Keeping the Agents in the Dark: Competing Mechanisms, Private Disclosures, and the Revelation Principle*

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

June 28, 2025

Abstract

We study the design of market information in competing-mechanism games. We identify a new dimension, private disclosures, whereby the principals asymmetrically inform the agents of how their mechanisms operate. We show that private disclosures have two important effects. First, they can raise a principal's payoff guarantee against her competitors' threats. Second, they can support equilibrium outcomes and payoffs that cannot be supported with standard mechanisms. These results call for a novel approach to competing mechanisms, which we develop to identify a canonical game and a canonical class of equilibria, thereby establishing a new revelation principle for this class of environments.

Keywords: Incomplete Information, Competing Mechanisms, Private Disclosures, Revelation Principle.
JEL Classification: D82.

*This paper supersedes a previous version circulated under the title "Keeping the Agents in the Dark: Private Disclosures in Competing Mechanisms." We thank Laura Doval, Piotr Dworczak, 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, Nuffield College, Séminaire Roy, Seminars in Economic Theory, TSE, UCLA, UC Riverside, Università di Bergamo, Universität Bonn, Université de Montréal, University of British Columbia, University of Chicago, University of Edinburgh, 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 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 and 2024 Conferences on Mechanism and Institution Design, the 2022 European Summer Symposium in Economic Theory, the 2022 Stony Brook International Conference on Game Theory, the 2022 Università di Bergamo Industrial Organization Winter Symposium, the 2023 Venice Workshop on the Design of Strategic Interactions, the 2025 Conference on Economic Design, and the 2025 SAET Conference for many useful discussions. This research has benefited from financial support of the ANR (Programme d'Investissements d'Avenir CHESS ANR-17-EURE-0010), the MIUR (PRIN 2015, 2022), the NSF (SES 2315652), 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 for the allocation 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 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 in the hands of several principals who non-cooperatively design mechanisms to deal with several privately-informed agents. These games have been used to study several market environments, including competing auctions, competitive search, and competing vertical structures.

We depart from classical mechanism-design theory in letting each principal inform the agents asymmetrically about her effective decision rule, namely, the mapping describing how the principal's decision depends on the messages she 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. In the resulting competing-mechanism game, 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 sends to the agents and the messages she receives from them into a (possibly random) decision. In practice, private disclosures may correspond to information disclosed asymmetrically by an auctioneer to the bidders about the reserve price or other elements of her auction, or contract details disclosed asymmetrically by a manufacturer to retailers specifying how her supply responds to the retailers' market information.

We show that allowing for such private communication from the principals to the agents can significantly affect the set of equilibrium allocations, even in pure incomplete-information environments in which the agents take no payoff-relevant actions. The general lesson is that the restriction to standard mechanisms is unwarranted in competitive settings. This calls for a novel approach, which we develop in this paper.

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

First, we show that private disclosures may permit the principals to raise their payoff guarantees, that is, to increase their minimum equilibrium payoff relative to what they can achieve with standard mechanisms. A direct implication is that 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 the message spaces are sufficiently rich for the principals to post *recommendation* mechanisms, whereby each agent can recommend a direct mechanism and report his type to each principal. In line with Yamashita (2010), we show that, without private disclosures, a version of the folk theorem holds: any feasible payoff vector yielding each principal a payoff above an appropriate min-max-min bound can be supported in equilibrium using standard mechanisms.[1] The example is a zero-sum game between two principals in which the min-max-min payoff for one of the principals is her lowest feasible payoff. We show that, with private disclosures, this principal can guarantee herself a payoff strictly above her min-max-min bound, regardless of the mechanism posted by the other principal and of the continuation equilibrium played by the agents. Indeed, by privately informing one of the agents of her effective decision rule while keeping the others in the dark, the principal perfectly aligns this agent's preferences with hers, making the agent an ally who is no longer willing to participate in the collective behavior necessary to deliver the principal's min-max-min payoff. The upshot of this example is that equilibrium outcomes and payoffs of standard competing-mechanism games—in particular those supported by recommendation mechanisms à la Yamashita (2010)—need not be robust, in the sense that they fail to be supportable once the principals can engage in private disclosures.

Next, we show that private disclosures may enable the principals to achieve equilibrium outcomes and payoffs that cannot be supported with standard mechanisms, no matter how rich the message spaces are. We provide an example of a competing-mechanism game in which two principals seek to perfectly coordinate their decisions with the agents' types. Achieving the desired correlation while respecting the agents' incentives requires that (a) 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 (b) such information not create common knowledge among the agents about the first principal's decision before they communicate with the second principal. The example illustrates the possibility to achieve both (a) and (b) with private disclosures and the necessity of both (a) and (b) when it comes to supporting the

[1]The min-max-min payoff of each principal is with respect to the other principals' mechanisms, her own mechanism, and the agents' continuation equilibrium.

desired correlation. In this example, private disclosures play the role of encrypted messages that, in isolation, are uninformative about a principal's decision, but, when combined with the signals disclosed to other agents, perfectly reveal it. The upshot is that standard mechanisms, and in particular the *universal mechanisms* of Epstein and Peters (1999), fail to support all equilibrium outcomes once private disclosures are accounted for.

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 findings are not theoretical curiosities: they have implications for how firms compete in markets. For example, the 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.[2] The second result, in turn, suggests that manufacturers may collude more effectively by asymmetrically informing common retailers of how their production responds to the retailers' private information about market conditions. This is because private disclosures allow the manufacturers to relax the retailers' incentive-compatibility constraints, thereby facilitating collusion among the manufacturers without resorting to illegal explicit agreements.

The possibility for the principals to design the agents' market information brings a new angle to mechanism-design theory and calls for a novel approach to the study of competing-mechanism games, which we develop in the second part of the paper. Our main contribution is to provide a new revelation principle for these games. This requires, in particular, to identify a canonical game, a corresponding protocol of communication between the parties, and a canonical class of equilibria.

We establish the result in two steps.

We first focus on games with a single round of disclosures followed by a single round of messages and identify the signal and message spaces that enable the players to convey all the information relevant for equilibrium allocations. Theorem 1 shows that any equilibrium outcome of any competing-mechanism game with private disclosures and rich signal and message spaces is also an equilibrium outcome of a game in which each principal asks each agent to report his *extended type*, which comprises his exogenous type and the signals received from the other principals. In equilibrium, principals play pure strategies, and agents truthfully report their extended types. The reason why, with private disclosures, attention can be restricted to equilibria in which the principals do not mix over their mechanisms is that any correlation generated by the agents using the realizations of the principals' mixed strategies as a correlation device for their reports can be replicated by the principals using

[2] In 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).

signals to correlate these reports. Similarly, any mixing by the agents over the messages sent to the principals on path can be replicated by the principals using the signals as a jointly-controlled lottery that none of them can individually manipulate (Aumann and Maschler (1995)). Finally, the reason why the agents' message spaces can be taken to coincide with their extended type spaces is that the latter contain all the information necessary to determine the principals' decisions.

We next examine whether the short-communication protocol in the game studied in Theorem 1, whereby the principals and the agents communicate only once using rich signal and message spaces, is itself without loss of generality. Theorem 2 shows that allowing for multiple rounds of disclosures and messages does not affect the set of equilibrium outcomes. That is, any equilibrium outcome of any long-communication game in which the principals gradually disclose elements of their mechanisms to the agents and repeatedly solicit private information from them is also an equilibrium outcome of the short-communication game. Conversely, any equilibrium outcome of the short-communication game is also an equilibrium outcome of any long-communication game. Thus, equilibria of the short-communication game with rich signal and message spaces are both *universal* and *robust*: gradual resolution of uncertainty is irrelevant for equilibrium outcomes.

The reason why attention can be confined to short communication is the following. In a long-communication game, each principal can correlate her decision with those of the other principals and the agents' types by gradually providing information to the agents about her decision rule in response to the information the agents provide about the other principals' decision rules. This correlation is key to implement, on path, the desired outcomes, and, off path, the necessary punishments for the deviators. In the short-communication game, such a correlation is replicated by (a) asking the agents to report their *message plans* for the long-communication game, and (b) generating an endogenous non-manipulable and non-informative common sampling variable that emulates the endogenous correlation in the original signals. The message plans describe the (possibly correlated) messages that each agent would have sent to the various principals in all rounds, as a function of the history of received signals. The auxiliary sampling variable is constructed by each principal drawing for each agent an auxiliary signal. These auxiliary signals are then suitably aggregated (across agents and principals) to generate the common auxiliary sampling variable, in a way that makes it impossible for any of the principals to manipulate its distribution. Furthermore, none of the agents can learn about the realization of this variable from his individual draws, and hence no agent can benefit from misreporting to any principal the new information received—here the assumption that there are two or more agents is crucial. We show how to encode each agent's message plan into a unidimensional statistic, and how to combine the message plans with the new variable to replicate the entire sequence of signals and messages in the long-communication game. The construction is used both on path to support the original equilibrium allocations and, off path, to inflict to any deviating principal the same punishment as in the long-communication game.

Perhaps surprisingly, the same construction also permits us to establish the robustness of the equilibrium allocations of the short-communication game. The difficulty stems from the need to identify which short-communication mechanism is "strategically equivalent" to each long-communication one. The identification is done by requiring equivalence between the decisions induced by the two mechanisms by varying the agents' behavior. Each agent then uses the messages he would have sent in the short-communication game to identify his message plan for the long-communication game. When the agents follow such plans, no principal can benefit from deviating to a long-communication mechanism.

Together, Theorems 1–2 identify a canonical game, a corresponding communication protocol, and a canonical class of equilibria, which allow to support any equilibrium outcome of any competing-mechanism game with arbitrarily long communication. In the canonical game, which entails short communication, the principals post extended-direct mechanisms that extract signals from the unit interval and ask each agent to report his extended type. In a canonical equilibrium, principals do not randomize over mechanisms, and agents truthfully report their extended types on path. These features point to an interpretation in line with that offered by the revelation principle in classical mechanism-design theory.

Related Literature This paper contributes to the theoretical foundations of competing-mechanism games. In a seminal work, McAfee (1993) shows that the equilibria of these games may require that the agents report to each principal both their exogenous types and their endogenous market information—that is, the mechanisms posted by the other principals. To address the resulting infinite-regress problem, Epstein and Peters (1999) construct a space of universal mechanisms and establish a revelation principle: any equilibrium outcome of any competing-mechanism game can be supported as an equilibrium outcome of the game in which the principals are restricted to posting universal mechanisms. Subsequent work focuses on providing explicit characterizations of the equilibrium outcomes of these games. In particular, Yamashita (2010) shows that, with three or more agents, every deterministic incentive-compatible allocation yielding each principal a payoff above some min-max-min bound can be supported in equilibrium. Crucially, these two results are established under the assumption that principals post standard mechanisms, which, despite their sophistication, both universal and recommendation mechanisms are instances of. By contrast, we allow for private disclosures and for gradual resolution of uncertainty, that is, long communication.

Along the way, we show that several other restrictions in these papers can be dispensed with—namely, to pure strategies, exclusive participation (Epstein and Peters (1999)), and three or more agents (Yamashita (2010)).

In a different context, Peters and Troncoso-Valverde (2013) show 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 noticeable feature of their 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 usual distinction between principals and agents and the usual restriction that the agents do not communicate among themselves and release no information before observing the mechanisms posted by the principals.

Private communication from the principals has received little attention in the literature and, when considered, has typically been confined to action recommendations, in line with the classical mechanism-design approach (Myerson (1982)). Attar, Campioni, and Piaser (2019) show that, in complete-information games, equilibrium allocations supported by standard mechanisms fail to be robust against a deviation to a mechanism with private recommendations. In a setting closer to that of the present paper, Attar, Campioni, and Piaser (2013) show that the principals can correlate their equilibrium decisions with the agents' types by recommending to the agents which reports to make in their competitors' (direct) mechanisms. However, in their example, the same correlation can also be obtained by letting the agents randomize over the messages they send to the principals. In both papers, private signals play a role similar to the one they play in single-principal settings. By contrast, we uncover two novel roles for private disclosures: raising the principals' individual payoff guarantees, and correlating their decisions with the agents' types in ways that cannot be achieved through standard mechanisms.

Private disclosures generate endogenous asymmetric information among the agents about a principal's effective decision rule. A similar role is played by the market information privately held by the agents when contracting is bilateral (see, for instance, Segal and Whinston (2003)). The literature on bilateral contracting, however, focuses on situations in which a single principal contracts with multiple agents but cannot commit to a public mechanism specifying how her decision responds to the messages she receives from them

(Rey and Tirole (1986), Hart and Tirole (1990), McAfee and Schwartz (1994), Segal (1999), Dequiedt and Martimort (2015), Akbarpour and Li (2020), Banchio, Skrzypacz, and Yang (2024)). We let the principals fully commit to their mechanisms and strategically choose to disclose private information to the agents in order to discipline how the agents behave with their competitors. Our approach suggests that private (or secret) contracting need not be the result of high transaction costs of processing information or of limits imposed to multilateral contracting, but rather of the optimal choice of a mechanism designer in a competitive environment.

In our setting, the principals cannot directly condition their decisions on other principals' 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 the principals can make commitments contingent on each other's decisions and/or mechanisms. The conclusions of our second example also extend to the case where the principals can condition their mechanisms on their competitors' mechanisms as in Peters' (2015) model of reciprocal contracting, suggesting that private disclosures may substitute for direct private communication between the principals.

In the common-agency case, where several principals contract with a single agent, the menu theorems of Peters (2001), Martimort and Stole (2002), and Pavan and Calzolari (2009, 2010) ensure that any equilibrium outcome of any game in which the principals compete by posting arbitrary message-contingent decision rules can be reproduced in a game in which the principals post menus of (possibly random) decisions and delegate to the agent the choice of the final allocation. In such settings, there is no role for private disclosures. However, certain outcomes and payoffs can be supported in equilibrium only by the principals disclosing information about their menus gradually (see Attar, Campioni, Mariotti, Pavan, and Renault (2025) for an example).

The paper is organized as follows. Section 2 introduces a general competing-mechanism game with private disclosures. Section 3 contains two examples that jointly show that the sets of equilibrium outcomes in games with and without private disclosures are not nested. Section 4 illustrates a new revelation principle for competing-mechanism games. In particular, Sections 4.1 and 4.2 focus on short-communication games and identify the sets of agents' messages and principals' signals that are rich enough to convey all their available information. Sections 4.3 and 4.4 consider multiple rounds of communication, and provide a canonical game together with a canonical class of equilibria. Section 5 concludes. The Online Supplement collects detailed proofs of the results.

2 The Model

Players We consider a setting in which several principals, indexed by $j = 1, \ldots, J$, contract with several agents, indexed by $i = 1, \ldots, I$, where $I \geq 2$ and $J \geq 2$. Throughout the paper, we use subscripts to refer to the principals and superscripts to refer to the 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 $\omega \equiv (\omega^i)_{i=1}^I$ is $\Omega \equiv \times_{i=1}^I \Omega^i$. Principals and agents commonly believe that the state ω is drawn from Ω according to the distribution $\mathbf{P}$.

Decisions and Payoffs Every principal j (she) takes a decision x_j in some finite set X_j. We let $v_j : X \times \Omega \to \mathbb{R}$ and $u^i : X \times \Omega \to \mathbb{R}$ be the vNM utility functions of principal j and of agent i, respectively, where $X \equiv \times_{j=1}^J X_j$ is the set of profiles of possible decisions for the principals. Agents take no payoff-relevant actions, making our setting one of pure incomplete information. We refer to $G \equiv (\Omega, \mathbf{P}, X, (u^i)_{i=1}^I, (v_j)_{j=1}^J)$ as the *primitive game*.[3]

Allocations and Outcomes An *allocation* is a function $z : \Omega \to \Delta(X)$ assigning a lottery over the set X to every state of the world. The *outcome* induced by an allocation z is the restriction of z to the set of states occurring with positive probability under $\mathbf{P}$.[4]

Standard Mechanisms A *standard mechanism* for principal j is a decision rule $\phi_j : M_j \to \Delta(X_j)$ assigning a lottery over principal j's decisions to every profile of messages $m_j \equiv (m_j^i)_{i=1}^I \in M_j$ she may receive from the agents, where $M_j \equiv \times_{i=1}^I M_j^i$ for some collection of nonempty sets M_j^i of messages from every agent i to principal j. We 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 to permit the agent to reveal his type to her. Unless otherwise stated, we also assume that the sets M_j^i are finite for all i and j.

Mechanisms with Private Disclosures A *mechanism with private disclosures* is one in which a principal can privately inform the agents of how her decision responds to their messages. The decision rule is then indexed by a family of parameters, one for each

[3]Notice that, unlike in the exclusive-competition model of Epstein and Peters (1999), an agent's payoff in G can depend on all the principals' decisions and on the other agents' types. Hence, the model also captures markets in which payoffs are interdependent and competition is non-exclusive.

[4]The distinction between allocations and outcomes is relevant when the agents' types are correlated.

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, and are then revealed to the agents in the form of private signals. In fine, different agents may have different information about how the principal's decision responds to their messages.

Formally, a mechanism with private disclosures for principal j is thus a pair $\gamma_j \equiv (\sigma_j, \phi_j)$ such that

(i) $\sigma_j \in \Delta(S_j)$ is a probability measure over the profiles of signals $s_j \equiv (s_j^i)_{i=1}^I \in S_j$ that principal j sends to the agents, where $S_j \equiv \times_{i=1}^I S_j^i$ for some collection of nonempty sets S_j^i of signals from principal j to every agent i;

(ii) $\phi_j : S_j \times M_j \to \Delta(X_j)$ is an extended decision rule assigning a lottery over principal j's decisions to every profile of signals $s_j \in S_j$ sent to the agents and every profile of messages $m_j \in M_j$ received from them.

Unless otherwise stated, we let the sets M_j^i and the sets S_j^i be finite for all i and j. The space of mechanisms with private disclosures for 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. The assumption that signal and message spaces are finite is relaxed in Section 4.

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

Timing and Strategies Given a primitive game G, the competing-mechanism game G^{SM} with private disclosures unfolds in three stages:

1. the principals simultaneously post mechanisms and accordingly send private signals to the agents 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 the payoffs accrue.

If S_j^i is a singleton for all i and j, the principals play a competing-mechanism game without private disclosures, which we denote by G^M. Epstein and Peters (1999) and Yamashita (2010) study different versions of this game, which share the feature that principals cannot asymmetrically inform the agents about their effective decision rules.

A mixed strategy for principal j in G^{SM} is a Borel probability measure $\mu_j \in \Delta(\Gamma_j)$. A strategy for agent i in G^{SM} is a Borel-measurable function $\lambda^i : \Gamma \times S^i \times \Omega^i \to \Delta(M^i)$ that assigns a lottery over agent i's messages $m^i \equiv (m^i_j)_{j=1}^J \in M^i \equiv \times_{j=1}^J M^i_j$ to the principals to every profile of mechanisms $\gamma \equiv (\gamma_j)_{j=1}^J \in \Gamma \equiv \times_{j=1}^J \Gamma_j$, signals $s^i \equiv (s^i_j)_{j=1}^J \in S^i \equiv \times_{j=1}^J S^i_j$, and type $\omega^i \in \Omega^i$ of agent i. The allocation $z_{\mu,\lambda} : \Omega \to \Delta(X)$ induced by the strategies $(\mu, \lambda) \equiv ((\mu_j)_{j=1}^J, (\lambda^i)_{i=1}^I)$ is then defined by

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

for all $(\omega, x) \in \Omega \times X$, where $S \equiv \times_{j=1}^J S_j$ and $M \equiv \times_{j=1}^J M_j$. For every profile of mechanisms $\gamma \in \Gamma$, a behavior strategy for agent i in the subgame γ of G^{SM} played by the agents is a function $\beta^i : S^i \times \Omega^i \to \Delta(M^i)$ assigning a lottery over the profiles of messages $m^i \in M^i$ to every profile of signals $s^i \in S^i$ and every type $\omega^i \in \Omega^i$. We let $z_{\gamma,\beta}$ be the allocation induced by the profile of behavior strategies $\beta \equiv (\beta^i)_{i=1}^I$ in the subgame γ; $z_{\gamma,\beta}$ 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 To any profile of mechanisms $\gamma \in \Gamma$ of G^{SM} corresponds a Bayesian game played by the agents, with type space $S^i \times \Omega^i$ and action space M^i for every agent i. In this subgame γ of G^{SM}, the agents' beliefs are pinned down by the prior distribution $\mathbf{P}$ over Ω and the distributions $(\sigma_j)_{j=1}^J$ from which the signals are drawn.

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

(i) for each $\gamma \in \Gamma$, $\lambda^*(\gamma) \equiv (\lambda^{*i}(\gamma))_{i=1}^I$ is a Bayes–Nash equilibrium (BNE) of the subgame γ played by the agents;

(ii) given the strategies λ^* for the agents, μ^* is a Nash equilibrium (NE) of the game played by the principals.

