{2}$, A2, without receiving further information from P2, strictly prefers to induce P1 to take decision x_{11} in both state. Because, as we just argued, she indeed has the possibility to do so, we conclude that she has no incentive to induce the distribution over X_1 that inflicts on P2 her min-max-min payoff of 5.

Lemma 2 more generally characterizes an interval of P2's equilibrium payoffs in G_1^M that cannot be supported when private disclosures are allowed for.

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

Lemma 2 constructs 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 richness of the signal spaces S_1^i and of the message spaces M_1^i used by P1 in G_1^{SM} . In particular, replacing all sums by the appropriate integrals in the proof of Lemma 2 reveals that this bound remains relevant even if some agent can send infinitely many messages to P1—provided, of course, an equilibrium still exists.

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 inflicts 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 collectively selecting a direct mechanism that punishes her in case she deviates.

¹³Otherwise, 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.

Lemmas 1–2 together imply our first main 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 disclosures. In particular, PBE payoff vectors of competing-mechanism games without private disclosures but with rich message spaces such that $D_j \times \Omega^i \subset M_j^i$ for all i and j need not be supportable once principals can engage into private disclosures.*

Proposition 1 points to the nonrobustness of equilibrium outcomes sustained by imposing that principals inform the agents symmetrically about the functioning of their mechanisms. As anticipated above, the practical relevance of the result comes from the fact that, in a competitive environment, a principal may have strong incentives to eliminate such outcomes, thus alleviating the competition from other principals. Private disclosures offer an effective tool to raise payoff guarantees in competitive settings.

4 Non-universality of Standard Mechanisms: How to Sustain New Payoffs with Private Disclosures

In the previous section, we have shown that equilibrium outcomes of G^M games in which principals can post mechanisms with potentially rich message spaces, need not be robust to the possibility for principals to engage into private disclosures. In this section, we address the dual question of whether G^{SM} games may admit equilibria whose outcomes and payoffs cannot be supported in G^M games, no matter how rich the message spaces are.

We provide an example showing that this is indeed the case. In this example, one of two principals exploits private disclosures to make the agents' messages to the other principal depend on information that correlates with her own decision. In the equilibrium we construct, every signal sent by this principal conveys no information about her decision to the agent who receives it; but, taken together, the signals sent by this principal perfectly reveal her decision. In equilibrium, the agents truthfully reveal their signals to the other principal, which enables the principals to perfectly correlate their decisions in a state-dependent way while respecting the agents' incentives. Thus private disclosures compensate for the lack of a direct communication channel between the principals. The key point is that the correlation between the principals' decisions and the agents' exogenous private information achieved in this equilibrium cannot be replicated, when private disclosures are not feasible, by the principals independently mixing over the mechanisms they post and/or the agents independently mixing over the messages they send. Indeed, in the absence of such disclosures,

correlation may only stem from two sources. First, randomization by the principals over their mechanisms may act as a sunspot enabling them to correlate their decisions even if agents do not mix in equilibrium, an observation first made by Peck (1997).¹⁴ In the example, however, incentive compatibility for the agents requires that the equilibrium joint distribution over the principals' decisions be constant across all the profiles of mechanisms that are selected with positive probability, which rules out this source of correlation from the outset. Second, an agent may be able to induce some correlation between the principals' decisions by correlating the messages he sends to them.¹⁵ However, any equilibrium supported by such a mixing by the agents must be immune to the temptation for the agents to *decorrelate* the messages they send to the principals, and hence the decisions they induce in the latter's mechanisms. In the example, this kind of deviations turns out to impose strong restrictions on the allocations that can be supported with standard mechanisms. As a result, no matter how rich the message spaces are, standard mechanisms do not permit one to support in equilibrium all the outcomes allowed for by private disclosures. Again, the example below is purposely simple but illustrates well how private disclosures can expand the set of equilibrium outcomes and permit the principals to obtain payoffs that they could not obtain with standard mechanisms.

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

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

Roadmap Our argument in the remainder of this section is twofold. We first construct in Section 4.1 an equilibrium of a competing-mechanism game with private disclosures in which principals' decisions in Example 2 are perfectly correlated with the state and in which P2 obtains her maximal feasible payoff of 10. We next argue in Section 4.2 that P2 cannot reach this payoff in any equilibrium of any competing-mechanism game without private

¹⁴Starting with Epstein and Peters (1999), the competing-mechanism literature has however disregarded this source of correlation by restricting attention to equilibria in which principals play pure strategies.

¹⁵As pointed out by Martimort and Stole (2002), this effect may be already at play when there is a single agent. Again, most of the competing-mechanism literature has disregarded this source of correlation by restricting attention to equilibria in which agents play pure strategies. See, however, Attar, Campioni, and Piaser (2013) for an example showing that this form of correlation can arise even if agents can only report their exogenous private information.

	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 .

disclosures. Taken together, these two results imply that standard mechanisms need not exhaust the range of outcomes and payoffs that can be supported when the principals can asymmetrically inform the agents of the decision rules they effectively follow.

4.1 Using Private Disclosures to Correlate Principals' Decisions while respecting Agents' Incentives

To illustrate the key ideas in the simplest possible manner, we consider a specific competing-mechanism game G_2^{SM} with private disclosures in which only P2 can send signals to the agents, and these signals are binary; that is, we let $S_1^1 = S_1^2 \equiv \{\emptyset\}$ and $S_2^1 = S_2^2 \equiv \{1, 2\}$. Furthermore, we consider the simplest possible message spaces that allow the agents to report their private information to the principals; that is, we let $M_1^i \equiv \Omega^i \times S_2^i$ and $M_2^i \equiv \Omega^i$ for all i .¹⁶ The following result then holds.

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

$$z(\omega_L) \equiv \alpha \delta_{(x_{13}, x_{21})} + (1 - \alpha) \delta_{(x_{14}, x_{22})}, \quad (2)$$

$$z(\omega_H) \equiv \alpha \delta_{(x_{12}, x_{21})} + (1 - \alpha) \delta_{(x_{11}, x_{22})}, \quad (3)$$

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

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

Our equilibrium construction relies on a mechanism with private disclosures for P2 whereby she commits to the decision x_{21} if the signals she sends to A1 and A2 match, and to the decision x_{22} otherwise. P2 chooses her joint probability distribution over profiles of signals in $S_2^1 \times S_2^2 = \{1, 2\} \times \{1, 2\}$ so as to keep both agents in the dark: regardless of the signal he receives from P2, every agent's posterior belief about P2's decision coincides with

¹⁶As the arguments below reveal, Lemma 3 does not hinge on these simplifying assumptions, and extends to games with richer signal and message spaces as long as $\Omega^i \times S_2^i \subset M_1^i$ and $\Omega^i \subset M_2^i$ for all i .

his prior belief. Private disclosures by P2 can thus be interpreted as *encryption keys*: taken in isolation, every signal sent by P2 is completely uninformative of her decision, whereas, taken together, the two signals sent by P2 are perfectly informative of her decision; we will return to this point in Section 5. P1's mechanism, in turn, can be designed to both elicit agents' information about their types and the signals they receive from P2, and, thanks to the encryption device that allows one to infer P2's decision from the signals P2 sends to the agents, to use this information to perfectly correlate P1's decision with P2's and the state of nature. The bulk of the argument consists in checking that the agents indeed have the incentives to report truthfully to P1. Notice in that respect that, if P2 were to inform the agents of her decision, then, after learning that P2 takes decision x_{21} , A2, when of type ω_L , would no longer be willing to induce P1 to take decision x_{13} . By claiming that his type is ω_H , type ω_L of A2 could instead induce P1 to take the decision x_{12} with certainty, obtaining a payoff of 5 instead of the payoff of 3 he obtains by being truthful.

The construction also reveals that, for P2 to obtain her maximum feasible payoff of 10 while maintaining the agents' incentives, it is essential that both principals randomize over their decisions, albeit in a perfectly correlated manner. From a purely technical viewpoint, the task of correlating the principals' decisions can be fully delegated to the agents by letting the agents randomize over the messages they send to the principals, while letting the principals respond deterministically to the messages they receive from the agents. Though technically feasible, however, such a delegation is not incentive-compatible. The desired equilibrium correlation between the principals' decisions requires that some information be passed on from one principal to the other. As principals cannot directly communicate with each other, this is possible only through private disclosures, agents' incentives being confined to passing the required information truthfully. The analysis in Section 4.2 and the discussion in Section 5 confirm this intuition by showing that private signals are indispensable, no matter how rich the message spaces are.

4.2 Indispensability of Private Signals

We now argue that the outcome (2)–(3) in Lemma 3, for $\alpha = \frac{2}{3}$, cannot be supported in any equilibrium of any game in which the principals are restricted to posting standard mechanisms, irrespective of the richness of the message spaces. More generally, the maximal feasible payoff of 10 for P2 cannot be supported in any equilibrium of any such game. Thus private disclosures are indispensable to support the above outcome and its associated payoff for P2. To show this, we consider a general competing-mechanism game G_2^M without private disclosures, and with arbitrary message spaces M_j^i that we no longer require to be finite.

This general formulation notably allows us to capture the case where every principal j 's message spaces are large enough—namely, uncountable Polish spaces—to encode the agents' information about the mechanism posted by her opponent, as in Epstein and Peters (1999)'s construction of universal mechanisms.

The structure of the argument can be sketched as follows.

Suppose, by way of contradiction, that there exists a distribution over pairs of standard mechanisms and a pair of continuation equilibrium strategies for the agents such that P2 obtains her maximum feasible payoff of 10. Then, because the principals' decisions must be perfectly correlated in both states, every pair of mechanisms posted by the principals must respond deterministically to the messages sent by the agents on path.

The desired correlation should thus be induced by the players' independent mixing behavior—that is, principals randomizing over the mechanisms they post and/or agents randomizing over the messages they send in those mechanisms. In either case, the correlation between the principals' decisions must ultimately obtain as an implication of incentive-compatible choices by the agents. Hence we only have to show that, if private disclosures are not allowed for, and regardless of the source of randomness, there exist no continuation equilibrium strategies for the agents that induce the desired correlation. The proof of this result consists of two steps.

Observe first that, because only A2 observes the state, when the distribution over the principals' decisions in any subgame reached on the equilibrium path is state-dependent, A2 must weakly prefer the distribution of messages he is supposed to carry out in each state to the one he is supposed to carry out in the other state. We show that this restricts the joint distribution over the principals' decisions to be constant across all such subgames. The proof relies on the possibility for A2 to decorrelate the decisions he is able to induce in the principals' mechanisms by drawing the message he sends to P1 from his continuation equilibrium strategy in state ω_H and by independently drawing the message he sends to P2 from his continuation equilibrium strategy in state ω_L . It turns out that, in any subgame reached on the equilibrium path, A2 can increase his payoff in state ω_L by behaving in this way, unless the joint distribution over the principals' decisions in this subgame is given by (2)–(3) for $\alpha = \frac{2}{3}$. Because the distribution over the principal's decisions must be the same regardless of the mechanisms they post on the equilibrium path, principals' mixing behavior is irrelevant for inducing the desired correlation.

The upshot of the properties discussed so far is that, if the outcome (2)–(3) for $\alpha = \frac{2}{3}$ can be supported in equilibrium without private disclosures, then it can be supported by each principal posting with probability 1 a deterministic mechanism, and the agents randomizing

over the messages they send to the principals. The last step of the proof shows that inducing this outcome is inconsistent with the agents' incentives. Specifically, we consider another way for A2 to decorrelate the decisions he induces in the principals' mechanisms, which consists in independently drawing twice from his continuation equilibrium strategy in state ω_H , and then in using the first and the second of these draws to determine his messages to P1 and P2, respectively. We show that, for A2 to weakly prefer the distribution over the principals' decisions he is supposed to induce in state ω_L to that induced by this alternative strategy, the messages that A2 sends in state ω_H must have no influence on the principals' decisions when combined with those sent with positive probability by A1. As a result, A1 should have full control over the final decisions in state ω_H . This, however, in turn implies that A1 has a profitable deviation, because he can induce the high-payoff decision profile (x_{11}, x_{22}) in this high-probability state. The following result then holds.

Lemma 4 *There exists no PBE of G_2^M in which P2 obtains her maximum feasible payoff of 10. In particular, there exists no PBE of G_2^M that supports the outcome (2)–(3) for $\alpha = \frac{2}{3}$.*

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

Lemmas 3–4 together imply our second main result.

Proposition 2 *PBE outcomes and PBE payoff vectors of competing-mechanism games with private disclosures need not be supported in any PBE of any competing-mechanism game without private disclosures—including, in particular, the game in which principals can post universal mechanisms—and this, even if the principals or the agents play mixed strategies in equilibrium.¹⁷*

Proposition 2 shows that the universal mechanisms of Epstein and Peters (1999) fail to be canonical when principals can engage into private disclosures. The latter enable the principals to coordinate their responses to the information privately held by the agents while respecting the agents' incentive compatibility. In so doing, private disclosures also enable the principals to attain payoffs that they would not be able to attain with standard mechanisms. Practically, in competitive settings, informing common retailers asymmetrically of how a manufacturer's output responds to the private information they hold may make it easier for

¹⁷Recall that Epstein and Peters (1999) restrict attention to equilibria in which principals and agents play pure strategies.

manufacturers to collude. More generally, private disclosures supplement for the possibility for the principals to communicate directly between themselves after consulting with the agents, something that may be precluded by antitrust legislation or by the mere nature of the relevant interactions.

Together, Propositions 1 and 2 imply that the sets of equilibrium outcomes and payoffs of competing-mechanism games with and without private disclosures are not nested. As an aside, Proposition 2 also implies that the folk theorem of Yamashita (2010) does not extend to stochastic allocations. Indeed, the allocation $(z(\omega_L), z(\omega_H))$ defined by (2)–(3) for $\alpha = \frac{2}{3}$ is certainly incentive-compatible; moreover, it yields P2 her maximum feasible payoff of 10, which is certainly at least equal to her min-max-min payoff, as defined by Yamashita (2010) over recommendation mechanisms. Nevertheless, Lemma 4 implies that this allocation cannot be supported in an equilibrium of G_2^M , even when $D_j \times \Omega^i \subset M_j^i$ for all i and j , so that recommendation mechanisms are feasible.