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

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

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 Two Key Roles of Private Disclosures

This section motivates the need to account for private disclosures when modeling competing mechanisms. We do so by means of two examples, each of them casted in a prominent strategic setting.

The first example considers a zero-sum game between two principals. In this context, we show how a principal may guarantee herself a higher payoff at the expense of her competitor by posting mechanisms with private disclosures rather than standard mechanisms. As a consequence, the 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 the principals post standard mechanisms with rich message spaces typically lend themselves to folk-theorem types of results.

The second example considers a pure coordination game between two principals to 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. The example provides a positive answer, thereby showing that the universal mechanisms of Epstein and Peters (1999) fail to support all equilibrium outcomes when the principals can engage into private disclosures.

3.1 Private Disclosures Can Raise a Principal's Payoff Guarantee

Consider a G^M game in which every message space M_j^i is sufficiently rich to enable agent i to recommend to every principal j any deterministic direct mechanism $d_j : \Omega \to X_j$, and to make a report about his type. That is, let D_j denote the finite set of all such direct mechanisms, and 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

$$\phi_j^r(m_j^1, \dots, m_j^I) \equiv \begin{cases} d_j(\omega^1, \dots, \omega^I) & \text{if } \operatorname{card}\{i : d_j^i = d_j\} \geq I - 1 \\ \overline{x}_j & \text{otherwise} \end{cases}, \tag{2}$$

where $\overline{x}_j$ 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 $(\overline{d}_j, \overline{\omega}_j^i)$ of $D_j \times \Omega^i$, once again applying (2). 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. Because $I \geq 3$, no agent is pivotal in selecting the direct mechanism. Deviations can then be punished by having the agents change the direct mechanisms recommended to the non-deviating principals.

In Example 1, we first confirm this folk-theorem result, computing the relevant min-max-min bounds. We then show that a principal, by informing the agents asymmetrically of how her decision responds to their messages, can raise her payoff guarantee above her min-max-min bound. The intuition is that private disclosures may prevent the agents from coordinating on the appropriate punishment, even if the non-deviating principals post

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

Table 1: Payoffs in state (ω_L, ω_L).

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

Table 2: Payoffs in state (ω_H, ω_H).

recommendation mechanisms. In the example, a principal randomly draws her decision and privately discloses it to one of three agents, keeping the other two in the dark. By doing so, she aligns the selected agent's preferences with her own. 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. Because, in the example, the remaining agents have neither the information nor the incentives to carry out the appropriate punishments themselves, the 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.

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: the states (ω_L, ω_L) and (ω_H, ω_H) occur with equal probability $\frac{1}{2}$.[5]

The players' payoffs in the primitive game G_1 are represented in Tables 1 and 2: in each cell, the first two numbers are P1's and P2's payoffs, and the last one denotes the common payoff to A1, A2, and A3, respectively.

A Folk Theorem in Standard Mechanisms Let G_1^M be a game in which $D_j \times \Omega^i \subset M_j^i$ for all i and j, so that P1 and P2 can post recommendation mechanisms as in Yamashita (2010). 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. Lemma 1 below characterizes the set of equilibrium payoffs for the principals in G_1^M.

Lemma 1 *Any payoff for P2 in* $[5, 5.5]$ *and, hence, any payoff for P1 in* $[-5.5, -5]$, *can be supported in a PBE of* G_1^M.

The arguments in the proof are similar to those in Yamashita (2010, Theorem 1), but account for stochastic mechanisms, which he does not consider. In particular, to keep P2 at her min-max-min payoff of 5, we exploit the agents' reports to implement a state-contingent

[5] That A1's and A2's types are perfectly correlated simplifies some of the derivations, but is not essential for our results.

punishment in P1's recommendation mechanism. The min-max-min payoff of -5.5 for P1, instead, obtains irrespective of the agents' message behavior with P2.

In a generalized version of Yamashita (2010), Peters and Troncoso-Valverde (2013) also establish a folk-theorem result. In the game they study, any outcome that satisfies players' incentive compatibility and individual rationality 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.[6]

As pointed out by Peters (2014), a principal's min-max-min payoff is typically sensitive to the set of available mechanisms. In Example 1, this issue does not arise: indeed, 5 is P2's lowest feasible payoff, and P1 can achieve the bound of -5.5 by committing to a constant decision. Szentes (2009) develops an example of a zero-sum game between two principals under complete information, and argues that recommendation mechanisms allow to support equilibrium payoffs *below* the min-max-min payoff, suggesting, as the relevant bound, the principals' max-min-min payoff. In his example, the discrepancy between these two bounds arises because, as in Yamashita (2010), he does not allow for random mechanisms. This is not the case in Example 1: allowing for stochastic mechanisms effectively guarantees that, for each principal, the min-max-min and max-min-min payoffs coincide. Hence, there is no controversy about which bound is relevant to establish Yamashita's (2010) result.

Non-Robustness to Private Disclosures We now show that in any enlarged game in which the principals can post mechanisms with private disclosures, P2 can guarantee herself a payoff strictly higher than her min-max-min payoff of 5.

Lemma 2 *Let G_1^{SM} be any game with private disclosures such that $D_j \times \Omega^i \subset M_j^i$ for all i and j. Then, G_1^{SM} admits a PBE. Moreover, if $\operatorname{card} S_2^1 \geq 2$, then P2's payoff in any PBE of G_1^{SM} is at least equal to $5 + \frac{1}{6}$.*

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 guarantees her a payoff strictly above the 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

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

informing A1 of her decision, P2 can exploit the fact that, in state (ω_L, ω_L), 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.[7] Moreover, because A3 does not observe the state, his message to P1 must be the same in each state. As a result, A2 must de facto have full control over P1's decision. However, when P2 is expected to take decision x_{21} with probability 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 states, which yields P2 a payoff strictly above 5. Hence, if P2 informs A1—and only A1—of her decision, then, no matter the mechanism posted by P1 and the agents' continuation equilibrium, her payoff is strictly higher than 5. More generally, Lemma 2 characterizes an interval of P2's equilibrium payoffs in G_1^M that cannot be supported when private disclosures are accounted for.[8]

Because A1's and A2's preferences are perfectly aligned, the reader may wonder why P2 informs these agents asymmetrically. The reason is that, if both agents have the same information about P2's decision, they can discipline each other, implementing incentive-compatible punishments for P2 as in Yamashita (2010). For example, if both 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.[9] Thus, private disclosures offer an effective tool to raise payoff guarantees in competitive settings. Lemmas 1–2 together imply the following result.

Proposition 1 *PBE outcomes of competing-mechanism games without private disclosures need not be robust to the possibility for the principals to post mechanisms with private disclosures. In particular, PBE payoff vectors of competing-mechanism games without private*

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

[8]Notice that the only role played by the finiteness of the M_j^i and the S_j^i spaces is to guarantee the existence of an equilibrium in every subgame (γ_1, γ_2). 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.

[9]Claim 1 in the Online Supplement handles the opposite polar case in which P1 posts a recommendation mechanism and P2 sends uninformative signals to the agents. We show that there exists a BNE of the agents' game in which P2 obtains her minimum feasible payoff of 5. This hinges on identifying a one-to-one correspondence between the *babbling equilibria* in which the agents ignore the signals they receive from P2, and the equilibria of a subgame in which P2 commits to a standard mechanism that averages over the same profiles of uninformative signals.

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 the principals can engage into private disclosures.

The result extends to settings in which each principal can condition her decision directly on the other principal's decisions and/or mechanisms, as in Kalai, Kalai, Lehrer, and Samet (2010), Peters and Szentes (2012), Szentes (2015), and 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 only observed by A1 and A2, P1 must ultimately let them determine which decisions to implement in response to a deviation by P2. Now, suppose that P2 posts a mechanism whereby she selects a decision at random and only informs A1 of her decision. Because, whenever P2 selects x_{22} in state (ω_L, ω_L) and x_{21} in state (ω_H, ω_H), 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. Hence, by using private disclosures, P2 can guarantee herself a payoff greater than 5 even if P1 can resort to contractible contracts or post a reciprocal mechanism.

In addition, the result in Proposition 1 remains true even if the principals and/or the agents have access to randomizing devices that can 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 the agents' messages. In fact, the result remains true even if the agents' messages can be coordinated by a mediator who first elicits information from the agents and then sends them private recommendations. This is because, given the mechanism posted by P2, there is no way for the mediator to extract from A1 information about the state and P2's decision and use the information to keep P2's payoff down to 5. Thus, the task of punishing P2 must be fully delegated to A2, which we have shown to be impossible.

3.2 Private Disclosures Can Support New Equilibrium Payoffs

Example 2 shows that 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 consider a pure coordination game between two principals, in which, in the absence of private disclosures, the agents' incentives prevent an efficient correlation of the principals' decisions in equilibrium. By using private disclosures, instead, one of the principals can make the agents' messages to his opponent depend on information that correlates with her own decision. In equilibrium, every signal sent by this principal conveys in isolation no information about her decision to

	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.

the agent who receives it; but, taken together, the signals perfectly reveal her decision. In addition, the agents truthfully report 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.

Example 2 Let $I = J \equiv 2$. 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 in the primitive game G_2 are represented in Tables 3 and 4, in which the first payoff is P1's and P2's, and the last two payoffs are A1's and A2's, respectively; $\zeta < 0$ is an arbitrary loss for the principals. In this game, achieving an expected payoff of 10 requires principals to coordinate on a different decision in each state, which only A2 privately observes. This, however, conflicts with the agents' incentives whenever private disclosures are not accounted for.

Private Disclosures as Encryption Keys To illustrate the key ideas in the simplest possible manner, we consider a game G_2^{SM} 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 types and the signals received from the other principal; that is, we let $M_1^i \equiv \Omega^i \times S_2^i$ and $M_2^i \equiv \Omega^i$ for all i.[10] The following result then holds.

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

$$z(\omega_L) \equiv \alpha\delta_{(x_{13}, x_{21})} + (1 - \alpha)\delta_{(x_{14}, x_{22})}, \tag{3}$$

$$z(\omega_H) \equiv \alpha\delta_{(x_{12}, x_{21})} + (1 - \alpha)\delta_{(x_{11}, x_{22})}, \tag{4}$$

in which the principals obtain their maximum feasible payoff of 10, *can be supported in a PBE of G_2^{SM}.*

[10]As 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.

Observe 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 of $\frac{23}{3}$ if he is of type ω_H.

In equilibrium, P2 posts a mechanism with private disclosures that selects the decision x_{21} if the signals she sends to A1 and A2 match, and the decision x_{22} otherwise. P2 chooses her joint probability distribution over profiles of signals in $S_2 = \{1, 2\} \times \{1, 2\}$ so as to keep both agents in the dark: regardless of the signal he receives, every agent's posterior belief about P2's decision coincides with his prior belief. These private disclosures can thus be interpreted as *encryption keys*: in isolation, every signal sent by P2 is completely uninformative of her decision; but, taken together, the two signals perfectly reveal her decision. P1's mechanism, in turn, is designed to elicit both the agents' information about their types and the signals received from P2, and to use this information to perfectly correlate her decision with P2's and the state of nature. Given the equilibrium mechanisms, the agents have the incentives to report truthfully to P1.

Notice that, for the principals to obtain their maximum feasible payoff of 10 while respecting 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 could be fully delegated to the agents by letting them 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. The construction crucially exploits that this information is not directly observable to the agents. In fact, 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 induce P1 to take decision x_{12} with certainty, obtaining a payoff of 5 instead of the payoff of 3 he obtains by being truthful. Because the principals cannot communicate directly with each other, the desired correlation can only be generated through private disclosures, as we show next.

Indispensability of Private Disclosures We now argue that the outcome (3)–(4) 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, no matter how rich the message spaces are. More generally, the maximal feasible payoff of 10 for the principals cannot be supported in any equilibrium of any such game. Thus, private disclosures are indispensable to support the above outcome and its associated payoffs. 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 enables us to capture the case in which every principal j's message spaces are large enough—namely, uncountable Polish spaces—to encode the agents' market information about her competitor's mechanism, as in Epstein and Peters' (1999) 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 the principals obtain their 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, by the principals randomizing over the mechanisms they post and/or the agents randomizing over the messages they send in those mechanisms. In either case, the correlation between the principals' decisions must ultimately obtain as a result of incentive-compatible choices by the agents. We establish that private disclosures are indispensable by showing that, in their absence, there is no continuation equilibrium that induces the desired correlation. The proof of this result consists of two steps.

First, 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 de-correlate 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 (3)–(4) for $\alpha = \frac{2}{3}$. Because the distribution over the principals' decisions must be the same regardless of the mechanisms they post on the equilibrium path, the principals' mixing behavior is irrelevant for inducing the desired correlation.

Second, we show that the outcome (3)–(4) for $\alpha = \frac{2}{3}$ is inconsistent with the agents' incentives. Specifically, we consider another way for A2 to de-correlate 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 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 decisions taken 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 the principals obtain their maximum feasible payoff of* 10. *In particular, there exists no PBE of G_2^M that supports the outcome (3)–(4) for $\alpha = \frac{2}{3}$.*

Lemmas 3–4 together imply the following.

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 is so even if the principals or the agents play mixed strategies in equilibrium.*[11]

Proposition 2 shows that the universal mechanisms of Epstein and Peters (1999) fail to support all equilibrium outcomes when the principals can engage into private disclosures.[12] The latter enable the principals to coordinate their responses to the information privately held by the agents while respecting the agents' incentives. In so doing, private disclosures also enable the principals to obtain payoffs that they are not able to obtain with standard mechanisms only.

The proof of Lemma 4 does not suppose that the principals' choices are independent. Hence, it is impossible to support the outcome (3)–(4) for $\alpha = \frac{2}{3}$ and the corresponding payoff of 10 for the principals in any game without private disclosures even if the principals can correlate their choices of standard mechanisms, by means, for instance, of a public randomization device. Similarly, this outcome, and the payoff vector associated with it,

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

[12] Proposition 2 also implies that Yamashita's (2010) folk theorem does not hold for stochastic allocations. Indeed, the allocation defined by (3)–(4) for $\alpha = \frac{2}{3}$ is certainly incentive-compatible; moreover, it yields the principals their maximum feasible payoff of 10, which is certainly at least equal to their 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.

cannot be supported in equilibrium even if G_2^M is enriched by allowing each principal to post mechanisms conditional on her competitor's mechanisms. In other words, the result in Proposition 2 continues to hold if one allows principals to post reciprocal mechanisms, as in Peters (2015). Finally, it should be noted that the outcome (3)–(4) for $\alpha = \frac{2}{3}$ could be supported without private disclosures if the principals had access to a direct communication channel. The role of private disclosures in Example 2 is to pass information from one principal on to the other without the agents being able to interpret it. Private disclosures are thus a substitute for direct private communication between the principals.

Together, Propositions 1 and 2 imply that the sets of equilibrium payoffs of games with and without private disclosures are not nested. This suggests that modeling competition among multiple mechanism designers requires to explicitly consider private disclosures. In the remainder of the paper, we develop a novel approach to analyze competition over mechanisms when principals design the agents' market information. Our objective is to derive a revelation principle for these settings. In this respect, identifying a language rich enough to represent the strategic role of private disclosures in any G^{SM} game is paramount but need not suffice. Indeed, when a single principal contracts with one or several agents, the standard revelation principle (Myerson (1979, 1982)) ensures that one can safely restrict attention to a single round of communication between the contracting parties. It is not a priori clear whether this also holds in games in which several principals contract with several agents and communication goes both ways. We address these issues in the next section.

4 A Generalized Revelation Principle

In this section, we lay down the methodology needed to derive a generalized revelation principle for competing-mechanism games with private disclosures. In line with the classical approach in mechanism-design theory (Myerson (1979, 1982, 1986), Forges (1986)), we fix the primitive game G, and consider the impact of alternative communication protocols on equilibrium outcomes.

To develop our construction, we first take as given the communication protocol that underlies Examples 1 and 2. That is, after mechanisms are posted, there is a single round of disclosures by the principals, followed by a single round of messages from the agents. In this context, we identify the signal and message spaces that allow to convey all relevant information. This allows us to introduce the class of extended-direct mechanisms, which permits one to support any equilibrium outcome of any game with rich signal and message spaces as a pure-strategy truthful equilibrium outcome.

We then allow the principals to engage in long communication, namely, to disclose signals to the agents and solicit messages from them over an arbitrary (possibly infinite) sequence of rounds. We establish that any equilibrium outcome of any such long-communication game is also a pure-strategy truthful equilibrium outcome of the game with short communication in which principals offer extended-direct mechanisms. We also show that the reverse is true: any equilibrium outcome of the short-communication game is robust in the sense that it is also an equilibrium outcome of any game of long communication with rich signal and message spaces.

Together, these results provide a new revelation principle for games with competing mechanism designers.

4.1 Competing Mechanisms with Rich Communication

Whereas the restriction to finite signal and message spaces enabled us to simplify the construction of equilibria in Lemmas 1–3, it cannot be maintained when searching for a general class of mechanisms. With this objective in mind, we henceforth consider games in which signal and message spaces are *rich*. Indeed, as we know from Epstein and Peters (1999), even if private disclosures are not accounted for, every agent's message spaces must typically contain a continuum of messages for him to be able to report his market information to the principals.[13] This suggests that, at the very least, any message space should embed a copy of $[0, 1]$. When considering private disclosures, the agent's behavior in a principal's mechanism responds to the signals he receives from the other principals, which further requires to enrich his space of available messages. In addition, from every principal's viewpoint, the space of available signals should be rich enough to let her asymmetrically inform the agents of how her decision responds to the agents' messages (as in Example 1), and to encrypt information and pass it on to the other principals to correlate their decisions with hers and the agents' types while respecting the agents' incentives (as in Example 2).

We now formally extend the model of Section 2 to allow for such rich communication.

Definition 1 *Given a primitive game* G, $G^{\hat{S}\hat{M}}$ *is the game with private disclosures in which, for all* i *and* j, $\hat{S}^i_j = \hat{M}^i_j \equiv [0, 1]$ *and* $\hat{\Gamma}_j$ *is the corresponding space of mechanisms* $\hat{\gamma}_j \equiv (\hat{\sigma}_j, \hat{\phi}_j)$ *such that* $\hat{\sigma}_j \in \Delta(\hat{S}_j)$ *is a Borel probability measure and* $\hat{\phi}_j : \hat{S}_j \times \hat{M}_j \to \Delta(X_j)$ *is a Borel-measurable function.*

A pure strategy for principal j in $G^{\hat{S}\hat{M}}$ is simply an element of $\hat{\Gamma}_j$. Following Aumann (1964), we define mixed strategies for principal j using an exogenous randomizing device,

[13] Even if principal j's competitors are restricted to post mechanisms with finitely many messages, from principal j's viewpoint, the agents' relevant market information belongs to $\times_{k \neq j} \Delta(X_k)^{M_k}$, a space that is Borel-isomorphic to $[0, 1]$.

modeled as a sampling space $\Xi_j \equiv [0,1]$, endowed with its Borel σ-field $\mathcal{B}([0,1])$ and Lebesgue measure $\mathrm{d}\xi_j$. A mixed strategy for principal j in $G^{\hat{S}\hat{M}}$ is a mapping assigning a mechanism in $\hat{\Gamma}_j$ to any realization of the sampling variable ξ_j. Formally, it is described by a pair $\hat{\mu}_j \equiv (\hat{\mathfrak{s}}_j, \hat{\mathfrak{f}}_j)$ of Borel-measurable functions $\hat{\mathfrak{s}}_j : \Xi_j \to \Delta(\hat{S}^i_j)$ and $\hat{\mathfrak{f}}_j : \Xi_j \times \hat{S}_j \times \hat{M}_j \to \Delta(X_j)$. Every draw ξ_j from Ξ_j determines a signal distribution $\hat{\sigma}^{\xi_j}_j \equiv \hat{\mathfrak{s}}_j(\xi_j) \in \Delta(\hat{S}_j)$ and an extended decision rule $\hat{\phi}^{\xi_j}_j \equiv \hat{\mathfrak{f}}_j(\xi_j, \cdot, \cdot) : \hat{S}_j \times \hat{M}_j \to \Delta(X_j)$, which together pin down a mechanism $\hat{\gamma}^{\xi_j}_j \equiv (\hat{\sigma}^{\xi_j}_j, \hat{\phi}^{\xi_j}_j) \in \hat{\Gamma}_j$. Slightly abusing notation, we shall use $\hat{\mu}_j$ to denote both principal j's mixed strategy in $G^{\hat{S}\hat{M}}$ and the deterministic mapping $\xi_j \mapsto \hat{\gamma}^{\xi_j}_j$ from Ξ_j to $\hat{\Gamma}_j$ corresponding to such a strategy.