Finally, the proof of Lemma 4 does not make use of the property that the principals choose their mechanisms independently. The result in Lemma 4 thus carries over to the case where G_2^M is augmented by arbitrarily rich public randomizing devices that the principals may use to correlate their choices of mechanisms. On the other hand, the result does hinge on the principals not having access to private randomizing devices whose realizations are not known to the agents at the time they send their messages to the principals; we discuss this issue further in Section 5.3

5 Discussion

In this section, we put our results in perspective by discussing the different roles that private disclosures play in our examples and by examining the consequences for our results of alternative informational and contracting assumptions.

5.1 Informative versus Uninformative Signals

Each of our examples is designed to illustrate a specific role of private disclosures; we now discuss these different roles in turn.

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

they perfectly reveal it.

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

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

$$\sum_{s_2^{-i} \in S_2^{-i}} \sigma_2(s_2^{-i} | s_2^i) \phi_2(x_2 | s_2^i, s_2^{-i}, m_2) = \sum_{s_2 \in S_2} \sigma_2(s_2) \phi_2(x_2 | s_2, m_2) \quad (4)$$

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

Lemma 5 *In G_1^{SM} , if P1 posts a recommendation mechanism ϕ_1^r , then, for every mechanism γ_2 of P2 that has uninformative signals, there exists a BNE of the subgame (ϕ_1^r, γ_2) in which P2 obtains her minimum feasible payoff of 5.*

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

a nonbabbling equilibrium of the agents' subgame in which they truthfully report to P1 the uninformative signals they receive from P2.

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

5.2 Contractible Contracts and Reciprocal Mechanisms

Proposition 1 is established under the assumption that no principal can directly condition the decisions she takes and/or the mechanism she chooses on the other principals' decisions and/or mechanisms. However, the result extends to settings in which such conditioning is feasible, as in the literature on contractible contracts (Kalai, Kalai, Lehrer, Samet (2010), Peters and Szentes (2012), Szentes (2015)) and reciprocal mechanisms (Peters (2015)). To see this, observe that, in Example 1, the only way to inflict on P2 her minimum feasible payoff of 5 is for P1 to take decision x_{11} in state (ω_L, ω_L) and decision x_{12} in state (ω_H, ω_H) with probability 1. However, because the state is privately observed by A1 and A2, P1 must ultimately let them determine which decisions to implement in response to a deviation by P2, were P2 to post a mechanism with private disclosures. Now suppose, for instance, that, as in the proof of Lemma 2, P2 posts a mechanism whereby she selects a decision at random and only informs A1 of her decision. Because, when P2 selects decision x_{22} , this mechanism perfectly aligns A1's preferences with P2's in each state, P1's mechanism must not be responsive to A1's messages on pain of moving P2's payoff away from 5; notice that this remains true even if P1 can condition the decision she takes and/or the mechanism she chooses on P2's decision and/or mechanism. Thus P1 must entirely delegate to A2 the task of making her decision contingent on the state. Yet, by construction, A2 does not know which decision P2 is committed to; moreover, the additional possibility for P1 to, for instance, condition her decision on P2's is of little use if P2's payoff is to be kept down to 5, as this requirement uniquely pins down P1's decision in each state. It follows that P1 still

has no way to reward A2 for truthfully reporting the state to her, and, as in the proof of Lemma 2, that A2 has a profitable deviation. We conclude that, even if P1 can resort to contractible contracts or post a reciprocal mechanism, it is impossible for her to induce A1 and A2 to carry out the punishments necessary to block P2’s deviation.

In this respect, the case of Proposition 2 is less clear cut. On the one hand, in Example 2, the proof of Lemma 4 goes through unaltered even if each principal can condition her choice of a mechanism on the other principal’s mechanism—indeed, the argument is valid for any pair of mechanisms posted by the principals, no matter how this pair is generated. This implies that Proposition 2 extends to settings in which principals can post reciprocal mechanisms (Peters (2015)). On the other hand, Lemma 4 does hinge on each principal being unable to directly condition her decision on the other principal’s decision. Indeed, if this were possible, then the outcome (2)–(3) for $\alpha = \frac{2}{3}$ could trivially be supported even without private disclosures.

5.3 Additional Randomizing Devices

The constructions in Epstein and Peters (1999) and Yamashita (2010) do not allow for additional randomizing devices that would directly enable the principals and/or the agents to correlate their choices. We have closely followed these authors in that respect; indeed, the whole point of Example 2 is to show that such correlation can endogenously emerge in equilibrium when principals can post mechanisms with private disclosures, but not when they can only post standard mechanisms. It is nevertheless natural to ask to which extent our findings are robust to the availability of additional randomizing devices.

In Example 1, the destabilizing role of private disclosures remains relevant regardless of whether and at which stage of the game the principals and the agents have access to randomizing devices—they may be used to correlate the principals’ choices of mechanisms, the messages sent by the agents to the principals, or the decisions taken by the principals in response to these messages. To see this, we can focus on the agents’ behavior because, as pointed out in Section 5.1, the key role of private signals in the game G_1^{SM} of Example 1 is that they can be used by P2 to eschew equilibrium outcomes supported by standard mechanisms in which she obtains her minimum feasible payoff of 5. In that respect, enabling the agents to observe the realization of a *public* randomization device—a sunspot—would not change the conclusion of Lemma 2 as the mechanism for P2 constructed in the proof guarantees her a payoff strictly above 5 regardless of the continuation equilibrium played by the agents. Less obviously, this mechanism is still effective regardless of the mechanism posted by P1 even if the agents’ behavior in the corresponding subgame is coordinated by

a mediator who can first elicit information from the agents and then send them private recommendations. In line with the discussion in Section 5.2, the point is that, once A1 knows that P2 has selected decision x_{22} , he has, state-by-state, the same preferences as P2. This means that the mediator cannot extract from A1 information about the state and P2's decision and use that information to keep P2's payoff down to 5. But then the task of punishing P2 must be fully delegated to A2, which we know is impossible.

The situation in Example 2 is slightly different. As mentioned in our discussion of Proposition 2, it is impossible to support the outcome (2)–(3) for $\alpha = \frac{2}{3}$ in the game G_2^M even if the principals can correlate their choices of standard mechanisms. Moreover, the proof of Lemma 4 goes through unaltered even if the agents can coordinate their messages by means of a public randomization device; to see this, we need only notice that, for any pair of mechanisms posted by P1 and P2, and for any realization of the sunspot that enables the agents to correlate their messages, A1 and A2 must play an equilibrium in the continuation game. Thus Proposition 2 extends to settings in which the principals and the agents have access to rich public randomizing devices whose realizations are observed by all the players before committing to their choices. Things would be different if the principals could use randomization devices whose realizations are not known to the agents at the time they send their messages to the principals. Indeed, because the only role of private disclosures in the example is to pass on information from one principal to the other without changing the agents' beliefs, such private disclosures can be dispensed with if the principals 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. The value of the example is rather to show that, in the setting that has been at the center stage of the literature—in which direct communication between the principals is not feasible, private randomizing devices are not available, and principals cannot condition their decisions on other principals' decisions—private disclosures have important implications for the set of equilibrium outcomes and for the validity of the universal revelation principle established in this setting.