Letting $\hat{\Gamma} \equiv \times_{j=1}^J \hat{\Gamma}_j$, $\hat{S}^i \equiv \times_{j=1} \hat{S}^i_j$, and $\hat{M}^i \equiv \times_{j=1} \hat{M}^i_j$, a strategy for agent i is a function $\hat{\lambda}^i : \Xi^i \times \hat{\Gamma} \times \hat{S}^i \times \Omega^i \to \hat{M}^i$, where $\Xi^i \equiv [0,1]$ is a sampling space for player i, endowed with its Borel σ-field $\mathcal{B}([0,1])$ and Lebesgue measure $\mathrm{d}\xi^i$. In line with Aumann (1961, 1963), however, the difficulty is that, without further restrictions, it is impossible to define a measurable structure over the spaces of decision rules for the principals with respect to which the agents' strategies can be taken as measurable,[14] which is needed for the allocation induced by a strategy profile $(\hat{\mu}, \hat{\lambda}) \equiv ((\hat{\mu}_j)_{j=1}^J, (\hat{\lambda}^i)_{i=1}^I)$ to be well-defined by integrals of the form

$$z_{\hat{\mu},\hat{\lambda}}(x\,|\,\omega) \equiv \int_{\times_{j=1}^J \Xi_j} \int_{\times_{i=1}^I \Xi^i} \int_{\times_{j=1}^J S_j} \prod_{j=1}^J \hat{\phi}^{\xi_j}_j (\hat{s}_j, (\hat{\lambda}^{i,\xi^i}_j((\hat{\gamma}^{\xi_k}_k)_{k=1}^J, \hat{s}^i, \omega^i))_{i=1}^I)(x_j) \bigotimes_{j=1}^J \hat{\sigma}^{\xi_j}_j(\mathrm{d}\hat{s}_j) \bigotimes_{i=1}^I \mathrm{d}\xi^i \bigotimes_{j=1}^J \mathrm{d}\xi_j \quad (5)$$

for all $(\omega, x) \in \Omega \times X$.[15] To avoid this pitfall, we shall assume that the principals randomize only over countably many extended decision rules on path, and only require every agent i's strategy to be $(\mathcal{B}([0,1]) \otimes \hat{\mathbf{\Sigma}} \otimes \hat{\mathcal{S}}^i \otimes 2^{\Omega^i}, \hat{\mathcal{M}}^i)$-measurable, where $\hat{\mathbf{\Sigma}} \equiv \bigotimes_{j=1}^J \mathcal{B}(\Delta(\hat{S}_j))$ and $\hat{\mathcal{S}}^i = \hat{\mathcal{M}}^i \equiv \bigotimes_{j=1}^J \mathcal{B}([0,1])$. This formulation ensures that allocations and payoffs are well-defined, while still enabling each principal to randomize over uncountably many signal structures; moreover, within a signal structure, randomization over signals sent to the agents generates a large set of effective decision rules.

Remark Because any uncountable Polish space equipped with the Borel σ-field generated by a compatible metric is Borel-isomorphic to $([0,1], \mathcal{B}([0,1]))$, any game in which all the principals' signal and message spaces are uncountable Polish spaces is strategically equivalent to the game $G^{\hat{S}\hat{M}}$ defined above, in the sense that the sets of PBE outcomes of the two games

[14] Doval and Skreta (2022) encounter a similar difficulty in their analysis of mechanism design with limited commitment.

[15] Because one wishes the agents to be able to distinguish finely between extended decision rules, it will not do to endow the spaces of such rules with the cylindrical σ-field, for which singletons are not measurable.

coincide. Therefore, when referring to general games with rich signal and message spaces, we hereafter have $G^{\hat{S}\hat{M}}$ in mind.

4.2 Extended-Direct Mechanisms and P-Truthful Equilibria

Below we show that all equilibrium outcomes of $G^{\hat{S}\hat{M}}$ can also be supported in a game in which each principal posts an *extended-direct* mechanism, which asks each agent to report his *extended* type, that is, his exogenous type and the signals he received from the other principals. Specifically, we focus on the class of *p-truthful* equilibria, in which the principals play pure strategies and the agents truthfully report their extended types on path.

Definition 2 *Given a primitive game G, $G^{\mathring{S}\mathring{M}}$ is the game with private disclosures in which every principal j posts an extended-direct mechanism $\mathring{\gamma}_j \equiv (\mathring{\sigma}_j, \mathring{\phi}_j)$, where $\mathring{S}^i_j \equiv [0,1]$ and $\mathring{M}^i_j \equiv \Omega^i \times [0,1]^{J-1}$ for all i.*

Hence, in $G^{\mathring{S}\mathring{M}}$, a mechanism for principal j draws a profile of signals $\mathring{s}_j \equiv (\mathring{s}^i_j)^I_{i=1} \in \mathring{S}_j \equiv [0,1]^I$ according to the distribution $\mathring{\sigma}_j \in \Delta(\mathring{S}_j)$, privately discloses the component $\mathring{s}^i_j$ of $\mathring{s}_j$ to every agent i, asks every agent i to report his exogenous type ω^i along with the signals $\mathring{s}^i_{-j} \equiv (\mathring{s}^i_k)_{k\neq j} \in [0,1]^{J-1}$ the agent privately received from the other principals, and finally selects a decision according to $\mathring{\phi}_j : \mathring{S}_j \times \mathring{M}_j \to \Delta(X_j)$, where $\mathring{M}_j \equiv \times^I_{i=1} \mathring{M}^i_j$. The players' strategies $(\mathring{\mu}, \mathring{\lambda}) \equiv ((\mathring{\mu}_j)^J_{j=1}, (\mathring{\lambda}^i)^I_{i=1})$ in $G^{\mathring{S}\mathring{M}}$ are required to satisfy the same measurability conditions as in $G^{\hat{S}\hat{M}}$ and the allocation $z_{\mathring{\mu},\mathring{\lambda}}$ induced by the strategies $(\mathring{\mu}, \mathring{\lambda})$ is defined in analogy with (5). For all i and j, we denote by $q^i_j \equiv (\omega^i, \mathring{s}^i_{-j}) \in \Omega^i \times [0,1]^{J-1}$ agent i's extended type vis-à-vis principal j.

Definition 3 *A PBE $(\mathring{\mu}^*, \mathring{\lambda}^*)$ of $G^{\mathring{S}\mathring{M}}$ is p-truthful if*

(i) for each j, $\mathring{\mu}^*_j$ is a pure strategy that selects with probability 1 an extended-direct mechanism $\mathring{\gamma}^*_j \equiv (\mathring{\sigma}^*_j, \mathring{\phi}^*_j)$;

(ii) on path, that is, in the subgame $\mathring{\gamma}^* \equiv (\mathring{\gamma}^*_j)^J_{j=1}$, every agent i truthfully reports q^i_j to every principal j.

The central theorem of this section establishes that p-truthful equilibria of $G^{\mathring{S}\mathring{M}}$ support all equilibrium outcomes of all short-communication games with rich signal and message spaces, including those supported by principals and agents playing mixed strategies.

Theorem 1 *For any primitive game G and for any PBE $(\hat{\mu}^*, \hat{\lambda}^*)$ of $G^{\hat{S}\hat{M}}$, there exists an outcome-equivalent p-truthful PBE $(\mathring{\mu}^*, \mathring{\lambda}^*)$ of $G^{\mathring{S}\mathring{M}}$; that is, $z_{\mathring{\mu}^*,\mathring{\lambda}^*} = z_{\hat{\mu}^*,\hat{\lambda}^*}$.*

The structure of the argument can be sketched as follows.

In a mixed-strategy equilibrium of $G^{\hat{S}\hat{M}}$, the agents can use the realizations of the principals' mixed strategies as a device to correlate their behavior within each mechanism. To replicate such a correlation in $G^{\mathring{S}\mathring{M}}$, every principal j encodes into the signal $\mathring{s}_j^i$ to every agent i the sampling variable ξ_j indexing the realization of her mixed strategy in $G^{\hat{S}\hat{M}}$. Moreover, in $G^{\hat{S}\hat{M}}$, given the principals' mechanisms $\hat{\gamma}$, each agent can, by himself, correlate the principals' decisions by randomizing over the messages he sends to the principals. Such a correlation is replicated in $G^{\mathring{S}\mathring{M}}$ by decomposing the sampling variable ξ^i indexing agent i's behavior into a collection of variables ξ_j^i, one for every principal j, with each ξ_j^i independently and uniformly drawn by principal j in $[0,1]$. When aggregated in a suitable way, the variables $(\xi_j^i)_{j=1}^J$ follow the same distribution as the original sampling variable ξ^i indexing agent i's behavior in $G^{\hat{S}\hat{M}}$. Furthermore, the aggregation is done so that none of the principals can unilaterally manipulate the distribution of the statistic aggregating $(\xi_j^i)_{j=1}^J$ and replicating ξ^i. The decomposition thus provides the principals with a way of generating ξ^i as the outcome of a jointly-controlled lottery that is non-manipulable by any of the principals (Aumann and Maschler (1995)). Every principal j then encodes ξ_j^i into the signal $\mathring{s}_j^i$ to agent i, along with the sampling variable ξ_j indexing the realization of her own mixed strategy, and the signal $\hat{s}_j^i$ the principal would have disclosed to the agent in $G^{\hat{S}\hat{M}}$.

Second, we use the original equilibrium strategies $(\hat{\mu}^*, \hat{\lambda}^*)$ in $G^{\hat{S}\hat{M}}$ to construct candidate equilibrium strategies $(\mathring{\mu}^*, \mathring{\lambda}^*)$ in $G^{\mathring{S}\mathring{M}}$.

Every principal j's strategy $\mathring{\mu}_j^*$ is a degenerate distribution—hence, a pure strategy—that selects with probability 1 the following extended-direct mechanism $\mathring{\gamma}_j^* \equiv (\mathring{\sigma}_j^*, \mathring{\phi}_j^*)$.

As for the distribution $\mathring{\sigma}_j^*$ over the agents' signals, principal j first draws ξ_j and $(\xi_j^i)_{i=1}^I$ uniformly in $[0,1]$, with all the draws made independently. She then draws the signals $\hat{s}_j = (\hat{s}_j^i)_{i=1}^I$ from the equilibrium distribution $\hat{\sigma}_j^{*\xi_j}$ corresponding to the realization ξ_j of the sampling variable indexing her equilibrium mixed strategy $\hat{\mu}_j^*$ in $G^{\hat{S}\hat{M}}$. She finally encodes the information $(\xi_j, \hat{s}_j^i, \xi_j^i)$ into the signal $\mathring{s}_j^i$ to agent i, with the encoding governed by an appropriate embedding $\kappa_j^i : \Xi_j \times \hat{S}_j^i \times \Xi_j^i \to \mathring{S}_j^i$.

For any profile of signals $\mathring{s}_j$ drawn from $\mathring{\sigma}_j^*$, the effective decision rule $\mathring{\phi}_j^*(\mathring{s}_j, \cdot) : \mathring{M}_j \to \Delta(X_j)$ then operates as follows. When the message $\mathring{m}_j^i = (\omega^i, \mathring{s}_{-j}^i)$ from agent i is such that every signal $\mathring{s}_k^i$, $k \neq j$, reported by the agent is in the range of the embedding κ_k^i, principal j uses the information $(\xi_k, \hat{s}_k^i, \xi_k^i)$ encoded into the signal $\mathring{s}_k^i$ reported by agent i, along with the information $(\xi_j, \hat{s}_j^i, \xi_j^i)$ encoded into principal j's signal to agent i and agent i's reported type ω^i to identify the message that agent i would have sent in $G^{\hat{S}\hat{M}}$. When, instead, the message $\mathring{m}_j^i = (\omega^i, \mathring{s}_{-j}^i)$ from agent i is such that the signal $\mathring{s}_k^i$ agent i claims to have received from some principal $k \neq j$ is not in the image of the embedding κ_k^i, principal j uses a different embedding $\rho^i_j : \hat{M}^i_j \to \mathring{M}^i_j$ to identify the message agent i would have sent in $G^{\hat{S}\hat{M}}$. The embeddings κ^i_j and ρ^i_j are carefully constructed so that there is no confusion about which message agent i would have sent in $G^{\hat{S}\hat{M}}$.[16]

Once the message every agent i would have sent in $G^{\hat{S}\hat{M}}$ is identified, principal j uses her original equilibrium extended decision rule $\hat{\phi}^{*\xi_j}_j$ in $G^{\hat{S}\hat{M}}$ to identify the decision that, given the signals sent by principal j to the agents and the messages received from them, she would have selected in $G^{\hat{S}\hat{M}}$. Both on and off the equilibrium path, principal j uses the information contained in the message from every agent i to identify the message that agent i would have sent in $G^{\hat{S}\hat{M}}$ when behaving according to his equilibrium strategy $\hat{\lambda}^{*i}$.

The remainder of the proof consists in establishing that (a) when every principal j posts her equilibrium mechanism $\mathring{\gamma}^*_j = (\mathring{\sigma}^*_j, \mathring{\phi}^*_j)$, it is optimal for every agent i to report to every principal j his extended type $q^i_j \equiv (\omega^i, \mathring{s}^i_{-j})$ (on-path truth-telling), and (b) when any of the principals deviates by offering an extended-direct mechanism $\mathring{\gamma}_j \neq \mathring{\gamma}^*_j$, it is optimal for every agent i to send to each deviating principal the analogue of the message he would have sent to her in $G^{\hat{S}\hat{M}}$ (appropriately translated to account for the difference in the language between $G^{\hat{S}\hat{M}}$ and $G^{\mathring{S}\mathring{M}}$), and to send to each non-deviating principal a message that reveals to each principal the message the agent would have sent in $G^{\hat{S}\hat{M}}$. We finally show that the agents' strategies $\mathring{\lambda}^*$ so constructed induce a BNE in every subgame of $G^{\mathring{S}\mathring{M}}$ and that, given these strategies, no principal has an incentive to unilaterally deviate from $\mathring{\gamma}^*$.

Remark In proving Theorem 1, the richness of the sets of private disclosures serves two main purposes. First, it allows one to generate all possible correlations arising from players' independent mixing; second, it guarantees that the agents have enough messages to adapt their reports to any possible unilateral deviation by a principal.

4.3 Long Communication

The above analysis leaves open an important question, namely, whether the restriction to short communication is warranted in games with multiple mechanism designers.[17] The concern is that allowing for several rounds of communication in which principals disclose signals to the agents and solicit messages from them may generate new strategic effects and further enrich equilibrium analysis. This is notably because a principal may strategically delay the revelation of relevant information to the agents so as to maintain greater flexibility over her decision and/or better tailor her decision to the decisions taken by the other

[16]If $\mathring{m}^i_j$ is neither consistent with principals $k \neq j$ using the embeddings κ^i_k to encode $(\xi_k, \hat{s}^i_k, \xi^i_k)$ into the signal $\mathring{s}^i_k$ to agent i, nor is a translation of a message agent i may have sent in the original game according to the embedding ρ^i_j, principal j replaces agent i's message with a default message in $\hat{M}^i_j$.

[17]We thank Bart Lipman to encourage us to explore this possibility

principals. For instance, she may use uninformative signals in the first round, let her opponents draw their signals and, only then, depending on the agents' first-period reports, determine her signals' distribution and her decision in response to the information reported by the agents about the other principals' behavior in both rounds. The central result of this and the next section is that the class of extended-direct mechanisms and the corresponding class of p-truthful equilibria play a fundamental role beyond the specific short-communication protocol considered in Theorem 1.

To tackle these issues, we extend the game $G^{\hat{S}\hat{M}}$ to $T > 1$ rounds of communication, with T potentially infinite. Allowing for protocols with infinitely many rounds permits us to capture the possibility that a principal expecting the other principals to finalize their decisions in finite time may want to postpone her decision till the uncertainty generated by the other mechanisms has been resolved.

The long-communication game $G^{\hat{S}\hat{M}T}$ unfolds as follows. In every round $t < \infty$, $1 \leq t \leq T$, every principal j simultaneously sends a signal $\hat{s}^i_j(t) \in \hat{S}^i_j(t) \equiv [0, 1]$ to every agent i, and then receives a message $\hat{m}^i_j(t) \in \hat{M}^i_j(t) \equiv [0, 1]$ from every agent i. Denoting by $\hat{S}_j(t) \equiv \times_{i=1}^I \hat{S}^i_j(t)$ and $\hat{M}_j(t) \equiv \times_{i=1}^I \hat{M}^i_j(t)$ principal j's round-t signal and message spaces, we can recursively define the space of principal j's private histories at the outset of round t by $\hat{H}_j(1) \equiv \{\emptyset\}$ and $\hat{H}_j(t) \equiv \hat{H}_j(t-1) \times \hat{S}_j(t-1) \times \hat{M}_j(t-1)$ for $t > 1$. A mechanism for principal j in $G^{\hat{S}\hat{M}T}$ is then a sequence $(\hat{\sigma}_j(t))_{t=1}^T$ of Borel-measurable transition probabilities $\hat{\sigma}_j(t) : \hat{H}_j(t) \to \Delta(\hat{S}_j(t))$ describing how the round-t distribution of principal j's signals responds to the past signals sent by the principal to the agents and the messages received from them, together with a Borel-measurable extended decision rule $\hat{\phi}^T_j : \hat{H}^T_j \to \Delta(X_j)$, in which $\hat{H}^T_j \equiv \hat{H}_j(T) \times \hat{S}_j(T) \times \hat{M}_j(T)$ if $T < \infty$ and with $\hat{H}^T_j$ defined as the (projective) limit of the family $(\hat{H}_j(t))_{t=1}^\infty$ if $T = \infty$.

The game $G^{\hat{S}\hat{M}T}$ starts with every principal j committing to a mechanism $\hat{\gamma}^T_j \equiv ((\hat{\sigma}_j(t))_{t=1}^T, \hat{\phi}^T_j)$. Then, each agent, upon observing the posted mechanisms, the first-round signals he receives from the principals, and his exogenous type, sends a first-round message to each principal. In every subsequent round $t < \infty$, $1 < t \leq T$, each agent receives a new signal from every principal j, drawn from the corresponding round-t distribution $\hat{\sigma}_j(t)(\cdot \mid \hat{h}_j(t))$, which depends on principal j's private history $\hat{h}_j(t)$ at the outset of round t. The agent then sends a message to each principal as a function of his private history up to round t. Eventually, the principals' decisions are determined by the extended decision rules $(\hat{\phi}^T_j)_{j=1}^J$ and the entire sequences of signals and messages sent and received by every principal j. Following the reasoning in Section 4.1, any long-communication game G^{SMT} in which, in every round $t < \infty$, $1 \leq t \leq T$, all the principals' signal and message spaces are uncountable Polish spaces is strategically equivalent to $G^{\hat{S}\hat{M}T}$, which we can thus take as the representant of this whole class of games.

The formalism introduced in Section 4.1 can be straightforwardly extended to reformulate players' strategies $\hat{\lambda}^T \equiv (\hat{\lambda}^{iT})_{i=1}^I$ in $G^{\hat{S}\hat{M}T}$.[18] To avoid defining belief processes for the agents during the communication phase of $G^{\hat{S}\hat{M}T}$, we adopt an equilibrium concept that is stronger than BNE, but weaker than PBE. Specifically, a strategy profile $(\hat{\mu}^{*T}, \hat{\lambda}^{*T})$ is a semiperfect Bayesian equilibrium (SPBE) of $G^{\hat{S}\hat{M}T}$ whenever

(i) for each $\hat{\gamma}^T \in \hat{\Gamma}^T$, $\lambda^{*T}(\hat{\gamma}^T)$ is a BNE of the subgame $\hat{\gamma}^T$ played by the agents;

(ii) given the continuation equilibrium strategies $\hat{\lambda}^{*T}$ for the agents, $\hat{\mu}^{*T}$ is a NE of the game played by the principals.

Assessing the relevance of long communication amounts to answering the following questions: (a) Can multiple rounds of communication enable the principals to dynamically use signals to correlate their decisions with the agents' exogenous private information in a way a single round of communication would not? (b) Can equilibrium outcomes of $G^{\mathring{S}\mathring{M}}$ be destabilized if principals can engage into such long communication with the agents? We shall answer both questions in the negative, establishing the irrelevance of long communication for equilibrium analysis. This is a key feature of the revelation principle for competing-mechanism games that we derive in the next section.

4.4 $G^{\mathring{S}\mathring{M}}$ as a Canonical Game

In this section, we establish that p-truthful PBE of the short-communication game $G^{\mathring{S}\mathring{M}}$ in which principals commit to extended-direct mechanisms support all SPBE outcomes of any $G^{\hat{S}\hat{M}T}$ game, and, conversely, that p-truthful PBE outcomes of $G^{\mathring{S}\mathring{M}}$ survive the introduction of long communication. That is, following the terminology introduced by Epstein and Peters (1999), extended-direct mechanisms are *universal*, and, in addition, any p-truthful PBE outcome of $G^{\mathring{S}\mathring{M}}$ is *robust*.

Together with Theorem 1, this result provides a generalized revelation principle for competing-mechanism games, based on $G^{\mathring{S}\mathring{M}}$, its short-communication protocol, and the class of p-truthful PBE. In line with recent contributions to the literature (Doval and Skreta (2022), Sugaya and Wolitzky (2021)), we will say that $G^{\mathring{S}\mathring{M}}$ is a *canonical game*, and that any p-truthful PBE of $G^{\mathring{S}\mathring{M}}$ is a *canonical equilibrium*.

Theorem 2 *For any primitive game G and each $1 < T \leq \infty$, the following holds*:

(i) (Universality) For any SPBE $(\hat{\mu}^{*T}, \hat{\lambda}^{*T})$ of $G^{\hat{S}\hat{M}T}$ such that the agents' behavior strategies $\hat{\lambda}^*(\hat{\gamma}^T)$ in any subgame $\hat{\gamma}^T$ of $G^{\hat{S}\hat{M}T}$ are of uniformly bounded Young class,

[18] We refer to the Online Supplement for a detailed presentation.

there exists an outcome-equivalent p-truthful PBE $(\mathring{\mu}^*, \mathring{\lambda}^*)$ *of* $G^{\mathring{S}\mathring{M}}$*; that is,* $z_{\mathring{\mu}^*,\mathring{\lambda}^*} = z_{\hat{\mu}^{*T},\hat{\lambda}^{*T}}$.

(ii) (Robustness) If $G^{\hat{S}\hat{M}T}$ *admits an SPBE, and the agents' behavior strategies* $\hat{\lambda}^*(\hat{\gamma}^T)$ *in any subgame* $\hat{\gamma}^T$ *of* $G^{\hat{S}\hat{M}T}$ *and in any SPBE of* $G^{\hat{S}\hat{M}T}$ *are of uniformly bounded Young class, then, for any p-truthful PBE* $(\mathring{\mu}^*, \mathring{\lambda}^*)$ *of* $G^{\mathring{S}\mathring{M}}$*, there exists an outcome-equivalent SPBE* $(\hat{\mu}^{*T}, \hat{\lambda}^{*T})$ *of* $G^{\hat{S}\hat{M}T}$*; that is,* $z_{\hat{\mu}^{*T},\hat{\lambda}^{*T}} = z_{\mathring{\mu}^*,\mathring{\lambda}^*}$.

Young classes, which are indexed by ordinals α less than the first uncountable ordinal ω_1, are intimately related to the usual Baire hierarchy, and overall exhaust the set of Borel-measurable functions (Kechris (1995, Chapter II, §24)). The ordinals $\alpha < \omega_1$ index the complexity of such functions. As we will explain below, bounding the complexity of the agents' strategies as required in items (i)–(ii) of Theorem 2 affords us to encode into the message spaces of the short-communication game $G^{\mathring{S}\mathring{M}}$ the message plans used by the agents to respond to the principals' mechanisms in the long-communication game $G^{\hat{S}\hat{M}T}$—and this in a measurable way, ensuring that allocations and payoffs remain well-defined. We view this restriction as a mere technical convenience, and as innocuous for all practical purposes.[19]

The structure of the argument can be sketched as follows.

Universality We start from an SPBE of $G^{\hat{S}\hat{M}T}$ and construct a p-truthful equilibrium of $G^{\mathring{S}\mathring{M}}$ that supports the same outcome.

First, one can safely focus on SPBE of long-communication games in which principals play pure strategies and agents truthfully report their (incremental) information in all rounds. This follows from the analogue of Theorem 1 for an auxiliary long-communication game $G^{\hat{S}\mathring{M}T}$. In this game, as in $G^{\hat{S}\hat{M}T}$, the signal spaces coincide with the unit interval, and, in analogy with $G^{\mathring{S}\mathring{M}}$, the message spaces enable every agent i to report to every principal j (a) his extended type in the first round, that is, $\mathring{M}^i_j(1) \equiv \Omega^i \times [0,1]^{J-1}$, and (b) the additional signals received from the other principals in every subsequent round, that is, $\mathring{M}^i_j(t) \equiv [0,1]^{J-1}$ for all $t < \infty$, $1 < t \leq T$. For every SPBE $(\hat{\mu}^{*T}, \hat{\lambda}^{*T})$ of $G^{\hat{S}\hat{M}T}$, there exists an outcome-equivalent p-truthful SPBE $(\mathring{\mu}^{*T}, \mathring{\lambda}^{*T})$ of $G^{\hat{S}\mathring{M}T}$.

Second, given any collection of transition probabilities $((\mathring{\sigma}_j(t))_{t=1}^T)_{j=1}^J$, all the signals past round 1 can be generated by letting every principal j make $I \times (T-1)$ independent draws from the uniform distribution over $[0,1]$. Indeed, for all j and $t < \infty$, $1 < t \leq T$, and for any private history $\mathring{h}_j(t)$ of principal j in round t, one can use the cdf of the distribution $\mathring{\sigma}_j(t)(\cdot \mid \mathring{h}_j(t))$ to construct the marginal over principal j's round-t signal to

[19]For instance, the Dirichlet function $\mathbb{1}_{\mathbb{Q} \cap [0,1]}$ is of Baire class 2, and hence of Young class 3. By contrast, Theorem 2 allows for strategies of arbitrary complexity $\alpha < \omega_1$. We refer to the Online Supplement for a detailed definition of Young classes.

agent 1 and a collection of conditional distributions, one for each agent $i > 1$. We then map independent draws into signals using the corresponding quantile functions for each agent. The decisions implemented in $G^{\mathring{S}\mathring{M}T}$ given the mechanisms $\mathring{\gamma}^T \equiv (\mathring{\gamma}_j^T)_{j=1}^J$ then become deterministic functions of the round-1 signals $\mathring{s}(1) = ((\mathring{s}(1))_{i=1}^I)_{j=1}^J$, the messages sent by the agents in all rounds, and $I \times J \times (T-1)$ independent draws from a uniform distribution over $[0, 1]$.

Third, we exploit the independence of these draws to show that the correlation between the principals' decisions and the agents' private information that may be generated by the players' communication past round 1 can be generated by a single random variable, uniformly distributed over $[0, 1]$. To this purpose, we introduce the concept of a *message plan*. A message plan $\mathring{\pi}^{iT} \equiv (\mathring{\pi}^i(t))_{t=1}^T$ for agent i specifies, for each round $t < \infty$, $1 \leq t \leq T$, a profile of messages $\mathring{m}^i(t) \equiv (\mathring{m}_j^i(t))_{j=1}^J$ to the principals as a function of his exogenous type ω^i, the sampling variable ξ^i indexing his behavior strategy, and all the signals $(\mathring{s}^i(\tau))_{\tau=1}^t$ received from all the principals up to round t, where $\mathring{s}^i(\tau) \equiv (\mathring{s}_j^i(\tau))_{j=1}^J$.[20] We show that, given any profile of mechanisms $\mathring{\gamma}^T$, any profile of round-1 signals $\mathring{s}(1)$, any profile of agents' exogenous types ω, any profile of message plans $\mathring{\pi}^T \equiv (\mathring{\pi}^{iT})_{i=1}^I$, and any profile of sampling variables $(\xi^i)_{i=1}^I$ for the agents, the correlation between the principals' decisions and the agents' types generated by the agents following the message plans $\mathring{\pi}^T$ in $\mathring{\gamma}^T$ can be captured by a single random variable ξ_0 drawn uniformly from $[0, 1]$.

The upshot of the analysis so far is that any equilibrium correlation supported via long communication can also be supported via short communication by introducing an additional, uniformly-distributed random variable, and by having the agents report their message plans. The bulk of the proof consists in formalizing these ideas.

First, we construct ξ_0 as a jointly-controlled lottery. Every principal j independently draws I auxiliary random variables $(\xi_j^i)_{i=1}^I$ from the uniform distribution over $[0, 1]$, and then, as in the proof of Theorem 1, constructs the fractional part of their sums, $\xi_j \equiv \{\sum_{i=1}^I \xi_j^i\}$, which is uniformly distributed over $[0, 1]$ and independent of $(\xi_j^i)_{i=1}^I$.[21] As the random variables $(\xi_j)_{j=1}^J$ are also independent, $\xi_0 \equiv \{\sum_{j=1}^J \xi_j\}$ is in turn uniformly distributed over $[0, 1]$ and independent of $((\xi_j^i)_{i=1}^I)_{j=1}^J$. In particular, no principal j can manipulate ξ_0 and, given $(\xi_j^i)_{j=1}^J$, no agent i can infer anything about ξ_0. Hence ξ_0 can be generated by every principal j privately informing every agent i of ξ_j^i, and then every agent i truthfully reporting to every principal j the auxiliary variables ξ_{-j}^i from the other principals, enabling all the principals to reconstruct ξ_0. Notice that this step of the argument crucially relies on there

[20] Notice that, unlike a strategy, a message plan condition neither on the principals' mechanisms nor on the agent's own past behavior.

[21] The fractional part $\{x\}$ of a non-negative real number x is the difference $x - \lfloor x \rfloor$ between x and its integer part $\lfloor x \rfloor$.

being at least two agents.

Next, we show how the richness of the agents' message spaces in $G^{\mathring{S}\mathring{M}}$ permits one to encode into every agent i's message $\mathring{m}_j^i$ to every principal j a description of any message plan $\mathring{\pi}^{iT}$ the agent can follow in $G^{\mathring{S}\mathring{M}T}$, along with his type ω^i, his sampling variable ξ^i indexing his behavior strategy in $G^{\mathring{S}\mathring{M}T}$, the round-1 signals $\mathring{s}_{-j}^i(1)$ from the other principals, and the auxiliary variables ξ_{-j}^i from the other principals. This is where the assumption that the agents' behavior strategies in any subgame of $G^{\hat{S}\hat{M}T}$ are of uniformly bounded Young class is used: it enables one to encode into $[0,1]$ the agents' message plans in such a way that, in any round $t < \infty$, $1 \leq t \leq T$, the action of a message plan on any agent i's sampling variable ξ^i, private history of signals $\mathring{h}^i(t)$, and type ω^i is a Borel-measurable function of $(\xi^i, \mathring{h}^i(t), \omega^i)$ and of a code for the message plan. Finally, every principal j can encode into her signal to every agent i in $G^{\mathring{S}\mathring{M}}$ any round-1 signal she may send to him in $G^{\mathring{S}\mathring{M}T}$, along with the auxiliary variable ξ_j^i used to construct the jointly-controlled lottery ξ^0.

Equipped with these results, the remaining steps of the argument show how to construct a PBE $(\mathring{\mu}^{\#}, \mathring{\lambda}^{\#})$ of $G^{\mathring{S}\mathring{M}}$ that is outcome-equivalent to the SPBE $(\mathring{\mu}^{*T}, \mathring{\lambda}^{*T})$ of $G^{\mathring{S}\mathring{M}T}$.

In $(\mathring{\mu}^{\#}, \mathring{\lambda}^{\#})$, every principal j selects with probability 1 the mechanism $\mathring{\gamma}_j^{\#} \equiv (\mathring{\sigma}_j^{\#}, \mathring{\phi}_j^{\#})$ defined as follows. (a) To construct the distribution $\mathring{\sigma}_j^{\#}$, principal j first draws the auxiliary variables $(\xi_j^i)_{i=1}^I$ independently and uniformly from $[0,1]$, and then draws the signals $\mathring{s}_j(1) \equiv (\mathring{s}_j^i(1))_{i=1}^I$ from the round-1 distribution $\mathring{\sigma}_j^*(1)$ of her equilibrium mechanism $\mathring{\gamma}_j^{*T} \equiv (\mathring{\sigma}_j^{*T}, \mathring{\phi}_j^{*T})$ in $G^{\mathring{S}\mathring{M}T}$.[22] Every signal $\mathring{s}_j^i$ is then obtained by encoding the information $(\mathring{s}_j^i(1), \xi_j^i)$ using an appropriate embedding, and the distribution $\mathring{\sigma}_j^{\#}$ is generated through the corresponding push-forward of $\mathring{\sigma}_j^*(1)$ and I uniform distributions. (b) The extended decision rule $\mathring{\phi}_j^{\#}$ is constructed as follows. From the message $\mathring{m}_j^i$ received from agent i, the principal extracts the agent's exogenous type ω^i, the sampling variable ξ^i indexing the agent's behavioral strategy in $G^{\mathring{S}\mathring{M}T}$, the agent's message plan $\mathring{\pi}^{iT}$, the round-1 signals $\mathring{s}_{-j}^i(1)$ and the auxiliary variables ξ_{-j}^i received from the other principals. Next, the principal extracts from her own signal $\mathring{s}_j^i$ to agent i the round-1 signal $\mathring{s}_j^i(1)$ she would have sent in $G^{\mathring{S}\mathring{M}T}$ and the auxiliary variable ξ_j^i. She then reconstructs the variable ξ_0 from the agents' reports and then uses all the decoded information to implement the same decision she would have implemented in $G^{\mathring{S}\mathring{M}T}$ under the mechanism $\mathring{\gamma}_j^{*T}$. The possibility to encode the agents' message plans into $[0,1]$ allows us to eschew the admissibility problem pointed out by Aumann (1961, 1963), and ensures that $\mathring{\phi}_j^{\#}$ is Borel-measurable, as requested.

When every principal j posts the mechanism $\mathring{\gamma}_j^{\#}$, every agent i reports to every principal j a message that encodes his true type ω^i, the sampling variable ξ^i indexing his behavioral strategy in $G^{\mathring{S}\mathring{M}T}$, the round-1 signals $\mathring{s}_{-j}^i(1)$ and the auxiliary variables ξ_{-j}^i decoded from

[22] Recall that $\mathring{\mu}_j^{*T}$ is pure, selecting with probability 1 the mechanism $\mathring{\gamma}_j^{*T}$.

the other principals' signals $\mathring{s}^i_{-j}$, and a plan $\mathring{\pi}^{iT}$ that specifies truthful reporting at each round. Similarly, when any of the principals posts a mechanism $\mathring{\gamma}_j \equiv (\mathring{\sigma}_j, \mathring{\phi}_j) \neq \mathring{\gamma}_j^{\#}$, each agent behaves as if, in the long-communication game $G^{\mathring{S}\mathring{M}T}$, the non-deviating principals offered their equilibrium mechanisms, and every deviating principal j offered a mechanism $\mathring{\gamma}_j^T$ that is equivalent to $\mathring{\gamma}_j$—that is, whose round-1 signal distribution $\mathring{\sigma}_j(1)$ is the same as $\mathring{\sigma}_j$, whose round-t distribution $\mathring{\sigma}_j(t)$, for $t < \infty$, $1 < t \leq T$, is arbitrary, and whose extended decision rule $\mathring{\phi}_j^T$ is invariant in the signals and messages past round 1 and agrees with $\mathring{\phi}_j$ over the round-1 signals and messages. Specifically, each agent sends to each deviating principal the same message he would have sent in $G^{\mathring{S}\mathring{M}T}$ and to each non-deviating principal a message that, when decoded, reveals the message plan the agent would have followed in the subgame of $G^{\mathring{S}\mathring{M}T}$ described above.

Given this construction, it is then straightforward to show that, in each state ω, the distribution over the principals' decisions under the strategies $(\mathring{\mu}^{\#}, \mathring{\lambda}^{\#})$ in $G^{\mathring{S}\mathring{M}}$ is the same as that induced by the strategies $(\mathring{\mu}^{*T}, \mathring{\lambda}^{*T})$ in $G^{\mathring{S}\mathring{M}T}$. In addition, we show that the agents' strategies $\mathring{\lambda}^{\#}$ sustain a BNE in every subgame of $G^{\mathring{S}\mathring{M}}$. Finally, we verify that, given the agents' strategy profile $\mathring{\lambda}^{\#}$, the strategy profile $\mathring{\mu}^{\#}$ is a NE of the principals' game. Now, the equilibrium $(\mathring{\mu}^{\#}, \mathring{\lambda}^{\#})$ of $G^{\mathring{S}\mathring{M}}$ need not be p-truthful; in fact, to encode into their messages all the information necessary to permit the principals to reconstruct the outcome in $G^{\mathring{S}\mathring{M}T}$, the agents typically do not truthfully report their type and the signals received from the other principals. However, applying Theorem 1, one can construct a p-truthful PBE $(\mathring{\mu}^{*}, \mathring{\lambda}^{*})$ of $G^{\mathring{S}\mathring{M}}$ that is outcome-equivalent to $(\mathring{\mu}^{\#}, \mathring{\lambda}^{\#})$, which completes the proof of universality.

Robustness We start from a p-truthful PBE of $G^{\mathring{S}\mathring{M}}$ and construct an SPBE of $G^{\hat{S}\hat{M}T}$ that supports the same outcome.

First, in $G^{\hat{S}\hat{M}T}$, every principal j posts with probability 1 a mechanism $\hat{\gamma}_j^{*T} = (\hat{\sigma}_j^{*T}, \hat{\phi}_j^{*T})$ that has the following properties. (a) The round-1 signal distribution $\hat{\sigma}_j^*(1)$ is identical to the signal distribution $\mathring{\sigma}_j^*$ in the equilibrium mechanism $\mathring{\gamma}_j^* = (\mathring{\sigma}_j^*, \mathring{\phi}_j^*)$ of the short-communication game $G^{\mathring{S}\mathring{M}}$; moreover, for each $t < \infty$, $1 < t \leq T$, the round-t transition probability $\hat{\sigma}_j^*(t)$ is degenerate. (b) The extended decision rule $\hat{\phi}_j^{*T}$ is invariant in signals and messages past round 1 and agrees with $\mathring{\phi}_j^*$ over round-1 signals and messages, modulo an appropriate Borel isomorphism between every message space $\hat{M}_j^i(1) = [0, 1]$ in $G^{\hat{S}\hat{M}T}$ and the corresponding message space $\mathring{M}_j^i = \Omega^i \times [0, 1]^{J-1}$ in $G^{\mathring{S}\mathring{M}}$.

When each principal in $G^{\hat{S}\hat{M}T}$ posts the above mechanism $\hat{\gamma}_j^{*T}$, each agent sends the analogue of the message he would have sent in the subgame of $G^{\mathring{S}\mathring{M}}$ in which each principal posts the equilibrium mechanism $\mathring{\gamma}_j^*$. The message that each agent sends to each principal is thus identified by the agent's strategy $\mathring{\lambda}^{*i}$ in $G^{\mathring{S}\mathring{M}}$, together with the Borel isomorphism between $\hat{M}^i_j(1)$ and $\mathring{M}^i_j$. When, instead, some principal in $G^{\hat{S}\hat{M}T}$ deviates, each agent maps each deviating principal's mechanism into a "strategically-equivalent" analogue in $G^{\mathring{S}\mathring{M}}$. The difficulty in identifying such mechanisms originates in the fact that (a) the length of communication differs across the two games, and (b) the signals are invariant in the agents' messages in one game but not in the other. The identification obtains by requiring the decisions that each agent can induce by varying his message plan (holding the other agents' message plans fixed) to correspond to those he could have induced through the corresponding messages in the short-communication mechanism. Each agent's reporting strategy in the long-communication game $G^{\hat{S}\hat{M}T}$ is then constructed from his equilibrium strategy in the short-communication game $G^{\mathring{S}\mathring{M}}$ as follows. The agent's round-1 report to each non-deviating principal is determined by the strategy $\mathring{\lambda}^{*i}$ he would have followed had he observed the strategically-equivalent deviations in $G^{\mathring{S}\mathring{M}}$, whereas the reports made in subsequent rounds have no strategic effect, and are therefore arbitrary. Instead, his reports to any deviating principal are determined, in any round $t < \infty$, $1 \leq t \leq T$, by the agent following the message plan encoded in the messages he would have sent in $G^{\mathring{S}\mathring{M}}$ following his equilibrium strategy $\mathring{\lambda}^{*i}$.

This construction guarantees that the outcome supported by the strategies $(\hat{\gamma}^{*T}, \hat{\lambda}^{*T})$ in $G^{\hat{S}\hat{M}T}$ coincides with that supported by the equilibrium strategies $(\mathring{\gamma}^{*}, \mathring{\lambda}^{*})$ in $G^{\mathring{S}\mathring{M}}$. To complete the proof, we show that the strategies $(\hat{\gamma}^{*T}, \hat{\lambda}^{*T})$ constitute an equilibrium in $G^{\hat{S}\hat{M}T}$. We first prove that the strategies $(\hat{\lambda}^{*iT})_{i=1}^{I}$ for the agents sustain a BNE in any subgame of $G^{\hat{S}\hat{M}T}$. We next prove that, given the strategies $\hat{\lambda}^{*T}$, the mechanisms $(\hat{\gamma}^{*T}_j)_{j=1}^{J}$ sustain a NE for the principals, which completes the proof of robustness.