6 Concluding Remarks

This paper explores a novel dimension in the design of mechanisms in settings in which multiple principals contract with multiple agents—namely, the possibility for the principals to asymmetrically inform the agents about the functioning of their mechanisms, that is, of how their decisions respond to the agents' messages. We show that private disclosures

enable the principals to guarantee themselves higher payoffs relative to what they could guarantee with standard mechanisms, thus protecting them against competition from the other principals. They also permit the principals to relax the agents' incentive-compatibility constraints, thus making it possible to sustain outcomes and payoffs that cannot be sustained with standard mechanisms, no matter how rich message spaces are allowed to be. These results call for a novel approach in the study of competing-mechanism games and have important implications for the conclusions established through the analysis of such games—for instance, the merits of secret reserve prices in competing auction or the effectiveness of implicit collusion. The exploration of such implications in specific markets of interest represents an important next step for this literature and appears to be a very promising avenue for future research.

Appendix

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

Step 1 We first show that the outcome

$$z(\omega_L, \omega_L) \equiv \delta_{(x_{11}, x_{21})}, \quad z(\omega_H, \omega_H) \equiv \delta_{(x_{12}, x_{22})}, \quad (5)$$

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

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

$$d_1^*(\omega) \equiv \begin{cases} x_{11} & \text{if } \omega = (\omega_L, \omega_L) \\ x_{12} & \text{otherwise} \end{cases} \quad \text{and} \quad d_2^*(\omega) \equiv \begin{cases} x_{21} & \text{if } \omega = (\omega_L, \omega_L) \\ x_{22} & \text{otherwise} \end{cases} \quad (6)$$

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

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

$$\bar{p} \equiv \max_{m_2 \in M_2} p(m_2) \quad (7)$$

and select a message profile $\bar{m}_2 \equiv (\bar{m}_2^1, \bar{m}_2^2, \bar{m}_2^3) \in M_2$ that achieves the maximum in (7); similarly, let

$$\underline{p} \equiv \min_{(m_2^1, m_2^2) \in M_2^1 \times M_2^2} p(m_2^1, m_2^2, \bar{m}_2^3) \quad (8)$$

and select a message profile $(\underline{m}_2^1, \underline{m}_2^2) \in M_2^1 \times M_2^2$ for A1 and A2 that, given \bar{m}_2^3 , achieves the minimum in (8). That \bar{p} , \bar{m}_2 , \underline{p} , and $(\underline{m}_2^1, \underline{m}_2^2)$ are well-defined for any given ϕ_2 follows from the fact that the set M_2 is finite. We now prove that there exist BNE strategies for the agents in the subgame (ϕ_1^r, ϕ_2) such that P2 obtains a payoff of 5, so that the deviation is not profitable. We consider two cases in turn.

Case 1: $\bar{p} \geq \frac{1}{2}$ Suppose first that ϕ_2 is such that $\bar{p} \geq \frac{1}{2}$. We claim that the subgame (ϕ_1^r, ϕ_2) admits a BNE that satisfies the following properties: (i) all agents recommend the direct mechanism d_1^* to P1, as if P2 did not deviate; (ii) A1 and A2 truthfully report their types to P1; (iii) A3 sends message \bar{m}_2^3 to P2; (iv) P2 obtains a payoff of 5. As for (i), the argument is that unilaterally sending a different recommendation to P1 is of no avail as no agent is pivotal. As for (iii), sending \bar{m}_2^3 to P2 is optimal for A3 given that his payoff is constant over $X \times \Omega$. Consider then (ii). Suppose first that the state is (ω_L, ω_L) . Because $\bar{p} \geq \frac{1}{2}$, $8\bar{p} + (1 - \bar{p}) \geq 4.5$. From Table 1, and by definition of d_1^* and \bar{m}_2 , it thus follows that, if A2 reports ω_L to P1 and sends \bar{m}_2^2 to P2, and if A3 sends \bar{m}_2^3 to P2, then A1 best responds by reporting ω_L to P1 and sending \bar{m}_2^1 to P2; notice, in particular, that, because $\bar{\omega}_1^1 = \omega_L$, if A1 sends a message $m_1^1 \notin D_1 \times \Omega^1$ to P1, then P1 takes the same decision as if A1 truthfully reported his type to her. The argument for A2 is identical. Suppose next that the state is (ω_H, ω_H) . If either A1 or A2 truthfully reports his type to P1, then, by definition of d_1^* , the other informed agent A2 or A1 cannot induce P1 to take a decision other than x_{12} . These properties, along with the fact that the set M_2 is finite, imply that the subgame (ϕ_1^r, ϕ_2) admits a BNE satisfying (i)–(iii). In this BNE, P1 takes decision x_{11} in state (ω_L, ω_L) and decision x_{12} in state (ω_H, ω_H) , yielding a payoff of 5 to P2, as required by (iv). The claim follows.

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

$$d_1(\omega) \equiv \begin{cases} x_{12} & \text{if } \omega = (\omega_H, \omega_H) \\ x_{11} & \text{otherwise} \end{cases} \quad (9)$$

to P1; (ii) A1 and A2 truthfully report their types to P1; (iii) A3 sends message \bar{m}_2^3 to P2; (iv) P2 obtains a payoff of 5. The arguments for (i) and (iii) are the same as in Case 1.

Consider then (ii). Suppose first that the state is (ω_L, ω_L) . If either A1 or A2 truthfully reports his type to P1, then, by definition of d_1 , the other informed agent A2 or A1 cannot induce P1 to take a decision other than x_{11} . Suppose next that the state is (ω_H, ω_H) . Because $\underline{p} \leq \bar{p} < \frac{1}{2}$, $\underline{p} + 8(1 - \underline{p}) > 4.5$. From Table 2, and by definition of d_1 and $(\underline{m}_2^1, \underline{m}_2^2)$, it thus follows that, if A2 reports ω_H to P1 and sends \underline{m}_2^2 to P2, and if A3 sends \bar{m}_2^3 to P2, then A1 best responds by reporting ω_H to P1 and sending \underline{m}_2^1 to P2; notice, in particular, that, because $\bar{\omega}_1^1 = \omega_L$, if A1 sends a message $m_1^1 \notin D_1 \times \Omega^1$ to P1, then P1 takes the same decision as if A1 misrepresented his type. The argument for A2 is identical. These properties, along with the fact that the set M_2 is finite, 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.

Step 2 We start with a definition. An *extended recommendation mechanism* $\tilde{\phi}_j^r : M_j \rightarrow \Delta(X_j)$ for principal j implements the same decisions as the recommendation mechanism ϕ_j^r in (1), except if at least $I - 1$ agents send messages $m_j^i \equiv (d_j^0, \omega^i) \in D_j \times \Omega^i$ to principal j , for some fixed direct mechanism $d_j^0 \in D_j$, in which case principal j disregards d_j^0 and implements a (possibly stochastic) direct mechanism $\tilde{d}_j : \Omega \rightarrow \Delta(X_j)$; again, if some agent i sends a message $m_j^i \notin D_j \times \Omega^i$ to principal j , then $\tilde{\phi}_j^r$ treats this message as if it coincided with some fixed element $(\bar{d}_j, \bar{\omega}_j^i)$ of $D_j \times \Omega^i$, for some $\bar{d}_j \neq d_j^0$.

We now construct a family of PBEs of G_1^M , indexed by P2's payoff $v \in (5, 6]$, in which P2 posts the same recommendation mechanism ϕ_2^r as in Step 1 of the proof and P1 posts an extended recommendation mechanism $\tilde{\phi}_1^r$. Again, because P1's payoff is constant over $X \times \Omega$, she has no profitable deviation. If P2 deviates to some arbitrary standard mechanism $\phi_2 : M_2 \rightarrow \Delta(X_2)$, then we require that the agents' strategies implement the same punishments for P2 as in Step 1 of the proof. We suppose in particular that the direct mechanism d_1^0 differs from the direct mechanisms d_1^* and d_1 , defined by (6) and (9), which may be recommended by the agents to P1 following a deviation by P2; recall that these punishments inflict on P2 her minimum feasible payoff of 5. We consider two cases in turn.

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

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

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

We now show that, for each $\xi \in [\frac{1}{2}, 1)$, the subgame $(\tilde{\phi}_1^r, \phi_2^r)$ admits a BNE in which: (i)

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

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

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

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

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

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

We now show that, for each $\xi \in (\frac{1}{2}, 1]$, the subgame $(\tilde{\phi}_1^r, \phi_2^r)$ admits a BNE in which: (i) each agent recommends to P1 the direct mechanism d_1^0 , and recommends to P2 the direct

mechanism d_2^{**} that selects the decision x_{21} regardless of A1's and A2's reports; (b) A1 and A2 truthfully report their types to P1—because P2's decision is fixed, the messages they send to P2 are irrelevant. The corresponding payoff for P2 in the subgame $(\tilde{\phi}_1^r, \phi_2^r)$ is $v = 5 + \xi \in (5.5, 6]$ as ξ varies in $(\frac{1}{2}, 1]$, as desired. Because A3's payoff is constant over $X \times \Omega$, we only need to focus on A1's and A2's incentives.

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

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

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

Proof of Lemma 2. We first show that a PBE exists. We next establish the desired bound on P2's equilibrium payoff.

Existence of a PBE Because, for each j , the sets S_j and M_j are finite, the space $\Gamma_j \equiv \Delta(S_j) \times \Delta(X_j)^{S_j \times M_j}$ of mechanisms for principal j in G_1^{SM} is compact, and every subgame $(\gamma_1, \gamma_2) \in \Gamma_1 \times \Gamma_2$ is finite; moreover, the agents' information structures and payoffs are continuous functions of (γ_1, γ_2) . Hence the BNE of the subgame (γ_1, γ_2) form a nonempty compact set $B^*(\gamma_1, \gamma_2)$, and the correspondence $B^* : \Gamma_1 \times \Gamma_2 \rightarrow \prod_{i=1}^3 \Delta(M^i)^{S^i \times \Omega^i}$ is upper hemicontinuous (Milgrom and Weber (1985, Theorem 2)) and, therefore, admits a Borel-measurable selection $b^* \equiv (b^{1*}, b^{2*}, b^{3*})$ by Kuratowski and Ryll-Nardzewski's selection theorem (Aliprantis and Border (2006, Theorem 18.13)); the corresponding strategy for every

agent i in G_1^{SM} is defined by $\lambda^{i*}(m^i | \gamma_1, \gamma_2, s^i, \omega^i) \equiv b^{i*}(\gamma_1, \gamma_2)(m^i | s^i, \omega^i)$. Now, suppose that P1 posts the mechanism γ_1^* that equiprobably randomizes between decisions x_{11} and x_{12} regardless of the signals P1 sends to the agents and the messages she receives from them. Then, from Tables 1–2, P2 obtains an expected payoff of 5.5 regardless of the mechanism she posts. Because P1’s payoff is constant over $X \times \Omega$, it follows that, for each $\gamma_2^* \in \Gamma_2$, $(\gamma_1^*, \gamma_2^*, \lambda^{1*}, \lambda^{2*})$ is a PBE of G_1^{SM} .

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

$$\inf_{\gamma_1 \in \Gamma_1} \inf_{\beta \in B^*(\gamma_1, \gamma_2(\sigma))} \sum_{\omega \in \Omega} \sum_{x \in X} \mathbf{P}[\omega] z_{\gamma_1, \gamma_2(\sigma), \beta}(x | \omega) v_2(x, \omega) \geq 5 + \frac{(1-\sigma)\mathbf{P}[(\omega_L, \omega_L)]\mathbf{P}[(\omega_H, \omega_H)]}{1-\sigma\mathbf{P}[(\omega_L, \omega_L)]}, \quad (10)$$

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 private disclosures 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.

Therefore, 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 (10).

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 $(\gamma_1, \gamma_2(\sigma), \beta)$, P2’s payoff is $5 + \varepsilon$, where

$$0 \leq \varepsilon < \frac{(1-\sigma)\mathbf{P}[(\omega_L, \omega_L)]\mathbf{P}[(\omega_H, \omega_H)]}{1-\sigma\mathbf{P}[(\omega_L, \omega_L)]}. \quad (11)$$

Observe that the mechanism $\gamma_2(\sigma)$ implements decisions in X_2 that are independent of any messages P2 may receive from the agents and, hence, of any signals sent by γ_1 . Thus the only

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

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

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

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

$$\pi_{11}(\omega_L, \omega_L, s_2^1) \equiv \sum_{(m_1^1, m_1^2, m_1^3) \in M_1} \beta^1(m_1^1 | s_2^1, \omega_L) \beta^2(m_1^2 | \omega_L) \beta^3(m_1^3) \phi_1(x_{11} | m_1^1, m_1^2, m_1^3) \quad (13)$$

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

$$\pi_{12}(\omega_H, \omega_H) \equiv \sigma \pi_{12}(\omega_H, \omega_H, 1) + (1 - \sigma) \pi_{12}(\omega_H, \omega_H, 2), \quad (14)$$

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

$$\pi_{12}(\omega_H, \omega_H, s_2^1) \equiv \sum_{(m_1^1, m_1^2, m_1^3) \in M_1} \beta^1(m_1^1 | s_2^1, \omega_H) \beta^2(m_1^2 | \omega_H) \beta^3(m_1^3) \phi_1(x_{12} | m_1^1, m_1^2, m_1^3) \quad (15)$$

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

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

or, equivalently,

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

which implies

$$\pi_{11}(\omega_L, \omega_L) \geq 1 - \frac{\varepsilon}{\mathbf{P}[(\omega_L, \omega_L)]} \quad \text{and} \quad \pi_{12}(\omega_H, \omega_H) \geq 1 - \frac{\varepsilon}{\mathbf{P}[(\omega_H, \omega_H)]} \quad (16)$$

as both $\pi_{11}(\omega_L, \omega_L)$ and $\pi_{12}(\omega_H, \omega_H)$ are at most equal to 1. Notice that (11) ensures that the right-hand side of each inequality in (16) is strictly positive, and thus can be interpreted as a probability as it is at most equal to 1. Similarly, it follows from (12) and from the first inequality in (16) that

$$\pi_{11}(\omega_L, \omega_L, 2) \geq 1 - \frac{\varepsilon}{(1 - \sigma)\mathbf{P}[(\omega_L, \omega_L)]}. \quad (17)$$