Theorem 2 provides a new revelation principle for competing-mechanism games. It covers a richer framework than the one considered in Epstein and Peters (1999). First, it accounts for private disclosures, which cannot be disregarded if one aims at supporting all possible equilibrium outcomes (Propositions 1–2). Second, it accommodates long communication, which is particularly relevant in the presence of private disclosures. Third, it allows all players to play mixed strategies. Fourth, it establishes robustness of the equilibrium allocations of short-communication games to alternative communication structures involving arbitrarily many rounds of exchanges of signals and messages between the principals and the agents. Together, Theorems 1–2 identify a canonical game and a canonical class of equilibria. That is, any equilibrium outcome—including those supported by the principals mixing over their mechanisms or the agents mixing over their messages, and/or the principals and the agents engaging in long communication—can be supported in a p-truthful equilibrium of the $G^{\mathring{S}\mathring{M}}$ game. These results emphasize that competing-principal games share important properties with classical settings in which there is a single mechanism designer (Myerson (1979, 1982)). First, there is no loss of generality in considering communication protocols that involve only one round of communication between the parties. Second, the agents' on-path behavior can be interpreted in terms of truthful reports of the private information they hold in the game, namely, their extended types, in the equilibrium mechanisms. Third, the results apply to arbitrary primitive games, including those in which an agent's payoff depends on the entire profile of types as well as on the decisions implemented in all the principals' mechanisms.

5 Concluding Remarks

This paper explores a novel dimension of mechanism design for settings in which several principals contract with several agents—namely, the possibility for the principals to design the agents' market information by means of private disclosures, which asymmetrically inform the agents of how a principal's decision responds to the agents' messages. The relevance of private disclosures is established by means of two examples, which motivate the need for a novel approach to competing-mechanism games.

We provide a general framework that incorporates private disclosures and allows for rich communication between the principals and the agents. Specifically, we identify the class of extended-direct mechanisms, whereby a principal asks every agent to report his exogenous type along with the signals received from the other principals, and the class of p-truthful equilibria, in which principals play pure strategies and each agent reports truthfully to each principal on path.

We show that extended-direct mechanisms are universal and that the corresponding p-truthful equilibrium outcomes are robust. Thus no additional outcome can be supported by allowing the principals to offer arbitrary indirect mechanisms, including long-communication ones, while any p-truthful equilibrium outcome survives against deviations to alternative, possibly long-communication mechanisms. These results provide a new revelation principle for competing-mechanism games.

References

[1] Akbarpour, M., and S. Li (2020): "Credible Auctions: A Trilemma," *Econometrica*, 88(2), 425–467.

[2] Attar, A., E. Campioni, T. Mariotti, and A. Pavan (2025): "Gradual Resolution of Uncertainty in Common Agency," Mimeo, Toulouse School of Economics.

[3] Attar, A., E. Campioni, and G. Piaser (2013): "Two-Sided Communication in Competing Mechanism Games," *Journal of Mathematical Economics*, 49(1), 62–70.

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

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

[6] Aumann, R.J. (1963): "On Choosing a Function at Random," in *Ergodic Theory*, ed. by F.W. Wright. Cambridge, MA: Academic Press, 1–20.

[7] Aumann, R.J. (1964): "Mixed and Behavior Strategies in Infinite Extensive Games," in *Advances in Game Theory*, Annals of Mathematics Studies Vol. 52, ed. by M. Dresher, L.S. Shapley, and A.W. Tucker. Princeton: Princeton University Press, 627–650.

[8] Aumann, R.J., and M.B Maschler (1995): *Repeated Games with Incomplete Information.* Cambridge, MA: MIT Press.

[9] Banchio, M., A. Skrzypacz, and F. Yang (2024): "Credible Informed Seller," Mimeo, Stanford University.

[10] Dequiedt, V., and D. Martimort (2015): "Vertical Contracting with Informational Opportunism," *American Economic Review*, 105(7), 2141–2182.

[11] Doval, L., and V. Skreta (2022): "Mechanism Design with Limited Commitment," *Econometrica*, 90(4), 1463–1500.

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

[13] Forges, F. (1986): "An Approach to Communication Equilibria," *Econometrica*, 54(6), 1375–1385.

[14] Hart, O., and J. Tirole (1990): "Vertical Integration and Market Foreclosure," *Brookings Papers on Economic Activity: Microeconomics*, 1990, 205–286.

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

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

[17] Kechris, A.S. (1995): *Classical Descriptive Set Theory*. New York: Springer Verlag.

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

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

[20] McAfee, R.P., and M. Schwartz (1994): "Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity," *American Economic Review*, 84(1), 210–230.

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

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

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

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

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

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

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

[31] Rey, P., and J. Tirole (1986): "The Logic of Vertical Restraints," *American Economic Review*, 76(5), 921–939.

[32] Segal, I. (1999): "Contracting with Externalities," *Quarterly Journal of Economics*, 114(2), 337–388.

[33] Segal, I., and M.D. Whinston (2003): "Robust Predictions for Bilateral Contracting with Externalities," *Econometrica*, 71(3), 757–791.

[34] Sugaya, T., and A. Wolitzky (2022): "The Revelation Principle in Multistage Games," *Review of Economic Studies*, 88(3), 1503–1540.

[35] Szentes, B. (2009): "A Note on 'Mechanism Games with Multiple Principals and Three or More Agents' by T. Yamashita," Mimeo, Department of Economics, University College London.

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

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

Online Supplement for ‘Keeping the Agents in the Dark: Competing Mechanisms, Private Disclosures, and the Revelation Principle’

Andrea Attar* Eloisa Campioni†
Thomas Mariotti‡ Alessandro Pavan§

June 28, 2025

Abstract

This supplement provides complete proofs of Lemmas 1–4 in Section 3, and of Theorems 1–2 in Section 4.

*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.

S.1 Proofs for Section 3

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 $(5, 5.5]$ and, correspondingly, any payoff for P1 in $[-5.5, -5)$, can also be supported in a PBE of G_1^M, which completes the proof.

Step 1 We first prove that the outcome

$$z(\omega_L, \omega_L) \equiv \delta_{(x_{11}, x_{21})}, \quad z(\omega_H, \omega_H) \equiv \delta_{(x_{12}, x_{22})}, \tag{S.1}$$

in which P2 obtains her minimum feasible payoff of 5 and P1 her maximum feasible payoff of -5, can be supported in a PBE of G_1^M. To this end, we first show that, if P1 and P2 post recommendation mechanisms, then there exists a BNE supporting the outcome (S.1). We next show that, in every subgame in which P1 posts her equilibrium recommendation mechanism, there exists a 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.

On Path Suppose that both P1 and P2 post recommendation mechanisms ϕ_1^r and ϕ_2^r. We assume that, for each j, $\overline{\omega}_j^1 = \overline{\omega}_j^2 \equiv \omega_L$, 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

$$d_1^*(\omega) \equiv \begin{cases} x_{11} & \text{if } \omega = (\omega_L, \omega_L) \\ x_{12} & \text{otherwise} \end{cases} \quad \text{and} \quad d_2^*(\omega) \equiv \begin{cases} x_{21} & \text{if } \omega = (\omega_L, \omega_L) \\ x_{22} & \text{otherwise} \end{cases} \tag{S.2}$$

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, while A3 has no information to report. To see this, we only need to observe that these strategies implement the outcome (S.1), which yields A1, A2, and A3 their maximum feasible payoff of 8 in each state. These strategies thus form a BNE of the subgame (ϕ_1^r, ϕ_2^r). The claim follows.

Off Path Because P1 obtains her maximum feasible payoff of -5, she has no profitable deviation. Suppose then that P2 deviates to some arbitrary standard mechanism $\phi_2 : M_2 \to \Delta(X_2)$, and let $p(m_2)$ be the probability that the lottery $\phi_2(m_2)$ assigns to decision x_{21} when the agents send the messages $m_2 \equiv (m_2^1, m_2^2, m_2^3) \in M_2$ to P2. Now, let

$$\overline{p} \equiv \max_{m_2 \in M_2} p(m_2) \tag{S.3}$$

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

$$\underline{p} \equiv \min_{(m_2^1, m_2^2) \in M_2^1 \times M_2^2} p(m_2^1, m_2^2, \overline{m}_2^3) \tag{S.4}$$

and select a message profile $(\underline{m}_2^1, \underline{m}_2^2) \in M_2^1 \times M_2^2$ for A1 and A2 that, given $\overline{m}_2^3$, achieves the minimum in (S.4). That $\overline{p}$, $\overline{m}_2$, $\underline{p}$, and $(\underline{m}_2^1, \underline{m}_2^2)$ are well-defined for any given ϕ_2 follows from the fact that 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: $\overline{p} \geq \frac{1}{2}$ Suppose first that ϕ_2 is such that $\overline{p} \geq \frac{1}{2}$. We claim that the subgame (ϕ_1^r, ϕ_2) admits a BNE that satisfies the following properties: (i) all agents recommend the direct mechanism d_1^* to P1, as if P2 did not deviate; (ii) A1 and A2 truthfully report their types to P1; (iii) A3 sends message $\overline{m}_2^3$ to P2; (iv) P2 obtains a payoff of 5.

As for (i), the argument is that unilaterally sending a different recommendation to P1 is of no avail as no agent is pivotal. Consider then (ii). Suppose first that the state is (ω_L, ω_L). Because $\overline{p} \geq \frac{1}{2}$, $8\overline{p} + (1 - \overline{p}) \geq 4.5$. From Table 1, and by definition of d_1^* and $\overline{m}_2$, it thus follows that, if A2 reports ω_L to P1 and sends $\overline{m}_2^2$ to P2, and if A3 sends $\overline{m}_2^3$ to P2, then A1 best responds by reporting ω_L to P1 and sending $\overline{m}_2^1$ to P2; notice, in particular, that, because $\overline{\omega}_1^1 = \omega_L$, if A1 sends a message $m_1^1 \notin D_1 \times \Omega^1$ to P1, then P1 takes the same decision as if A1 truthfully reported his type to her. The argument for A2 is identical. As for (iii), if A1 and A2 report ω_L to P1 and send $\overline{m}_2^1$ and $\overline{m}_2^2$ to P2, sending $\overline{m}_2^3$ to P2 is optimal for A3. 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 and the uninformed agent A3 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: $\overline{p} < \frac{1}{2}$ Suppose next that ϕ_2 is such that $\overline{p} < \frac{1}{2}$. We claim that the subgame (ϕ_1^r, ϕ_2) admits a BNE that satisfies the following properties: (i) all agents recommend the direct mechanism

$$d_1(\omega) \equiv \begin{cases} x_{12} & \text{if } \omega = (\omega_H, \omega_H) \\ x_{11} & \text{otherwise} \end{cases} \tag{S.5}$$

to P1; (ii) A1 and A2 truthfully report their types to P1; (iii) A3 sends message $\overline{m}_2^3$ to P2; (iv) P2 obtains a payoff of 5.

The arguments for (i) and (iii) are analogous to 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 and the uninformed agent A3 cannot induce P1 to take a decision other than x_{11}. Suppose next that the state is (ω_H, ω_H). Because $\underline{p} \leq \overline{p} < \frac{1}{2}$, $\underline{p} + 8(1-\underline{p}) > 4.5$. From Table 2, and by definition of d_1 and $(\underline{m}_2^1, \underline{m}_2^2)$, it thus follows that, if A2 reports ω_H to P1 and sends $\underline{m}_2^2$ to P2, and if A3 sends $\overline{m}_2^3$ to P2, then A1 best responds by reporting ω_H to P1 and sending $\underline{m}_2^1$ to P2; notice, in particular, that, because $\overline{\omega}_1^1 = \omega_L$, if A1 sends a message $m_1^1 \notin D_1 \times \Omega^1$ to P1, then P1 takes the same decision as if A1 misreported his type. The argument for A2 is identical. These properties, along with the 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* $\tilde{\phi}_j^r : M_j \to \Delta(X_j)$ for principal j implements the same decisions as the recommendation mechanism ϕ_j^r in (2), 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 $\tilde{d}_j : \Omega \to \Delta(X_j)$; again, if some agent i sends a message $m_j^i \notin D_j \times \Omega^i$ to principal j, then $\tilde{\phi}_j^r$ treats this message as if it coincided with some fixed element $(\overline{d}_j, \overline{\omega}_j^i)$ of $D_j \times \Omega^i$, for some $\overline{d}_j \neq d_j^0$.

Observe that P1 can guarantee herself a payoff of -5.5 by committing to play x_{11} with probability 1, regardless of the messages she receives. Thus there is no PBE of G_1^M in which P2's payoff exceeds 5.5. We now construct a family of PBEs of G_1^M, indexed by P2's payoff $v \in (5, 5.5]$, in which P1 and P2 post extended recommendation mechanisms $(\tilde{\phi}_1^r, \tilde{\phi}_2^r)$. As for $\tilde{\phi}_1^r$, we first suppose that the direct mechanism d_1^0 differs from the direct mechanisms d_1^* and d_1 defined by (S.2) and (S.5), which may be recommended by the agents to P1 following a deviation by P2. We then fix $\xi \in [\frac{1}{2}, 1)$ and let

$$\tilde{d}_1(\omega) \equiv \begin{cases} \tilde{x}_1^{\xi} & \text{if } \omega = (\omega_L, \omega_L) \\ \tilde{x}_1^{1-\xi} & \text{otherwise} \end{cases}, \tag{S.6}$$

where $\tilde{x}_1^{\xi} \equiv \xi \delta_{x_{11}} + (1-\xi)\delta_{x_{12}}$ and $\tilde{x}_1^{1-\xi} \equiv (1-\xi)\delta_{x_{11}} + \xi \delta_{x_{12}}$. As for $\tilde{\phi}_2^r$, we let

$$\tilde{d}_2(\omega) \equiv \tilde{x}_2^{\frac{1}{2}}, \quad \omega \in \Omega, \tag{S.7}$$

where $\tilde{x}_2^{\frac{1}{2}} \equiv \frac{1}{2}\delta_{x_{21}} + \frac{1}{2}\delta_{x_{22}}$. Finally, we assume that, for each j, $\overline{\omega}_j^1 = \overline{\omega}_j^2 \equiv \omega_L$, so that, if some agent $i = 1, 2$ sends a message $m_j^i \notin D_j \times \Omega^i$ to principal j, then $\tilde{\phi}_j^r$ treats this message as if agent i reported to principal j to be of type ω_L.

On Path Suppose that P1 and P2 post the recommendation mechanisms $\tilde{\phi}_1^r$ and $\tilde{\phi}_2^r$ defined by (S.6)–(S.7). We claim that, for each $\xi \in [\frac{1}{2}, 1)$, the subgame $(\tilde{\phi}_1^r, \tilde{\phi}_2^r)$ admits a

BNE that satisfies the following properties: (i) each agent recommends to P1 the direct mechanism d_1^0 and recommends to P2 the direct mechanism d_2^* defined in (S.2); (ii) A1 and A2 truthfully report their types to P1 and P2. The corresponding payoff for P2 in the subgame $(\tilde{\phi}_1^r, \tilde{\phi}_2^r)$ is $v = 6 - \xi \in (5, 5.5]$ as ξ varies in $[\frac{1}{2}, 1)$, as desired.

As for (i), the argument is again that unilaterally sending a different recommendation to P1 is of no avail as no agent is pivotal. Consider then (ii). Suppose first that the state is (ω_L, ω_L), 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. By (i), we only need to study A1's reporting decisions. Because $\overline{\omega}_j^1 = \overline{\omega}_j^2 = \omega_L$, 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. (a) If A1 truthfully reports his type to P1 and P2, then P1 implements the lottery $\tilde{x}_1^{\xi}$, P2 takes decision x_{21}, and A1 obtains a payoff of $8\xi + 4.5(1 - \xi)$. (b) If A1 truthfully reports his type to P1 and misreports his type to P2, then P1 implements the lottery $\tilde{x}_1^{\xi}$, P2 takes decision x_{22}, and A1 obtains a payoff of $\xi + 4.5(1 - \xi) < 8\xi + 4.5(1 - \xi)$. (c) 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 $\xi \geq \frac{1}{2}$. (d) 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. Suppose next that the state is (ω_H, ω_H), 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. The claim follows.

Off Path If P1 deviates to some arbitrary standard mechanism $\phi_1 : M_1 \to \Delta(X_1)$, then we require that, in the subgame $(\phi_1, \tilde{\phi}_2^r)$, the agents recommend d_2^0 to P2—again, no agent is pivotal in this recommendation. As a result, the payoff to A1, A2, and A3 is constant and equal to 4.5 in each state no matter P1's decision. We may thus assume that the agents send the same message profile m_1 to P1 in each state; this leads to a payoff of $\frac{1}{2}\{-5\phi_1(m_1)(x_{11}) - 6[1 - \phi_1(m_1)(x_{11})]\} + \frac{1}{2}\{-5[1 - \phi_1(m_1)(x_{11})] - 6\phi_1(m_1)(x_{11})\} = -5.5$ for P1, making her deviation unprofitable.

If P2 deviates to some arbitrary standard mechanism $\phi_2 : M_2 \to \Delta(X_2)$, then we require that, in the subgame $(\tilde{\phi}_1, \phi_2^r)$, the agents' strategies implement the same punishments for P2 as in Step 1 of the proof; this leads to a payoff of 5 for P2, making her deviation unprofitable. 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 metric compact, and every subgame $(\gamma_1, \gamma_2) \in \Gamma_1 \times \Gamma_2$ is finite; moreover, the agents' information structures and payoffs are continuous functions of (γ_1, γ_2). Hence the BNE of the subgame (γ_1, γ_2) form a nonempty compact set $B^*(\gamma_1, \gamma_2)$, and the correspondence $B^* : \Gamma_1 \times \Gamma_2 \twoheadrightarrow \times_{i=1}^3 \Delta(M^i)^{S^i \times \Omega^i}$ 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 the Kuratowski–Ryll-Nardzewski 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 implements the decision x_{11} with probability 1 regardless of the signals she sends to the agents and the messages she receives from them, and that P2 posts the mechanism γ_2^* that implements the random decision $\tilde{x}_2^{\frac{1}{2}}$ introduced in Step 2 of the proof of Lemma 1 regardless of the signals she sends to the agents and the messages she receives from them. Given γ_1^*, P2 obtains a payoff of 5.5 no matter the mechanism she posts; hence she has no profitable deviation. In turn, given γ_2^*, the payoff to A1, A2, and A3 is constant and equal to 4.5 in each state, no matter P1's decision. We may thus assume that, following a deviation by P1 from γ_1^*, the agents send the same message profile m_1 to P1 in each state regardless of the signals they receive from P1. In line with Step 2 of the proof of Lemma 1, this yields a payoff of -5.5 to P1 no matter her signal profile s_1; hence she has no profitable deviation. We conclude that G_1^{SM} admits a PBE.

A Tighter Payoff Bound for P2 We more generally show that, if A1's and A2's types are perfectly correlated and both states (ω_L, ω_L) and (ω_H, ω_H) occur with strictly positive probability, then P2 can guarantee herself a payoff strictly higher than 5 by using private disclosures. Specifically, we construct for each $\sigma \in (\frac{1}{2}, 1)$ 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,

$$\inf_{\gamma_1 \in \Gamma_1} \inf_{\beta \in B^*(\gamma_1, \gamma_2(\sigma))} \sum_{\omega \in \Omega} \sum_{x \in X} \mathbf{P}[\omega] z_{\gamma_1, \gamma_2(\sigma), \beta}(x \,|\, \omega) v_2(x, \omega)$$
$$\geq 5 + \frac{(1-\sigma)\mathbf{P}[(\omega_L, \omega_L)]\mathbf{P}[(\omega_H, \omega_H)]}{1 - \sigma \mathbf{P}[(\omega_L, \omega_L)]}, \qquad \text{(S.8)}$$

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 (S.8).