Again, (11) ensures that the right-hand side of (17) 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 (13) and (17), that, for any message $m_1^1 \in M_1^1$,

$$\sum_{(m_1^2, m_1^3) \in M_1^2 \times M_1^3} \beta^2(m_1^2 | \omega_L) \beta^3(m_1^3) \phi_1(x_{11} | m_1^1, m_1^2, m_1^3) \geq 1 - \frac{\varepsilon}{(1-\sigma)\mathbf{P}[(\omega_L, \omega_L)]}; \quad (18)$$

otherwise, by (17), A1 could induce P1 to take decision x_{11} with a probability strictly lower than $\pi_{11}(\omega_L, \omega_L, 2)$, yielding A1 a strictly higher payoff, a contradiction. Integrating (18) with respect to the measure $\sigma\beta^1(\cdot | 1, \omega_H) + (1-\sigma)\beta^1(\cdot | 2, \omega_H)$ then yields

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

This means 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 $1 - \frac{\varepsilon}{(1-\sigma)\mathbf{P}[(\omega_L, \omega_L)]}$. Because $4.5 > \sigma + 8(1-\sigma)$ as $\sigma > \frac{1}{2}$, A2 can thus guarantee himself a payoff at least equal to

$$4.5 \left\{ 1 - \frac{\varepsilon}{(1-\sigma)\mathbf{P}[(\omega_L, \omega_L)]} \right\} + [\sigma + 8(1-\sigma)] \frac{\varepsilon}{(1-\sigma)\mathbf{P}[(\omega_L, \omega_L)]}. \quad (19)$$

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 second inequality in (16), he obtains an expected payoff at most equal to

$$4.5 \frac{\varepsilon}{\mathbf{P}[(\omega_H, \omega_H)]} + [\sigma + 8(1-\sigma)] \left\{ 1 - \frac{\varepsilon}{\mathbf{P}[(\omega_H, \omega_H)]} \right\}. \quad (20)$$

Comparing (19) and (20), 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 (11), contradicting the assumption that $\beta \in B^*(\phi_1, \gamma_2(\sigma))$. Thus $\gamma_2(\sigma)$ satisfies (10), as claimed.

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

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

The result follows. ■

Proof of Lemma 3. Let P2 post the mechanism $\gamma_2^* \equiv (\sigma_2^*, \phi_2^*)$ such that

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

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

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

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 he receives from P2, every agent's posterior distribution about P2's decision coincides with his prior distribution; that is, each agent believes that P2 takes decision x_{21} with probability α and decision x_{22} with probability $1 - \alpha$. For the same reason, each agent believes that the other agent received the same signal as his with probability α and a different signal with probability $1 - \alpha$. Thus γ_2^* keeps both agents in the dark.

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

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

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 he received signal $s_2^1 = 1$ from P2, whereas A2 reports that his type is ω_L and that he received signal $s_2^2 = 1$ from P2. Observe from (21)–(22) that the outcome (2)–(3) is implemented in the subgame (γ_1^*, γ_2^*) if every agent reports truthfully to P1 his type and the signal he receives from P2. We now show that, if $\alpha = \frac{2}{3}$, then truthful reporting is consistent with a BNE of the subgame (γ_1^*, γ_2^*) . The proof consists of two steps.

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

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

$$\begin{aligned} & \frac{1}{4} [\alpha u^1(x_{13}, x_{21}, \omega_L) + (1 - \alpha) u^1(x_{14}, x_{22}, \omega_L)] \\ & + \frac{3}{4} [\alpha u^1(x_{12}, x_{21}, \omega_H) + (1 - \alpha) u^1(x_{11}, x_{22}, \omega_H)] = 3\alpha + 7.5(1 - \alpha). \end{aligned} \quad (23)$$

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

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

which is strictly less than the value in (23) for all $\alpha \in [0, 1]$.

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

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

$$\alpha u^2(x_{13}, x_{21}, \omega_L) + (1 - \alpha) u^2(x_{14}, x_{22}, \omega_L) = 3\alpha + 7.5(1 - \alpha). \quad (24)$$

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

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

which is at most equal to the value in (24) if $\alpha \leq \frac{8}{9}$.

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

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

which is at most equal to the value in (24) if $\alpha \leq \frac{2}{3}$.

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

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

which is at most equal to the value in (24) if $\alpha \geq \frac{1}{5}$.

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

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

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

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

which is at most equal to the value in (25) if $\alpha \geq \frac{1}{4}$.

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

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

which is at most equal to the value in (25) if $\alpha \geq \frac{2}{3}$.

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

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

which is at most equal to the value in (25) if $\alpha \geq \frac{2}{3}$.

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

Proof of Lemma 4. Let Φ_j be the space of standard mechanisms, endowed with an appropriate σ -field \mathcal{F}_j ; we refer to Online Appendix A for how to define these objects when the message spaces M_j^i are uncountably infinite, as is the case in Epstein and Peters (1999). The arguments below more generally show that there exist no joint probability measure $\mu \in \Delta(\Phi_1 \times \Phi_2)$ over $\mathcal{F}_1 \otimes \mathcal{F}_2$ and no equilibrium strategies $\lambda \equiv (\lambda^1, \lambda^2)$ for the agents that deliver a payoff of 10 to P2. In particular, we do not require that μ be a product measure. In other words, we allow the principals to coordinate their choice of a mechanism through arbitrary correlation devices. The proof is by contradiction, and consists of five 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 μ -positive probability, P2 would incur a loss ζ , and his overall payoff would be strictly less than 10, a contradiction. The above property implies that, 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 We now prove that, for μ -almost every ϕ , $\alpha_L^\phi = \alpha_H^\phi = \frac{2}{3}$. Notice first that, as A1 does not know which state prevails, it must be that, given A1's state-independent behavior strategy $\lambda^1(\phi)$, the state-dependent outcomes $z^\phi(\omega_L)$ and $z^\phi(\omega_H)$ are induced by A2's state-dependent behavior strategies $\lambda^2(\phi)(\cdot | \omega_L)$ and $\lambda^2(\phi)(\cdot | \omega_H)$. Then, for type ω_L of A2 to induce $z^\phi(\omega_L)$ instead of $z^\phi(\omega_H)$, it must be that

$$3\alpha_L^\phi + 7.5(1 - \alpha_L^\phi) \geq 5\alpha_H^\phi + 3.5(1 - \alpha_H^\phi). \quad (26)$$

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

$$9\alpha_H^\phi + 5(1 - \alpha_H^\phi) \geq 7\alpha_L^\phi + 9(1 - \alpha_L^\phi). \quad (27)$$