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

$$0 \leq \varepsilon < \frac{(1-\sigma)\mathbf{P}[(\omega_L, \omega_L)]\mathbf{P}[(\omega_H, \omega_H)]}{1 - \sigma\mathbf{P}[(\omega_L, \omega_L)]}. \tag{S.9}$$

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

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

$$\pi_{11}(\omega_L, \omega_L) \equiv \sigma\pi_{11}(\omega_L, \omega_L, 1) + (1-\sigma)\pi_{11}(\omega_L, \omega_L, 2), \tag{S.10}$$

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

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

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

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

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

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

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

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

or, equivalently,

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

which implies

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

as both $\pi_{11}(\omega_L, \omega_L)$ and $\pi_{12}(\omega_H, \omega_H)$ are at most equal to 1. Notice that (S.9) ensures that the right-hand side of each inequality in (S.12) is strictly positive, and thus can be interpreted as a probability as it is at most equal to 1. Similarly, it follows from (S.10) and from the first inequality in (S.12) that

$$\pi_{11}(\omega_L, \omega_L, 2) \geq 1 - \frac{\varepsilon}{(1 - \sigma)\mathbf{P}[(\omega_L, \omega_L)]}. \tag{S.13}$$

Again, (S.9) ensures that the right-hand side of (S.13) 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 (S.11) and (S.13), that, for any message $m_1^1 \in M_1^1$,

$$\sum_{(m_1^2, m_1^3) \in M_1^2 \times M_1^3} \beta^2(m_1^2 \,|\, \omega_L) \beta^3(m_1^3)\, \phi_1(x_{11} \,|\, m_1^1, m_1^2, m_1^3) \geq 1 - \frac{\varepsilon}{(1 - \sigma)\mathbf{P}[(\omega_L, \omega_L)]}; \tag{S.14}$$

otherwise, by (S.13), 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 (S.14) with respect to the measure $\sigma\beta^1(\cdot \,|\, 1, \omega_H) + (1 - \sigma)\beta^1(\cdot \,|\, 2, \omega_H)$ then yields

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

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

$$4.5\left\{1 - \frac{\varepsilon}{(1-\sigma)\mathbf{P}[(\omega_L,\omega_L)]}\right\} + [\sigma + 8(1-\sigma)]\,\frac{\varepsilon}{(1-\sigma)\mathbf{P}[(\omega_L,\omega_L)]}. \tag{S.15}$$

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 (S.12), he obtains an expected payoff at most equal to

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

Comparing (S.15) and (S.16), 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 (S.9), contradicting the assumption that $\beta \in B^*(\phi_1, \gamma_2(\sigma))$. Thus $\gamma_2(\sigma)$ satisfies (S.8), 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

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

which reduces to $5 + \frac{1}{6}$ when $\mathbf{P}[(\omega_L,\omega_L)] = \mathbf{P}[(\omega_H,\omega_H)] = \frac{1}{2}$. The result follows. ■

Uninformative Signals and Babbling Equilibria To strengthen the result of Lemma 2, Claim S.1 below proves that, if P2 were to use any signal structure that keeps all the agents in the dark, then it would be possible for P1 to post a mechanism inflicting on P2 her minimum feasible payoff of 5. To see this, consider the game G_1^{SM} of Section 3.1 and maintain the assumption 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

$$\sum_{s_2^{-i}\in S_2^{-i}} \sigma_2(s_2^{-i}|s_2^i)\phi_2(x_2|s_2^i, s_2^{-i}, m_2) = \sum_{s_2\in S_2} \sigma_2(s_2)\phi_2(x_2|s_2, m_2) \tag{S.17}$$

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 any information about P2's effective decision rule $\phi_2(\cdot\,|\,s_2, \cdot)$. The following result then holds.

Claim S.1 *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.

Proof of Claim S.1. We can with no loss of generality focus on A1's incentives. Suppose that, in the subgame (ϕ_1^r, γ_2), A2 and A3 play behavior strategies β^2 and β^3 that prescribe the same play for any signals s_2^2 and s_2^3 they may receive from P2, respectively; that is, for each $\omega^2 \in \Omega^2$, $\beta^2(\cdot \mid s_2^2, \omega^2)$ is independent of s_2^2, and similarly $\beta^3(\cdot \mid s_2^3)$ 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 $\omega^1 \in \Omega^1$, $\beta^1(\cdot \mid s_2^1)$ 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 (S.17), any such BNE of the subgame (ϕ_1^r, γ_2) can be straightforwardly turned into a BNE of the subgame $(\phi_1^r, \overline{\phi}_2)$ in which P1 posts the recommendation mechanism ϕ_1^r and P2 posts the standard mechanism $\overline{\phi}_2$ defined by

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

for all $m_2 \in M_2$ and $x_2 \in X_2$. By construction, the same outcome is implemented in either case. Conversely, any BNE of the subgame $(\phi_1^r, \overline{\phi}_2)$ can be 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 $\overline{\phi}_2$ is a standard mechanism, we know from Lemma 1 that the subgame $(\phi_1^r, \overline{\phi}_2)$ admits a BNE in which P2 obtains a payoff of 5. 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 $(s_2, m_2) \in S_2 \times M_2$,

$$\phi_2^*(s_2, m_2) \equiv \begin{cases} \delta_{x_{21}} & \text{if } s_2 \in \{(1,1),(2,2)\} \\ \delta_{x_{22}} & \text{if } s_2 \in \{(1,2),(2,1)\} \end{cases} \tag{S.18}$$

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 $(m_1^1, m_1^2) \in M_1$,

$$\phi_1^*(\emptyset, \emptyset, m_1) \equiv \begin{cases} \delta_{x_{13}} & \text{if } m_1 \in \{(1,\omega_L,1),(2,\omega_L,2)\} \\ \delta_{x_{14}} & \text{if } m_1 \in \{(1,\omega_L,2),(2,\omega_L,1)\} \\ \delta_{x_{12}} & \text{if } m_1 \in \{(1,\omega_H,1),(2,\omega_H,2)\} \\ \delta_{x_{11}} & \text{if } m_1 \in \{(1,\omega_H,2),(2,\omega_H,1)\} \end{cases}, \tag{S.19}$$

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 (S.18)–(S.19) that the outcome (3)–(4) 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

$$\frac{1}{4}\left[\alpha u^1(x_{13}, x_{21}, \omega_L) + (1-\alpha)u^1(x_{14}, x_{22}, \omega_L)\right] \\ + \frac{3}{4}\left[\alpha u^1(x_{12}, x_{21}, \omega_H) + (1-\alpha)u^1(x_{11}, x_{22}, \omega_H)\right] = 3\alpha + 7.5(1-\alpha). \tag{S.20}$$

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

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

which is strictly less than the value in (S.20) 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

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

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 (S.21) 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 (S.21) 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 (S.21) 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

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

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 (S.22) 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 (S.22) 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 (S.22) if $\alpha \geq \frac{2}{5}$.

The above analysis implies that it is a BNE for A1 and A2 to truthfully report their private information to P1 in the subgame (γ_1^*, γ_2^*) if and only if $\alpha = \frac{2}{3}$. In this continuation equilibrium, P1 and P2 obtain their maximum feasible payoff of 10. Hence, 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 a space of admissible standard mechanisms for principal j, endowed with an appropriate σ-field $\mathcal{F}_j$. We refer to Aumann (1961) 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

$$3\alpha_L^\phi + 7.5(1 - \alpha_L^\phi) \geq 5\alpha_H^\phi + 3.5(1 - \alpha_H^\phi). \quad \text{(S.23)}$$

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

$$9\alpha_H^\phi + 5(1 - \alpha_H^\phi) \geq 7\alpha_L^\phi + 9(1 - \alpha_L^\phi). \quad \text{(S.24)}$$

Summing (S.23)–(S.24) yields $\alpha_L^\phi \leq \alpha_H^\phi$, and reinserting this inequality in (S.23)–(S.24), we obtain

$$\alpha_L^\phi \leq \frac{2}{3} \leq \alpha_H^\phi. \tag{S.25}$$

Now, consider the alternative behavior strategy for A2 obtained from his state-dependent candidate equilibrium behavior strategies $\lambda^2(\phi)(\cdot \mid \omega_L)$ and $\lambda^2(\phi)(\cdot \mid \omega_H)$ by de-correlating 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 $\lambda^2(\phi)(\cdot \mid \omega_H)$ and $\lambda^2(\phi)(\cdot \mid \omega_L)$, respectively, and then sending m_1^2 to P1 and $\hat{m}_2^2$ to P2, thus using the distribution $\lambda^2(\phi)(\cdot \mid \omega_H)$ to determine his message to P1 and the distribution $\lambda^2(\phi)(\cdot \mid \omega_L)$ 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:

$$\begin{aligned}
\Pr(x_{11}, x_{21}) + \Pr(x_{11}, x_{22}) &= 1 - \alpha_H^\phi, \\
\Pr(x_{12}, x_{21}) + \Pr(x_{12}, x_{22}) &= \alpha_H^\phi, \\
\Pr(x_{11}, x_{21}) + \Pr(x_{12}, x_{21}) &= \alpha_L^\phi, \\
\Pr(x_{11}, x_{22}) + \Pr(x_{12}, x_{22}) &= 1 - \alpha_L^\phi.
\end{aligned}$$

It is easy to check that this system has not full rank, and admits a continuum of solutions indexed by $p \equiv \Pr(x_{11}, x_{21})$, 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 (S.24), we obtain $\alpha_L^\phi \geq \alpha_H^\phi$ and hence $\alpha_L^\phi = \alpha_H^\phi = \frac{2}{3}$ by (S.25), 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 de-correlating the two principals' decisions, but this time using only the candidate equilibrium behavior strategy $\lambda^2(\phi)(\cdot \mid \omega_H)$. 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 $\lambda^2(\phi)(\cdot \mid \omega_H)$ 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

$$(\phi_1(m_1^1, m_1^2), \phi_2(m_2^1, \hat{m}_2^2)) \in \{\delta_{(x_{11}, x_{22})}, \delta_{(x_{12}, x_{21})}\} \tag{S.26}$$

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 (S.26) 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

$$(\phi_1(m_1^1, m_1^2), \phi_2(m_2^1, m_2^2)) = (\phi_1(m_1^1, \hat{m}_1^2), \phi_2(m_2^1, \hat{m}_2^2)) \tag{S.27}$$

for $\lambda^1(\phi)$-almost every m_1. Because ϕ_1 and ϕ_2 are measurable, we can then conclude from Fubini's theorem (Bogachev (2007, Theorem 3.4.4)) that (S.27) indeed 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, (S.27) 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 $(m_1^2, m_2^2) \mapsto (\phi_1(m_1^1, m_1^2), \phi_2(m_2^1, m_2^2))$ is constant over a set of $\lambda^2(\phi)(\cdot\,|\,\omega_H)$-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. ■

S.2 Proofs for Section 4.2

Proof of Theorem 1. The proof consists of six steps.

Step 1: Additional Sampling Variables First, assume that every principal j, in addition to drawing ξ_j uniformly from $\Xi_j \equiv [0,1]$, also draws ξ_j^i uniformly from $\Xi_j^i \equiv [0,1]$, one for every agent i, with all the draws made independently. As we explain below, these second draws are used to generate a new random variable jointly controlled by the principals that replicates the original sampling variable ξ^i used by every agent i in $G^{\hat{S}\hat{M}}$. For all i and j, we then let $\kappa_j^i : \Xi_j \times \hat{S}_j^i \times \Xi_j^i \to \mathring{S}_j^i$ and $\rho_j^i : \hat{M}_j^i \to \mathring{M}_j^i$ be two Borel-measurable embeddings[1] such that

$$\operatorname{Im} \rho_j^i \cap \left\{ (\omega^i, (\mathring{s}_k^i)_{k \neq j}) \in \mathring{M}_j^i : \mathring{s}_k^i \in \operatorname{Im} \kappa_k^i \text{ for all } k \neq j \right\} = \emptyset. \tag{S.28}$$

The existence of such embeddings, which are necessarily non-surjective because of (S.28), follows from the fact that $\Xi_j \times \hat{S}_j^i \times \Xi_j^i = [0,1]^3$, $\hat{M}_j^i = [0,1]$, and $\mathring{S}_j^i = [0,1]$ are all uncountable Polish spaces;[2] we can with no loss of generality assume that $\operatorname{Im} \kappa_j^i = \mathcal{I}_\kappa$ and $\operatorname{Im} \rho_j^i = \Omega^i \times \mathcal{I}_\rho^{J-1}$, where $\mathcal{I}_\kappa$ and $\mathcal{I}_\rho$ are disjoint compact subintervals of $[0,1]$. We denote by $(\kappa_j^i)^{-1}$ and $(\rho_j^i)^{-1}$ the preimage mappings of κ_j^i and ρ_j^i over $\operatorname{Im} \kappa_j^i$ and $\operatorname{Im} \rho_j^i$, respectively. In particular, there exist Borel-measurable injections $a_j^i : \operatorname{Im} \kappa_j^i \to \Xi_j$, $b_j^i : \operatorname{Im} \kappa_j^i \to \hat{S}_j^i$, and $c_j^i : \operatorname{Im} \kappa_j^i \to \Xi_j^i$ such that $(\kappa_j^i)^{-1} = (a_j^i, b_j^i, c_j^i)$.

We are now ready to specify the p-truthful PBE $(\mathring{\mu}^*, \mathring{\lambda}^*)$ of $G^{\mathring{S}\mathring{M}}$ corresponding to the PBE $(\hat{\mu}^*, \hat{\lambda}^*)$ of $G^{\hat{S}\hat{M}}$. We first describe the principals' and the agents' strategies (Steps 2–3). We then argue that the allocation induced by $(\mathring{\mu}^*, \mathring{\lambda}^*)$ in $G^{\mathring{S}\mathring{M}}$ is the same as the one induced by $(\hat{\mu}^*, \hat{\lambda}^*)$ in $G^{\hat{S}\hat{M}}$ (Step 4). Finally, we show that $(\mathring{\mu}^*, \mathring{\lambda}^*)$ satisfies all the equilibrium requirements in $G^{\mathring{S}\mathring{M}}$ (Steps 5–6).

Step 2: Principals' Strategies Every principal j posts with probability 1 a mechanism

[1] That is, injections that yield Borel isomorphisms between their domains and their images.

[2] Indeed, by Kuratowski's theorem (see, for instance, Kechris (1995, Theorem 15.6)), any uncountable standard Borel space—that is, any uncountable Polish space endowed with the Borel σ-field generated by a compatible metric—is Borel-isomorphic to $([0,1], \mathcal{B}([0,1]))$.

$\mathring{\gamma}_j^* \equiv (\mathring{\sigma}_j^*, \mathring{\phi}_j^*)$ defined as follows.

We start with the distribution $\mathring{\sigma}_j^*$. Principal j first draws ξ_j and all the $(\xi_j^i)_{i=1}^I$ uniformly from $[0,1]$, with all the draws made independently. She then uses the draw ξ_j along with the function $\hat{\mu}_j^* : \Xi_j \to \hat{\Gamma}_j$ describing her equilibrium mixed strategy in $G^{\hat{S}\hat{M}}$ to identify the mechanism $\hat{\mu}_j^*(\xi_j) = (\hat{\sigma}_j^{*\xi_j}, \hat{\phi}_j^{*\xi_j})$ that she would have posted in $G^{\hat{S}\hat{M}}$. Next, principal j draws the signals $\hat{s}_j$ from $\hat{S}_j$ using the distribution $\hat{\sigma}_j^{\xi_j}$. Finally, she uses the embeddings κ_j^i described above to map each $(\xi_j, \hat{s}_j^i, \xi_j^i)$ into the corresponding signal $\mathring{s}_j^i = \kappa_j^i(\xi_j, \hat{s}_j^i, \xi_j^i)$ to disclose to every agent i in $G^{\mathring{S}\mathring{M}}$. Formally, the distribution $\mathring{\sigma}_j^*$ of $\mathring{s}_j$ is thus the push-forward of the measure $[\mathrm{d}\xi_j \otimes (\delta_{\xi_j} \otimes \hat{\sigma}_j^{\xi_j})] \otimes \bigotimes_{i=1}^I \mathrm{d}\xi_j^i$ by the mapping $(\kappa_j^i)_{i=1}^I : \Xi_j \times \hat{S}_j \times \bigtimes_{i=1}^I \Xi_j^i : (\xi_j, \hat{s}_j, (\xi_j^i)_{i=1}^I) \mapsto (\kappa_j^i(\xi_j, \hat{s}_j^i, \xi_j^i))_{i=1}^I$; that is, for each $A \in \mathcal{B}([0,1])$,

$$\begin{aligned} \mathring{\sigma}_j^*(A) &\equiv (\kappa_j^i)_{i=1}^I \sharp \left[\mathrm{d}\xi_j \otimes \left(\delta_{\xi_j} \otimes \hat{\sigma}_j^{\xi_j}\right)\right] \otimes \bigotimes_{i=1}^I \mathrm{d}\xi_j^i (A) \\ &\equiv \left[\mathrm{d}\xi_j \otimes \left(\delta_{\xi_j} \otimes \hat{\sigma}_j^{\xi_j}\right)\right] \otimes \bigotimes_{i=1}^I \mathrm{d}\xi_j^i \left(((\kappa_j^i)_{i=1}^I)^{-1}(A)\right), \end{aligned} \tag{S.29}$$

where δ_{ξ_j} is the Dirac measure centered on ξ_j. Notice that the set $\operatorname{supp} \mathring{\sigma}_j^* \cap \operatorname{Im} (\kappa_j^i)_{i=1}^I$ has $\mathring{\sigma}_j^*$-measure 1 and that, given the above construction, we can assume that

(A) every profile of signals $\mathring{s}_j$ sent by principal j to the agents belongs to $\operatorname{supp} \mathring{\sigma}_j^* \cap \operatorname{Im} (\kappa_j^i)_{i=1}^I$.

Next, consider the extended decision rule $\mathring{\phi}_j^*$. Let $\mathring{m}_j \equiv (\omega^i, \mathring{s}_{-j}^i)_{i=1}^I$ denote an arbitrary profile of messages received by principal j in $G^{\mathring{S}\mathring{M}}$, with $\mathring{s}_{-j}^i \equiv (\mathring{s}_k^i)_{k \neq j}$ for all i. We distinguish two cases.

Case 1 First, take any $(\mathring{s}_j, \mathring{m}_j)$ such that

(B) for each i, $\mathring{m}_j^i = (\omega^i, \mathring{s}_{-j}^i)$ is such that $\mathring{s}_k^i \in \operatorname{Im} \kappa_k^i$ for all $k \neq j$.

Condition (B) states that the messages principal j received from the agents are such that the signals $\mathring{s}_k^i$ reported by every agent i are consistent with the embeddings κ_k^i used by every principal $k \neq j$ to encode the information $(\xi_k, \hat{s}_k^i, \xi_k^i)$ into $\mathring{s}_k^i$.

Recall that, for all i, j, ξ^i, $\hat{\gamma}$, $\hat{s}^i$, and ω^i, $\hat{\lambda}_j^{*i,\xi^i}(\hat{\gamma}, \hat{s}^i, \omega^i)$ is the message agent i of type ω^i sends in equilibrium to principal j in $G^{\hat{S}\hat{M}}$, given the profile of mechanisms $\hat{\gamma}$, the profile of signals $\hat{s}^i$ received by agent i, and the realization ξ^i of his sampling variable. Now, condition (A) ensures that $\xi_j = a_j(\mathring{s}_j) \equiv a_j^i(\mathring{s}_j^i)$ is independent of i. Thus

$$\begin{aligned} &\mathring{\phi}_j^*(\mathring{s}_j, \mathring{m}_j) \\ &\quad \equiv \hat{\phi}_j^{*a_j(\mathring{s}_j)} \left((b_j^i(\mathring{s}_j^i))_{i=1}^I, \left(\hat{\lambda}_j^{*i,\left\{\sum_{k=1}^J c_k^i(\mathring{s}_k^i)\right\}} \left((\hat{\mu}_k^*(a_k^i(\mathring{s}_k^i)))_{k=1}^J, (b_k^i(\mathring{s}_k^i))_{k=1}^J, \omega^i \right) \right)_{i=1}^I \right), \end{aligned} \tag{S.30}$$

where $\{\cdot\}$ is the fractional part operator, is well-defined for all $(\mathring{s}_j, \mathring{m}_j)$ satisfying conditions (A)–(B). That is, the extended decision rule $\mathring{\phi}_j^*$ implements the same decisions as the extended decision rule $\hat{\phi}_j^{*a_j(\mathring{s}_j)}$ implements in $G^{\hat{S}\hat{M}}$ whenever principal j discloses the signals $\hat{s}_j = (b_j^i(\mathring{s}_j^i))_{i=1}^I$ to the agents and every agent i sends the message

$$\hat{m}_j^i = \hat{\lambda}_j^{*i,\left\{\sum_{k=1}^J c_k^i(\mathring{s}_k^i)\right\}}\big((\hat{\mu}_k^*(a_k^i(\mathring{s}_k^i)))_{k=1}^J, (b_k^i(\mathring{s}_k^i))_{k=1}^J, \omega^i\big) \tag{S.31}$$

to principal j.

Case 2 Second, take any $(\mathring{s}_j, \mathring{m}_j)$ such that condition (A) is satisfied but condition (B) is violated. The decision implemented by the mechanism $\mathring{\phi}_j^*$ is then given by the same expression as in (S.30) after replacing the message (S.31) of any agent i for whom $\mathring{m}_j^i = (\omega^i, \mathring{s}_{-j}^i)$ is such that $\mathring{s}_k^i \notin \operatorname{Im} \kappa_k^i$ for some $k \neq j$ with the message $\hat{m}_j^i = (\rho_j^i)^{-1}(\mathring{m}_j^i)$ if $\mathring{m}_j^i \in \operatorname{Im} \rho_j^i$ and with an arbitrarily fixed element $\hat{m}_{j,0}^i$ of $\hat{M}_j^i$ otherwise.

This completes the description of every principal j's candidate equilibrium mechanism $\mathring{\gamma}_j^* = (\mathring{\sigma}_j^*, \mathring{\phi}_j^*)$ in $G^{\mathring{S}\mathring{M}}$.

Notice that, because the functions a_j^i, b_j^i, c_j^i, and $(\rho_j^i)^{-1}$ are Borel-measurable for all i and j, the measurability restrictions imposed in Section 4.1 on the functions $\hat{\phi}_j^*$, $\hat{\lambda}^{*i}$, and $\hat{\mu}_j^*$ imply that $\mathring{\phi}_j^*$ is Borel-measurable, as requested. We let $\mathring{\gamma}^* \equiv (\mathring{\gamma}_j^*)_{j=1}^J$ be the profile of principals' candidate equilibrium mechanisms in $G^{\mathring{S}\mathring{M}}$.

Step 3: Agents' Strategies For all i and j, let us fix a Borel isomorphism $\tau_j^i : \hat{M}_j^i \to \mathring{M}_j^i$ (see Footnote 2). Then, to every mechanism $\mathring{\gamma}_j = (\mathring{\sigma}_j, \mathring{\phi}_j)$ of principal j in $G^{\mathring{S}\mathring{M}}$, we can associate a mechanism $\chi_j(\mathring{\gamma}_j) = (\hat{\sigma}_j, \hat{\phi}_j)$ in $G^{\hat{S}\hat{M}}$, defined by $\hat{\sigma}_j \equiv \mathring{\sigma}_j$ and

$$\hat{\phi}_j(\hat{s}_j, \hat{m}_j) \equiv \mathring{\phi}_j\Big(\hat{s}_j, \big(\tau_j^i(\hat{m}_j^i)\big)_{i=1}^I\Big) \tag{S.32}$$

for all $\hat{s}_j \in \hat{S}_j$ and $\hat{m}_j \in \hat{M}_j$. By construction, the mapping $\mathring{\gamma}_j \mapsto \chi_j(\mathring{\gamma}_j)$ is injective.

To construct every agent i's strategy $\mathring{\lambda}^{i*}$ in $G^{\mathring{S}\mathring{M}}$, we distinguish three cases according to the profile of mechanisms $\mathring{\gamma} \equiv (\mathring{\gamma}^j)_{j=1}^J$ posted by the principals.

Case 1 If $\mathring{\gamma} = \mathring{\gamma}^*$, that is, every principal j posts her candidate equilibrium mechanism $\mathring{\gamma}_j^*$, then agent i truthfully reports $q_j^i \equiv (\omega^i, \mathring{s}_{-j}^i)$ to every principal j; notice that, by condition (A), $\mathring{s}_{-j}^i \in \operatorname{Im}(\kappa_k^i)_{k \neq j}$ for all j.

Case 2 If $\mathring{\gamma}_j \neq \mathring{\gamma}_j^*$ but $\mathring{\gamma}_{-j} = \mathring{\gamma}_{-j}^*$, that is, principal j unilaterally deviates from $\mathring{\gamma}^*$, then every agent i's behavior in $G^{\mathring{S}\mathring{M}}$ is predicated on the behavior that agent i would have followed in the subgame of $G^{\hat{S}\hat{M}}$ in which principal j posts the mechanism $\chi_j(\mathring{\gamma}_j)$ and every principal $k \neq j$ posts the mechanism $\hat{\mu}_k^*(a_k^i(\mathring{s}_k^i))$. That is, we postulate that every agent i of type ω^i draws ξ^i uniformly from $[0,1]$ and then sends to principal j the message

$$\mathring{m}_j^i = \tau_j^i\Big(\hat{\lambda}_j^{*i,\xi^i}\big(\big(\chi_j(\mathring{\gamma}_j), \hat{\mu}_l^*(a_l^i(\mathring{s}_l^i))_{l\neq j}\big), \big(\mathring{s}_j^i, (b_l^i(\mathring{s}_l^i))_{l\neq j}\big), \omega^i\big)\Big), \tag{S.33}$$

and to every principal $k \neq j$ the message

$$\mathring{m}_k^i = \rho_k^i\Big(\hat{\lambda}_j^{*i,\xi^i}\big(\big(\chi_j(\mathring{\gamma}_j), \hat{\mu}_l^*(a_l^i(\mathring{s}_l^i))_{l\neq j}\big), \big(\mathring{s}_j^i, (b_l^i(\mathring{s}_l^i))_{l\neq j}\big), \omega^i\big)\Big). \tag{S.34}$$

Intuitively, by sending messages in $\operatorname{Im}\rho_k^i$ to a non-deviating principal $k \neq j$, agent i tells her to forget about the transformation used to induce truthful-reporting by the agents on path, and to implement the decision that principal k would have implemented off path in $G^{\hat{S}\hat{M}}$.

Case 3 Finally, if more than one principal deviate from $\mathring{\gamma}^*$, then every agent i's behavior in $G^{\mathring{S}\mathring{M}}$ is predicated on the behavior that agent i would have followed in the subgame of $G^{\hat{S}\hat{M}}$ in which every principal j posts the mechanism $\chi_j(\mathring{\gamma}_j)$. That is, we postulate that every agent i of type ω^i draws ξ^i uniformly from $[0,1]$ and then sends the message

$$\mathring{m}_j^i = \tau_j^i\Big(\hat{\lambda}_j^{*i,\xi^i}\big((\chi_j(\mathring{\gamma}_j))_{j=1}^J, \mathring{s}^i, \omega^i\big)\Big) \tag{S.35}$$

to every principal j.

This completes the description of every agent i's candidate equilibrium strategy $\mathring{\lambda}^{*i}$ in $G^{\mathring{S}\mathring{M}}$. Again, because the functions a_j^i, b_j^i, τ_j^i, and ρ_j^i are Borel-measurable for all i and j, the measurability restrictions imposed in Section 4.1 on the functions $\hat{\lambda}^{*i}$ and $\hat{\mu}_j^*$ imply that $\mathring{\lambda}^{*i}$ is $(\mathcal{B}([0,1]) \otimes \hat{\mathbf{\Sigma}} \otimes \hat{\mathcal{S}}^i \otimes 2^{\Omega^i}, \hat{\mathcal{M}}^i)$-measurable, as requested. We let $\mathring{\lambda}^* \equiv (\mathring{\lambda}^{*i})_{i=1}^I$ be the profile of agents' candidate equilibrium strategies in $G^{\mathring{S}\mathring{M}}$.

Step 4: Outcome Equivalence of $(\mathring{\mu}^*, \mathring{\lambda}^*)$ **and** $(\hat{\mu}^*, \hat{\lambda}^*)$ We now claim that the strategy profiles $(\mathring{\mu}^*, \mathring{\lambda}^*)$ and $(\hat{\mu}^*, \hat{\lambda}^*)$ are outcome-equivalent. Indeed, the allocation $z_{\mathring{\mu}^*,\mathring{\lambda}^*}$ induced by $(\mathring{\mu}^*, \mathring{\lambda}^*)$ in $G^{\mathring{S}\mathring{M}}$ satisfies

$$\begin{aligned}
&z_{\mathring{\mu}^*,\mathring{\lambda}^*}(x\,|\,\omega)\\
&= \int_{\times_{j=1}^J \mathring{S}_j} \prod_{j=1}^J \mathring{\phi}_j^*(\mathring{s}_j, ((\omega^i, \mathring{s}_{-j}^i))_{i=1}^I)(x_j) \bigotimes_{j=1}^J \mathring{\sigma}_j^*(\mathrm{d}\mathring{s}_j)\\
&= \int_{\times_{j=1}^J \mathring{S}_j} \prod_{j=1}^J \hat{\phi}_j^{*a_j(\mathring{s}_j)}\left((b_j^i(\mathring{s}_j^i))_{i=1}^I, \left(\hat{\lambda}_j^{*i,\{\sum_{k=1}^J c_k^i(\mathring{s}_k^i)\}}\big((\hat{\mu}_k^*(a_k(\mathring{s}_k)))_{k=1}^J, (b_k^i(\mathring{s}_k^i))_{k=1}^J, \omega^i\big)\right)_{i=1}^I\right)(x_j)\\
&\qquad \bigotimes_{j=1}^J \left[(\kappa_j^i)_{i=1}^I \sharp \left[\mathrm{d}\xi_j \otimes \big(\delta_{\xi_j} \otimes \hat{\sigma}_j^{\xi_j}\big)\right] \otimes \bigotimes_{i=1}^I \mathrm{d}\xi_j^i\right](\mathrm{d}\mathring{s}_j)\\
&= \int_{\times_{j=1}^J \Xi_j} \int_{\times_{j=1}^J \hat{S}_j} \int_{\times_{j=1}^J \times_{i=1}^I \Xi_j^i} \prod_{j=1}^J \hat{\phi}_j^{*\xi_j}\left(\hat{s}_j, \left(\hat{\lambda}_j^{*i,\{\sum_{k=1}^J \xi_k^i\}}\big((\hat{\mu}_k^*(\xi_k))_{k=1}^J, \hat{s}^i, \omega^i\big)\right)_{i=1}^I\right)(x_j)
\end{aligned}$$

$$\bigotimes_{j=1}^{J}\left[\left[\mathrm{d}\xi_j \otimes \left(\delta_{\xi_j} \otimes \hat{\sigma}_j^{\xi_j}\right)\right] \otimes \bigotimes_{i=1}^{I} \mathrm{d}\xi_j^i\right]$$

$$= \int_{\times_{j=1}^{J} \Xi_j} \int_{\times_{j=1}^{J} \hat{S}_j} \int_{\times_{i=1}^{I} \Xi^i} \prod_{j=1}^{J} \hat{\phi}_j^{*\xi_j}\left(\hat{s}_j, \left(\hat{\lambda}_j^{*i,\xi^i}\left((\hat{\mu}_k^*(\xi_k))_{k=1}^J, \hat{s}^i, \omega^i\right)\right)_{i=1}^{I}\right)(x_j)$$

$$\bigotimes_{i=1}^{I} \mathrm{d}\xi^i \bigotimes_{j=1}^{J} \hat{\sigma}_j^{\xi_j}(\mathrm{d}\hat{s}_j) \bigotimes_{j=1}^{J} \mathrm{d}\xi_j$$

$$= z_{\hat{\mu}^*,\hat{\lambda}^*}(x \,|\, \omega) \tag{S.36}$$

for all $(x, \omega) \in X \times \Omega$, where the first equality follows from the fact that every agent i reports $q_j^i \equiv (\omega^i, \mathring{s}_{-j}^i)$ truthfully to every principal j, the second equality follows from (S.29)–(S.30) along with the fact that $a_j(\mathring{s}_j) = a_j^i(\mathring{s}_j^i)$ is independent of i for all j and $\mathring{\sigma}_j$-almost every $\mathring{s}_j^i$, the third equality follows from the change-of-variable formula for push-forward measures (Bogachev (2007, Theorem 3.6.1)), the fourth equality follows from the fact that the random variable $\{\sum_{k=1}^J \xi_k^i\}$ jointly controlled by the principals is uniformly distributed over $[0, 1]$ (see, for instance, Peters and Troncoso-Valverde (2013, Appendix A.1)), and the last equality follows from (5). Thus $z_{\mathring{\mu}^*,\mathring{\lambda}^*} = z_{\hat{\mu}^*,\hat{\lambda}^*}$, as claimed.

Step 5: Equilibrium Properties of $\mathring{\lambda}^*$ We distinguish three cases. In each case, we study the incentives of some agent i, assuming that the other agents stick to their candidate equilibrium strategies $\mathring{\lambda}^{*-i}$.

Case 1 Suppose first that $\mathring{\gamma} = \mathring{\gamma}^*$. If agent i does not deviate from $\mathring{\lambda}^{*i}$, then the allocation implemented in $G^{\hat{S}\hat{M}}$ is given by (S.36). Now, according to Step 2, agent i may deviate in two ways from $\mathring{\lambda}^{*i}$ vis-à-vis any principal j. First, he may send to principal j a message $\tilde{\mathring{m}}_j^i \equiv (\tilde{\omega}^i, \tilde{\mathring{s}}_{-j}^i)$ such that condition (B) of Case 1 of Step 2 is satisfied. According to (S.30), this would amount, in $G^{\hat{S}\hat{M}}$, to play vis-à-vis principal j as if (i) he had observed mechanisms different from $(\hat{\mu}_k^*(a_k^i(\mathring{s}_k^i)))_{k=1}^J$, or (ii) he had received signals different from $(b_k^i(\mathring{s}_k^i))_{k=1}^J$, or (iii) he had observed a realization of the sampling variable different from $\{\sum_j^J c_j^i(\mathring{s}_j^i)\}$, or (iv) he had a type different from ω^i. Second, he may send to principal j a message $\tilde{\mathring{m}}_j^i \equiv (\tilde{\omega}^i, \tilde{\mathring{s}}_{-j}^i)$ such that condition (B) of Case 1 of Step 2 is not satisfied. According to (S.30) and Case 2 of Step 2, this would amount, in $G^{\hat{S}\hat{M}}$, to send to principal j the message $\hat{m}_j^i = (\rho_j^i)^{-1}(\tilde{\mathring{m}}_j^i)$ or the message $\hat{m}_{j,0}^i$. Because all these options are available in $G^{\hat{S}\hat{M}}$, and because, in the first case, it is inconsequential for agent i whether the sampling variable ξ^i is drawn by himself or by averaging over the components $(c_j^i(\mathring{s}_j^i))_{j=1}^I$ received from the principals, we conclude from the optimality of agents i's equilibrium strategy $\hat{\lambda}^{*i}$ in $G^{\hat{S}\hat{M}}$ that, when $\mathring{\gamma} = \mathring{\gamma}^*$ and the other agents follow their candidate equilibrium strategies $\mathring{\lambda}^{*-i}$ in $G^{\mathring{S}\mathring{M}}$, agent i can do no better than reporting $q_j^i = (\omega^i, \mathring{s}_{-j}^i)$ truthfully to every principal j.

Case 2 Suppose next that $\mathring{\gamma}_j \equiv (\mathring{\sigma}_j, \mathring{\phi}_j) \neq \mathring{\gamma}_j^*$ but $\mathring{\gamma}_{-j} = \mathring{\gamma}_{-j}^*$ for some j. We first claim that the strategy profiles $((\mathring{\gamma}_j, \mathring{\mu}_{-j}^*), \mathring{\lambda}^*)$ in $G^{\mathring{S}\mathring{M}}$ and $((\chi_j(\mathring{\gamma}_j), \hat{\mu}_{-j}^*), \hat{\lambda}^*)$ in $G^{\hat{S}\hat{M}}$ are outcome-equivalent. Indeed, letting $\chi_j(\mathring{\gamma}_j) \equiv (\hat{\sigma}_j, \hat{\phi}_j)$, where $\hat{\sigma}_j \equiv \mathring{\sigma}_j$ and $\hat{\phi}_j$ is given by (S.32), the allocation $z_{(\mathring{\gamma}_j, \mathring{\mu}_{-j}^*), \mathring{\lambda}^*}$ induced by $((\mathring{\gamma}_j, \mathring{\mu}_{-j}^*), \mathring{\lambda}^*)$ in $G^{\mathring{S}\mathring{M}}$ satisfies

$$
\begin{aligned}
& z_{(\mathring{\gamma}_j, \mathring{\mu}_{-j}^*), \mathring{\lambda}^*}(x \,|\, \omega) \\
&= \int_{\times_{l=1}^J \mathring{S}_l} \int_{\times_{i=1}^I \Xi^i} \mathring{\phi}_j \Bigg(\mathring{s}_j, \Big(\tau_j^i \Big(\hat{\lambda}_j^{*i,\xi^i} \big((\chi_j(\mathring{\gamma}_j), \hat{\mu}_l^*(a_l^i(\mathring{s}_l^i))_{l \neq j}), (\mathring{s}_j^i, (b_l^i(\mathring{s}_l^i))_{l \neq j}), \omega^i \big) \Big) \Big)_{i=1}^I \Bigg)(x_j) \\
& \quad \prod_{k \neq j} \mathring{\phi}_k^* \Bigg(\mathring{s}_k, \Big(\rho_k^i \Big(\hat{\lambda}_j^{*i,\xi^i} \big((\chi_j(\mathring{\gamma}_j), \hat{\mu}_l^*(a_l^i(\mathring{s}_l^i))_{l \neq j}), (\mathring{s}_j^i, (b_l^i(\mathring{s}_l^i))_{l \neq j}), \omega^i \big) \Big) \Big)_{i=1}^I \Bigg)(x_k) \\
& \quad \bigotimes_{i=1}^I \mathrm{d}\xi^i \bigotimes_{k \neq j} \mathring{\sigma}_k^*(\mathrm{d}\mathring{s}_k) \otimes \mathring{\sigma}_j(\mathrm{d}\mathring{s}_j) \\
&= \int_{\times_{l=1}^J \mathring{S}_l} \int_{\times_{i=1}^I \Xi^i} \hat{\phi}_j \Bigg(\mathring{s}_j, \Big(\hat{\lambda}_j^{*i,\xi^i} \big((\chi_j(\mathring{\gamma}_j), \hat{\mu}_l^*(a_l^i(\mathring{s}_l^i))_{l \neq j}), (\mathring{s}_j^i, (b_l^i(\mathring{s}_l^i))_{l \neq j}), \omega^i \big) \Big)_{i=1}^I \Bigg)(x_j) \\
& \quad \prod_{k \neq j} \hat{\phi}_k^{*a_k(\mathring{s}_k)} \Bigg((b_k^i(\mathring{s}_k^i))_{i=1}^I, \Big(\hat{\lambda}_j^{*i,\xi^i} \big((\chi_j(\mathring{\gamma}_j), \hat{\mu}_l^*(a_l^i(\mathring{s}_l^i))_{l \neq j}), (\mathring{s}_j^i, (b_l^i(\mathring{s}_l^i))_{l \neq j}), \omega^i \big) \Big)_{i=1}^I \Bigg)(x_k) \\
& \quad \bigotimes_{i=1}^I \mathrm{d}\xi^i \bigotimes_{k \neq j} \left\{ (\kappa_k^i)_{i=1}^I \,\sharp\, \big[\mathrm{d}\xi_k \otimes \big(\delta_{\xi_k} \otimes \hat{\sigma}_k^{\xi_k} \big) \big] \otimes \bigotimes_{i=1}^I \mathrm{d}\xi_k^i \right\} (\mathrm{d}\mathring{s}_k) \otimes \mathring{\sigma}_j(\mathrm{d}\mathring{s}_j) \\
&= \int_{\times_{k \neq j}^J \Xi_k} \int_{\times_{j=1}^J \hat{S}_j} \int_{\times_{i=1}^I \Xi^i} \hat{\phi}_j \Bigg(\hat{s}_j, \Big(\hat{\lambda}_j^{*i,\xi^i} \big((\chi_j(\mathring{\gamma}_j), \hat{\mu}_l^*(\xi_l)_{l \neq j}), \hat{s}^i, \omega^i \big) \Big)_{i=1}^I \Bigg)(x_j) \\
& \quad \prod_{k \neq j} \hat{\phi}_k^{*\xi_k} \Bigg(\hat{s}_k, \Big(\hat{\lambda}_k^{*i,\xi^i} \big((\chi_j(\mathring{\gamma}_j), (\hat{\mu}_l^*(\xi_l))_{l \neq j}), \hat{s}^i, \omega^i \big) \Big)_{i=1}^I \Bigg)(x_k) \\
& \quad \bigotimes_{i=1}^I \mathrm{d}\xi^i \bigotimes_{k \neq j} \hat{\sigma}_k^{\xi_k}(\mathrm{d}\hat{s}_k) \otimes \hat{\sigma}_j(\mathrm{d}\hat{s}_j) \bigotimes_{k \neq j} \mathrm{d}\xi_k \\
&= z_{(\chi_j(\mathring{\gamma}_j), \hat{\mu}_{-j}^*), \hat{\lambda}^*}(x \,|\, \omega)
\end{aligned}
\tag{S.37}
$$

for all $(x, \omega) \in X \times \Omega$, where the first equality follows from (S.33)–(S.34), the second equality follows from (S.29)–(S.30), (S.32), and the construction of every principal k's mechanism in Case 2 of Step 2, along with the fact that, for each $k \neq j$, $(\rho_k^i)^{-1} \circ \rho_k^i = \mathrm{Id}_{\hat{M}_k^i}$ as ρ_k^i is injective, and that $a_k(\mathring{s}_k) = a_k^i(\mathring{s}_k^i)$ is independent of i and $\mathring{\sigma}_k$-almost every $\mathring{s}_k^i$, the third equality follows from the change-of-variable formula for push-forward measures and the fact that $\hat{\sigma}_j = \mathring{\sigma}_j$, and the last equality follows from (5). Thus $z_{(\mathring{\gamma}_j, \mathring{\mu}_{-j}^*), \mathring{\lambda}^*} = z_{(\chi_j(\mathring{\gamma}_j), \hat{\mu}_{-j}^*), \hat{\lambda}^*}$, as claimed.