Summing (26)–(27) yields $\alpha_L^\phi \leq \alpha_H^\phi$, and reinserting this inequality in (26)–(27), we obtain

$$\alpha_L^\phi \leq \frac{2}{3} \leq \alpha_H^\phi. \quad (28)$$

Now, consider the alternative behavior strategy for A2 obtained from his state-dependent candidate equilibrium behavior strategies $\lambda^2(\phi)(\cdot | \omega_L)$ and $\lambda^2(\phi)(\cdot | \omega_H)$ by decorrelating the two principals' decisions. Formally, this amounts for A2 to independently drawing two message profiles $m^2 \equiv (m_1^2, m_2^2)$ and $\hat{m}^2 \equiv (\hat{m}_1^2, \hat{m}_2^2)$ from $\lambda^2(\phi)(\cdot | \omega_H)$ and $\lambda^2(\phi)(\cdot | \omega_L)$, respectively, and then sending m_1^2 to P1 and \hat{m}_2^2 to P2, thus using the distribution $\lambda^2(\phi)(\cdot | \omega_H)$ to determine his message to P1 and the distribution $\lambda^2(\phi)(\cdot | \omega_L)$ to determine his message to P2. Given A1's behavior strategy $\lambda^1(\phi)$, this alternative strategy induces a distribution \Pr over $(x_{11}, x_{12}, x_{21}, x_{22})$ with the following marginals:

$$\Pr(x_{11}, x_{21}) + \Pr(x_{11}, x_{22}) = 1 - \alpha_H^\phi,$$

$$\Pr(x_{12}, x_{21}) + \Pr(x_{12}, x_{22}) = \alpha_H^\phi,$$

$$\Pr(x_{11}, x_{21}) + \Pr(x_{12}, x_{21}) = \alpha_L^\phi,$$

$$\Pr(x_{11}, x_{22}) + \Pr(x_{12}, x_{22}) = 1 - \alpha_L^\phi.$$

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

$1 - \alpha_H^\phi - p$, and $\Pr(x_{12}, x_{22}) = p + \alpha_H^\phi - \alpha_L^\phi$. Now, if type ω_L of A2 were to play in this way, thus sending the messages m_1^2 and \hat{m}_2^2 according to the strategy described above, he would obtain an expected payoff of

$$p + 5(\alpha_L^\phi - p) + 3.5(1 - \alpha_H^\phi - p) + 8(p + \alpha_H^\phi - \alpha_L^\phi) = 3.5 + 0.5p + 4.5\alpha_H^\phi - 3\alpha_L^\phi.$$

Because this payoff must at most be equal to his equilibrium payoff of $3\alpha_L^\phi + 7.5(1 - \alpha_L^\phi)$ and $p \geq 0$, it follows that $4 \geq 4.5\alpha_H^\phi + 1.5\alpha_L^\phi$. Combining this inequality with (27), we obtain $\alpha_L^\phi \geq \alpha_H^\phi$ and hence $\alpha_L^\phi = \alpha_H^\phi = \frac{2}{3}$ by (28), as desired. As a result, in μ -almost every subgame ϕ , type ω_L of A2 obtains a payoff of 4.5.

Step 3 Now, fixing a subgame ϕ such that $\alpha_L^\phi = \alpha_H^\phi = \frac{2}{3}$, consider the alternative behavior strategy for A2 obtained by decorrelating the two principals' decisions, but this time using only the candidate equilibrium behavior strategy $\lambda^2(\phi)(\cdot | \omega_H)$. Formally, this amounts for A2 to independently drawing two message profiles $m^2 \equiv (m_1^2, m_2^2)$ and $\hat{m}^2 \equiv (\hat{m}_1^2, \hat{m}_2^2)$ from $\lambda^2(\phi)(\cdot | \omega_H)$ and then sending m_1^2 to P1 and \hat{m}_2^2 to P2, thus using the first draw to determine his message to P1 and the second draw to determine his message to P2. Given A1's behavior strategy $\lambda^1(\phi)$, this alternative strategy induces a distribution $\tilde{\Pr}$ over $(x_{11}, x_{12}, x_{21}, x_{22})$ with the same marginals as under the original strategy,

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

It is easy to check that this system too has not full rank, and admits a continuum of solutions indexed by $p \equiv \tilde{\Pr}(x_{11}, x_{21}) = \tilde{\Pr}(x_{12}, x_{22})$, which allows us to write $\tilde{\Pr}(x_{11}, x_{22}) = \frac{1}{3} - p$ and $\tilde{\Pr}(x_{12}, x_{21}) = \frac{2}{3} - p$. Now, if type ω_L of A2 were to play in this way, thus sending the messages m_1^2 and \hat{m}_2^2 according to the strategy described above, he would obtain an expected payoff of

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

Because this payoff must at most be equal to his equilibrium payoff of 4.5 and $p \geq 0$, it follows that $p = 0$. This implies that, for $\lambda^2(\phi)(\cdot | \omega_H) \otimes \lambda^2(\phi)(\cdot | \omega_H)$ -almost every (m^2, \hat{m}^2) , we have

$$(\phi_1(m_1^1, m_1^2), \phi_2(m_2^1, \hat{m}_2^2)) \in \{\delta_{(x_{11}, x_{22})}, \delta_{(x_{12}, x_{21})}\} \quad (29)$$

for $\lambda^1(\phi)$ -almost every m^1 . But, according to Step 1, for $\lambda^2(\phi)(\cdot|\omega_H) \otimes \lambda^2(\phi)(\cdot|\omega_H)$ -almost every (m^2, \hat{m}^2) , we have

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

for $\lambda^1(\phi)$ -almost every m_1 . Thus (29) implies that for $\lambda^2(\phi)(\cdot|\omega_H) \otimes \lambda^2(\phi)(\cdot|\omega_H)$ -almost every (m^2, \hat{m}^2) , we have

$$(\phi_1(m_1^1, m_1^2), \phi_2(m_2^1, m_2^2)) = (\phi_1(m_1^1, \hat{m}_1^2), \phi_2(m_2^1, \hat{m}_2^2)) \quad (30)$$

for $\lambda^1(\phi)$ -almost every m_1 . Because ϕ_1 and ϕ_2 are measurable, we can then conclude from Fubini's theorem (Aliprantis and Border (2006, Theorem 11.27)) that (30) holds for $\lambda^1(\phi) \otimes \lambda^2(\phi)(\cdot|\omega_H) \otimes \lambda^2(\phi)(\cdot|\omega_H)$ -almost every (m^1, m^2, \hat{m}^2) . Applying again Fubini's theorem, we obtain that for $\lambda^1(\phi)$ -almost every m_1 , (30) holds for $\lambda^2(\phi)(\cdot|\omega_H) \otimes \lambda^2(\phi)(\cdot|\omega_H)$ -almost every (m_2, \hat{m}_2) , so that the mapping $(m_2^2, \hat{m}_2^2) \mapsto (\phi_1(m_1^1, m_1^2), \phi_2(m_2^1, m_2^2))$ is constant over a set of $\lambda^2(\phi)(\cdot|\omega_H)$ -measure 1.

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

Proof of Lemma 5. Because A3's payoff is constant over $X \times \Omega$ and A1's and A2's payoff functions are identical, we can focus on A1's incentives. Suppose that, in the subgame (ϕ_1^r, γ_2) , A2 and A3 play behavior strategies β^2 and β^3 that prescribe the same play for any signals s_2^2 and s_2^3 they may receive from P2, respectively; that is, for each $\omega^2 \in \Omega^2$, $\beta^2(\cdot|s_2^2, \omega^2)$ is independent of s_2^2 , and similarly $\beta^3(\cdot|s_2^3)$ is independent of s_2^3 . Then, because every signal A1 receives from P2 is uninformative, A1 may as well best respond by playing a behavior strategy β^1 that prescribes the same play for any signal s_2^1 he may receive from P2; that is, for each $\omega^1 \in \Omega^1$, $\beta^1(\cdot|s_2^1)$ is independent of s_2^1 . Because all the message spaces M_j^i are finite, this implies that the subgame (ϕ_1^r, γ_2) admits a BNE in which all agents play behavior strategies that prescribe the same play for any signals they may receive from P2. According to (4), any such BNE of the subgame (ϕ_1^r, γ_2) can be straightforwardly turned into a BNE of the subgame $(\phi_1^r, \hat{\phi}_2)$ in which P1 posts the recommendation mechanism ϕ_1^r

and P2 posts the standard mechanism $\hat{\phi}_2$ defined by

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

for all $m_2 \in M_2$ and $x_2 \in X_2$. Notice that, by construction, the same outcome is implemented in either case. Conversely, any BNE of the subgame $(\phi_1^r, \hat{\phi}_2)$ can be straightforwardly turned into a BNE of the subgame (ϕ_1^r, γ_2) in which all agents play behavior strategies that prescribe the same play for any signals they may receive from P2, and which implements the same outcome. To conclude, observe that, as $\hat{\phi}_2$ is a standard mechanism, we know from the proof of Lemma 1 that the subgame $(\phi_1^r, \hat{\phi}_2)$ admits a BNE in which P2 obtains a payoff of 5. The result follows. ■

References

- [1] Akbarpour, M., and S. Li (2020): “Credible Auctions: A Trilemma,” *Econometrica*, 88(2), 425–467.
- [2] Aliprantis, C.D., and K.C. Border (2006): *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. Berlin, Heidelberg, New York: Springer.
- [3] Attar, A., E. Campioni, T. Mariotti, and G. Piaser (2021): “Competing Mechanisms and Folk Theorems: Two Examples,” *Games and Economic Behavior*, 125, 79–93.
- [4] Attar, A., E. Campioni, and G. Piaser (2013): “Two-Sided Communication in Competing Mechanism Games,” *Journal of Mathematical Economics*, 49(1), 62–70.
- [5] Attar, A., E. Campioni, and G. Piaser (2019): “Private Communication in Competing Mechanism Games,” *Journal of Economic Theory*, 183, 258–283.
- [6] Dequiedt, V., and D. Martimort (2015): “Vertical Contracting with Informational Opportunism,” *American Economic Review*, 105(7), 2141–2182.
- [7] Eeckhout, J., and P. Kircher (2010): “Sorting and Decentralized Price Competition,” *Econometrica*, 78(2), 539–574.
- [8] Epstein, L.G., and M. Peters (1999): “A Revelation Principle for Competing Mechanisms,” *Journal of Economic Theory*, 88(1), 119–160.
- [9] Hart, O., and J. Tirole (1990): “Vertical Integration and Market Foreclosure,” *Brookings Papers on Economic Activity: Microeconomics*, 1990, 205–286.
- [10] Hurwicz, L. (1973): “The Design of Mechanisms for Resource Allocation,” *American Economic Review*, 63(2), 1–30.
- [11] Kalai A., E. Kalai , E. Lehrer, and D. Samet (2010): “A Commitment Folk Theorem,” *Games and Economic Behavior*, 69(1), 127–137.
- [12] Laffont, J.-J., and D. Martimort (1997): “Collusion under Asymmetric Information,” *Econometrica*, 65(4), 875–911.
- [13] Martimort, D. and L. Stole (2002): “The Revelation and Delegation Principles in Common Agency Games,” *Econometrica*, 70(4), 1659–1673.
- [14] McAfee, R.P. (1993): “Mechanism Design by Competing Sellers,” *Econometrica*, 61(6), 1281–1312.

- [15] McAfee, R.P., and M. Schwartz (1994): “Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity,” *American Economic Review*, 84(1), 210–230.
- [16] Milgrom, P.R., and R.J. Weber (1985): “Distributional Strategies for Games with Incomplete Information,” *Mathematics of Operations Research*, 10(4), 619–631.
- [17] Moen, E.R. (1997): “Competitive Search Equilibrium,” *Journal of Political Economy*, 105(2), 385–411.
- [18] Myerson, R.B. (1979): “Incentive Compatibility and the Bargaining Problem,” *Econometrica*, 47(1), 61–73.
- [19] Myerson, R.B. (1982): “Optimal Coordination Mechanisms in Generalized Principal-Agent Problems,” *Journal of Mathematical Economics*, 10(1), 67–81.
- [20] Myerson, R.B. (1989): “Mechanism Design,” in *The New Palgrave: Allocation, Information, and Markets*, ed. by J. Eatwell, M. Milgate, and P. Newman. New York: Norton, 191–206.
- [21] von Negenborn, C., and M. Pollrich (2020): “Sweet Lemons: Mitigating Collusion in Organizations,” *Journal of Economic Theory*, 189(105074), 1–25.
- [22] Pavan, A. and G. Calzolari (2009): “Sequential Contracting with Multiple Principals,” *Journal of Economic Theory*, 144(2), 503–531.
- [23] Pavan, A. and G. Calzolari (2010): “Truthful Revelation Mechanisms for Simultaneous Common Agency Games,” *American Economic Journal: Microeconomics*, 2(2), 132–190.
- [24] Peck, J. (1997): “A Note on Competing Mechanisms and the Revelation Principle,” Mimeo, Ohio State University.
- [25] Peters, M. (1997): “A Competitive Distribution of Auctions,” *Review of Economic Studies*, 64(1), 97–123.
- [26] Peters, M. (2001): “Common Agency and the Revelation Principle,” *Econometrica*, 69(5), 1349–1372.
- [27] Peters, M. (2014): “Competing Mechanisms,” *Canadian Journal of Economics*, 47(2), 373–397.

- [28] Peters, M. (2015): “Reciprocal Contracting,” *Journal of Economic Theory*, 158, 102–126.
- [29] Peters, M., and S. Severinov (1997): “Competition among Sellers who Offer Auctions instead of Prices,” *Journal of Economic Theory*, 75(1), 141–179.
- [30] Peters, M., and B. Szentes (2012): “Definable and Contractible Contracts,” *Econometrica*, 80(1), 363–411.
- [31] Peters, M., and C. Troncoso-Valverde (2013): “A Folk Theorem for Competing Mechanisms,” *Journal of Economic Theory*, 148(3), 953–973.
- [32] Prat, A., and A. Rustichini (2003): “Games Played Through Agents,” *Econometrica*, 71(4), 989–1026.
- Rao, B.V. (1971): “Borel Structures for Function Spaces,” *Colloquium Mathematicum*, 23(1), 33–38.
- [33] Rahman, D., and I. Obara (2010): “Mediated Partnerships,” *Econometrica*, 78(1), 285–308.
- [34] Rey, P., and J. Tirole (1986): “The Logic of Vertical Restraints,” *American Economic Review*, 76(5), 921–939.
- [35] Segal, I. (1999): “Contracting with Externalities,” *Quarterly Journal of Economics*, 114(2), 337–388.
- [36] Segal, I., and M.D. Whinston, M. D. (2003): “Robust Predictions for Bilateral Contracting with Externalities,” *Econometrica*, 71(3), 757–791.
- [37] Szentes, B. (2015): “Contractible Contracts in Common Agency Problems,” *Review of Economic Studies*, 82(1), 391–422.
- [38] Virág, G. (2010): “Competing Auctions: Finite Markets and Convergence,” *Theoretical Economics*, 5(2), 241–274.
- [39] Wright, R., P. Kircher, B. Julien, B., V. Guerrieri (2021): “Directed Search and Competitive Search Equilibrium: A Guided Tour,” *Journal of Economic Literature*, 59(1), 90–148.
- [40] Xiong, S. (2013): “A Folk Theorem for Contract Games with Multiple Principals and Agents,” Mimeo, Department of Economics, Rice University.

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