If agent i does not deviate from $\mathring{\lambda}^{*i}$ following principal j's unilateral deviation to $\mathring{\gamma}_j$, then the allocation implemented in $G^{\mathring{S}\mathring{M}}$ is given by (S.37). The proof that agent i cannot be better off deviating from (S.33) vis-à-vis the deviating principal j, or from (S.34) vis-à-vis one or several of the non-deviating principals $k \neq j$ then proceeds as in Case 1. Specifically, any such deviation would amount, in $G^{\hat{S}\hat{M}}$, to play as if (i) he had observed mechanisms different from $(\chi_j(\mathring{\gamma}_j), (\hat{\mu}^*_k(a^i_k(\mathring{s}^i_k)))_{k \neq j})$, or (ii) he had received signals different from $(\mathring{s}^i_j, (b^i_k(\mathring{s}^i_k))_{k \neq j})$, or (iii) he had observed a realization of his sampling variable different from ξ^i, or (iv) he had a type different from ω^i. Because all these options are available in $G^{\hat{S}\hat{M}}$, we conclude from the optimality of agents i's equilibrium strategy $\hat{\lambda}^{*i}$ in $\hat{G}^{SM}$ that, when the other agents follow their equilibrium strategies $\mathring{\lambda}^{*-i}$ in $G^{\mathring{S}\mathring{M}}$, agent i can do no better than playing according to (S.33)–(S.34) vis-à-vis principals j and $k \neq j$.

Case 3 Suppose finally that $\mathring{\gamma}_j \equiv (\mathring{\sigma}_j, \mathring{\phi}_j) \neq \mathring{\gamma}^*_j$ for at least two principals j. Then, according to (S.32), the subgames $\mathring{\gamma}$ and $(\chi_j(\mathring{\gamma}_j))_{j=1}^J$ of $G^{\mathring{S}\mathring{M}}$ and $G^{\hat{S}\hat{M}}$ are strategically equivalent, up to relabeling of every message from agent i to principal j using the Borel isomorphism τ^i_j. It follows that letting the agents send, in $\mathring{\gamma}$, messages according to the translations (S.35) of their equilibrium messages in $(\chi_j(\mathring{\gamma}_j))_{j=1}^J$ forms a BNE of $\mathring{\gamma}$.

Step 6: Equilibrium Properties of $\mathring{\mu}^*$ There only remains to check that, given the agents' strategy profile $\mathring{\lambda}^*$, the strategy profile $\mathring{\mu}^*$ is a Nash equilibrium in the principals' game. By Step 4, the allocation induced by $\mathring{\mu}^*$ and $\mathring{\lambda}^*$ in $G^{\mathring{S}\mathring{M}}$ coincides with the allocation induced by $\hat{\mu}^*$ and $\hat{\lambda}^*$ in $G^{\hat{S}\hat{M}}$. Moreover, by Case 2 of Step 5, if some principal j unilaterally deviates from $\mathring{\mu}^*$ by posting a mechanism $\mathring{\gamma}_j$, the allocation induced by $(\mathring{\gamma}_j, \mathring{\mu}^*_{-j})$ and $\mathring{\lambda}^*$ in $G^{\mathring{S}\mathring{M}}$ coincides with the allocation induced by $(\chi_j(\mathring{\gamma}_j), \hat{\mu}^*_{-j})$ and $\hat{\lambda}^*$ in $G^{\hat{S}\hat{M}}$. Because $(\hat{\mu}^*, \hat{\lambda}^*)$ is a PBE of $G^{\hat{S}\hat{M}}$, it follows that no principal j can profitably deviate from $\mathring{\mu}^*_j$ in $G^{\mathring{S}\mathring{M}}$ given the other principals' strategy profile $\mathring{\mu}^*_{-j}$ and the agents' strategy profile $\mathring{\lambda}^*$. Thus $(\mathring{\mu}^*, \mathring{\lambda}^*)$ is a PBE of $G^{\mathring{S}\mathring{M}}$ that is outcome-equivalent to $(\hat{\mu}^*, \hat{\lambda}^*)$. Hence the result. ■

S.3 Proofs for Section 4.4

Strategies and Allocations We first rigorously define principals' and agents' strategies in the long-communication game $G^{\hat{S}\hat{M}T}$ with T communication rounds. The exposition closely follows that in Section 4.1 for the short-communication game $G^{\hat{S}\hat{M}}$.

A pure strategy for principal j in $G^{\hat{S}\hat{M}T}$ is simply an element of $\hat{\Gamma}^T_j$. A mixed strategy for principal j in $G^{\hat{S}\hat{M}T}$ is described, given a sampling space $\Xi_j \equiv [0,1]$, by a tuple $\hat{\mu}^T_j \equiv ((\hat{\mathfrak{s}}_j(t))_{t=1}^T, \hat{\mathfrak{f}}^T_j)$, where, for each $t < \infty$, $1 \leq t \leq T$, $\hat{\mathfrak{s}}_j(t) : \Xi_j \times \hat{H}_j(t) \to \Delta(\hat{S}_j(t))$ and $\hat{\mathfrak{f}}^T_j : \Xi_j \times \hat{H}^T_j \to \Delta(X_j)$ are Borel-measurable. Every draw ξ_j from Ξ_j determines for any such t a transition probability $\hat{\sigma}^{\xi_j}_j(t) \equiv \hat{\mathfrak{s}}_j(t)(\xi_j) : \hat{H}_j(t) \to \Delta(\hat{S}_j(t))$ and an extended decision rule $\hat{\phi}^{T\xi_j}_j \equiv \hat{\mathfrak{f}}^T_j(\xi_j, \cdot, \cdot) : \hat{H}^T_j \to \Delta(X_j)$, which together pin down a long-communication mechanism $\hat{\gamma}^{T\xi_j}_j \equiv ((\hat{\sigma}^{\xi_j}_j(t))_{t=1}^T, \hat{\phi}^{T\xi_j}_j) \in \hat{\Gamma}^T_j$ with T communication rounds. In line with Section 4.1, we shall assume that the principals randomize only over countably many extended decision rules and countably many transition probabilities past round 1 on path.

Letting $\hat{S}^i(t) \equiv \times_{j=1}^J \hat{S}^i_j(t)$ and $\hat{M}^i(t) \equiv \times_{j=1}^J \hat{M}^i_j(t)$, endowed with the product σ-fields $\hat{\mathcal{S}}^i(t) = \hat{\mathcal{M}}^i(t) \equiv \bigotimes_{j=1}^J \mathcal{B}([0,1])$, the set of agent i's private histories of signals in round t is $\hat{H}^i(t) \equiv \times_{\tau=1}^t \hat{S}^i(\tau) \times \times_{\tau=1}^{t-1} \hat{M}^i(\tau)$, endowed with the product σ-field $\hat{\mathcal{H}}^i(t) = \bigotimes_{\tau=1}^t \hat{\mathcal{S}}^i(\tau) \otimes \bigotimes_{\tau=1}^{t-1} \hat{\mathcal{M}}^i(\tau)$. Letting $\hat{\Gamma}^T \equiv \times_{j=1}^J \hat{\Gamma}^T_j$, a strategy for agent i is a sequence of functions $\hat{\lambda}^i(t) : \Xi^i \times \hat{\Gamma}^T \times \hat{H}^i(t) \times \Omega^i \to \hat{M}^i(t)$, one for each $t < \infty$, $1 \leq t \leq T$, where $\Xi^i \equiv [0,1]$ is a sampling space for player i, endowed with its Borel σ-field $\mathcal{B}([0,1])$ and Lebesgue measure $\mathrm{d}\xi^i$. We require every such function to be $(\mathcal{B}([0,1]) \otimes \hat{\mathbf{\Sigma}}(1) \otimes \hat{\mathcal{H}}^i(t), \hat{\mathcal{M}}^i(t))$-measurable, where $\hat{\mathbf{\Sigma}}(1) \equiv \bigotimes_{j=1}^J \mathcal{B}(\Delta(\hat{S}_j(1)))$. The allocation induced by the strategies $(\hat{\mu}^T, \hat{\lambda}^T) \equiv ((\hat{\mu}^T_j)_{j=1}^J, ((\hat{\lambda}^i(t))_{t=1}^T)_{i=1}^I)$ is defined by

$$\begin{aligned} &z_{\hat{\mu}^T,\hat{\lambda}^T}(x\,|\,\omega) \\ &\equiv \int_{\times_{j=1}^J \Xi_j} \int_{\times_{i=1}^I \Xi^i} \int_{\times_{t=1}^T \times_{j=1}^J \hat{S}_j(t)} \prod_{j=1}^J \hat{\phi}_j^{T\xi_j}\Big(\big((\hat{s}_j(t), (\hat{\lambda}_j^{i,\xi^i}(t)((\hat{\gamma}_k^{\xi_k})_{k=1}^J, \hat{h}^i(t), \omega^i))_{i=1}^I\big)_{t=1}^T\Big)(x_j) \\ &\qquad\qquad \bigotimes_{j=1}^J \bigotimes_{t=1}^T \hat{\sigma}_j^{\xi_j}(t)(\mathrm{d}\hat{s}_j(t)\,|\,\hat{h}_j(t)) \bigotimes_{i=1}^I \mathrm{d}\xi^i \bigotimes_{j=1}^J \mathrm{d}\xi_j \qquad \text{(S.38)} \end{aligned}$$

for all $(\omega, x) \in \Omega \times X$, where $\hat{h}_j(t)$ is related to the agents' strategies by $\hat{h}_j(1) \equiv \emptyset$ and $\hat{h}_j(t) \equiv (\hat{h}_j(t-1), \hat{s}_j(t-1), (\hat{\lambda}_j^{i,\xi^i}(t-1)((\hat{\gamma}_k^{\xi_k})_{k=1}^J, \hat{h}^i(t-1), \omega^i))_{i=1}^I)$ for all $t < \infty$, $1 < t \leq T$.

Young Classes Before we proceed with the proof of Theorem 2, the following technical reminder about Young classes may be helpful.[3] For any Polish space E and any family B of functions $b : E \to [0,1]$, we let $B\!\uparrow$ and $B\!\downarrow$ be the sets of all functions that are nondecreasing or decreasing limits of sequences of functions in B, respectively. The Young (1911, 1913) hierarchy can then described as follows. First, we let $\underline{B}_0(E) = \overline{B}_0(E) \equiv C(E, [0,1])$, the set of continuous functions from E to $[0,1]$. Then, for every ordinal $1 \leq \alpha < \omega_1$, where ω_1 is the first uncountable ordinal, we let $\underline{B}_\alpha(E) \equiv (\bigcup_{\beta<\alpha} \overline{B}_\beta(E))\!\uparrow$ and $\overline{B}_\alpha(E) \equiv (\bigcup_{\beta<\alpha} \underline{B}_\beta(E))\!\downarrow$. Hence, for instance, $\underline{B}_1(E)$ is the set of lower semicontinuous functions from E to $[0,1]$; likewise, $\overline{B}_1(E)$ is the set of upper semicontinuous functions from E to $[0,1]$. The Young classes $\underline{B}_\alpha(E)$ and $\overline{B}_\alpha(E)$ are related to the classic Baire classes $B_\alpha(E)$ (Kechris (1995, Definition 24.1)) by $B_\alpha(E) \subset \underline{B}_\alpha(E), \overline{B}_\alpha(E) \subset B_{\alpha+1}(E)$ and $B_\alpha(E) = \underline{B}_{\alpha+1}(E) \cap \overline{B}_{\alpha+1}(E)$ for all ordinals $\alpha < \omega_1$. In particular, $\bigcup_{\alpha<\omega_1} \underline{B}_\alpha(E) = \bigcup_{\alpha<\omega_1} \overline{B}_\alpha(E)$ exhausts the set of Borel-measurable functions $f : E \to [0,1]$. A standard result (see, for instance, Šupina and Uhrik (2019, Theorem 3.2)) states that $\underline{B}_\alpha(E)$ and $\overline{B}_\alpha(E)$ can alternatively be characterized as the sets of lower and upper $\mathbf{\Sigma}^0_\alpha(E)$-measurable functions, respectively, that is, functions $\underline{b} : E \to [0,1]$ and $\overline{b} : E \to [0,1]$ such that, for each $r \in [0,1]$, $\underline{b}^{-1}((r,1]) \in \mathbf{\Sigma}^0_\alpha(E)$ and

[3] See Šupina and Uhrik (2019) for an especially readable introduction.

$\overline{b}^{-1}([0,r)) \in \mathbf{\Sigma}^0_\alpha(E)$, where $\mathbf{\Sigma}^0_\alpha(E)$ is the α^{th} additive class in the hierarchy of Borel subsets of E (Kechris (1995, Chapter II, §11).

The following lemma is key to our results.

Lemma S.1 (Cichoń and Morayne (1988, Theorem 2.1)) *For every Polish space E and every ordinal $0 < \alpha < \omega_1$, there exist functions $\underline{U}_\alpha : [0,1] \times E \to [0,1]$ and $\overline{U}_\alpha : [0,1] \times E \to [0,1]$ in $\underline{B}_\alpha([0,1] \times E)$ and $\overline{B}_\alpha([0,1] \times E)$, respectively, that are universal for functions in $\underline{B}_\alpha(E)$ and $\overline{B}_\alpha(E)$, respectively, in the sense that $\underline{B}_\alpha(E) = \{\underline{U}_\alpha(\theta, \cdot) : \theta \in [0,1]\}$ and $\overline{B}_\alpha(E) = \{\overline{U}_\alpha(\theta, \cdot) : \theta \in [0,1]\}$. For all $\underline{f} \in \underline{B}_\alpha(E)$ and $\overline{f} \in \overline{B}_\alpha(E)$, we can therefore write $\underline{f} = \underline{U}_\theta(\cdot, \underline{\theta}_\alpha(\underline{f}))$ and $\overline{f} = \overline{U}_\theta(\cdot, \overline{\theta}_\alpha(\overline{f}))$ for some codes $\underline{\theta}_\alpha(\underline{f})$ and $\overline{\theta}_\alpha(\overline{f})$ in $[0,1]$.*

Proof of Theorem 2. We establish the two parts of the theorem in turn.

Part (i) (Universality) The proof consists of eleven steps.

Step 1: Reduction to $G^{\mathring{S}\mathring{M}T}$ Define, in analogy with $G^{\mathring{S}\mathring{M}}$, the game $G^{\mathring{S}\mathring{M}T}$ with T communication rounds, signal spaces $\mathring{S}^i_j(t) \equiv [0,1]$ for all i, j, and $t < \infty$, $1 \leq t \leq T$, and message spaces $\mathring{M}^i_j(1) \equiv \Omega^i \times \mathring{S}^i_{-j}(1) \equiv \Omega^i \times \bigtimes_{k \neq j} \mathring{S}^i_k(1)$ and $\mathring{M}^i_j(t) \equiv \mathring{S}^i_{-j}(t) \equiv \bigtimes_{k \neq j} \mathring{S}^i_k(t)$ for all i, j, and $t < \infty$, $1 < t \leq T$. For all i and j, we use $q^i_j(1) \equiv (\omega^i, \mathring{s}^i_{-j}(1)) \in \Omega^i \times \mathring{S}^i_{-j}(1)$ to denote agent i's true type and the signals he received from all the principals other than j in round 1, and similarly $q^i_j(t) \equiv \mathring{s}^i_{-j}(t) \in \mathring{S}^i_{-j}(t)$ in rounds $t < \infty$, $1 < t \leq T$. The following definition parallels Definition 3.

Definition S.1 *A SPBE $(\mathring{\mu}^{*T}, \mathring{\lambda}^{*T})$ of $G^{\mathring{S}\mathring{M}T}$ is p-truthful if*

(i) for each j, principal j's strategy $\mathring{\mu}^{*T}_j$ is pure, selecting with probability 1 a mechanism $\mathring{\gamma}^{*T}_j \equiv ((\mathring{\sigma}^*_j(t))^T_{t=1}, \mathring{\phi}^{*T}_j)$;

(ii) on path, that is, in the subgame $\mathring{\gamma}^* \equiv (\mathring{\gamma}^*_j)^J_{j=1}$, every agent i truthfully reports $q^i_j(t)$ to every principal j in any round $t < \infty$, $1 \leq t \leq T$.

Our first result follows along the same lines as Theorem 1.

Lemma S.2 *For any primitive game G and any number of communication rounds T, and for any SPBE $(\hat{\mu}^{*T}, \hat{\lambda}^{*T})$ of $G^{\hat{S}\hat{M}T}$, there exists an outcome-equivalent p-truthful SPBE $(\mathring{\mu}^{*T}, \mathring{\lambda}^{*T})$ of $G^{\mathring{S}\mathring{M}T}$; that is, $z_{\mathring{\mu}^{*T}, \mathring{\lambda}^{*T}} = z_{\hat{\mu}^{*T}, \hat{\lambda}^{*T}}$.*

Step 2: Independent Draws We next show that, for each j, every sequence of transition probabilities $\mathring{\sigma}^T_j \equiv (\mathring{\sigma}_j(t))^T_{t=2}$ for principal j's signals in $G^{\mathring{S}\mathring{M}T}$ past round 1 can be generated by making $I \times (T-1)$ independent draws $\varepsilon^T_j \equiv ((\varepsilon^i_j(t))^I_{i=1})^T_{t=2}$ from the uniform distribution over $[0,1]$. To see this, for each $t<\infty$, $1<t\leq T$, and for any round-t private history $\mathring{h}_j(t)$ of principal j in $G^{\mathring{S}\mathring{M}T}$, let $F_j(t)(\cdot\,|\,\mathring{h}_j(t))$ be the cdf associated with the measure $\mathring{\sigma}_j(t)(\cdot\,|\,\mathring{h}_j(t))$ over $\mathring{S}_j(t)$. By Jiřina's theorem (Bogachev (2007, Theorem 10.4.14)), we can disintegrate this cdf into its marginal $F^1_j(t)(\cdot \mid \mathring{h}_j(t))$ over $\mathring{S}^1_j(t)$, its conditional $F^2_j(t)(\cdot \mid \mathring{h}_j(t),\mathring{s}^1_j(t))$ over $\mathring{S}^2_j(t)$ given $\mathring{s}^1_j(t)$, and so on up to its conditional $F^I_j(t)(\cdot\,|\,\mathring{h}_j(t),\mathring{s}^1_j(t),\dots,\mathring{s}^{I-1}_j(t))$ over $\mathring{S}^I_j(t)$ given $(\mathring{s}^1_j(t),\dots,\mathring{s}^{I-1}_j(t))$; all these functions are jointly measurable with respect to their arguments and the conditioning variables. Given principal j's round-t draws $(\varepsilon^i_j(t))^I_{i=1}$ from the uniform distribution over $[0,1]$, we can then use the generalized inverses of these functions to recursively construct a family of signals $(\mathring{s}^i_j(t))^I_{i=1}$ as follows:

$$\begin{aligned}\mathring{s}^1_j(t) &\equiv F^{1-}_j(t)(\cdot\,|\,\mathring{h}_j(t))(\varepsilon^1_j(t)),\\ \mathring{s}^{i+1}_j(t) &\equiv F^{i+1-}_j(t)(\cdot\,|\,\mathring{h}_j(t),\mathring{s}^1_j(t),\dots,\mathring{s}^i_j(t))(\varepsilon^{i+1}_j(t)), \quad 1\leq i\leq I-1.\end{aligned} \tag{S.39}$$

Given any sequence $\mathring{\sigma}^T_j$, (S.39) enables us to recursively define Borel-measurable functions

$$\eta_{\mathring{\sigma}^T_j}(t):\mathring{S}_j(1)\times[0,1]^{I\times(T-1)}\times\prod_{\tau=1}^{t-1}\mathring{M}_j(\tau)\to\mathring{S}_j(t), \quad t<\infty,\ 1<t\leq T \tag{S.40}$$

that define principal j's signals past round 1 as functions of her round-1 signals $\mathring{s}_j(1)\in\mathring{S}_j(1)\equiv\times^I_{i=1}\mathring{S}^i_j(1)$, her independent draws ε^T_j, and the agents' messages. Notice for future reference that the functions $\eta_{\mathring{\sigma}^T_j}(t)$ do not depend on the distribution $\mathring{\sigma}_j(1)$, reflecting that our construction uses as inputs principal j's round-1 signals $\mathring{s}_j(1)$. Also notice that every function $\eta_{\mathring{\sigma}^T_j}(t)$ does not depend on the draws $(\varepsilon^i_j(\tau))^I_{i=1}$ for $\tau>t$, which we leave in the arguments of $\eta_{\mathring{\sigma}^T_j}(t)$ only to simplify notation; to this end, we also denote by

$$\eta^T_{\mathring{\sigma}^T_j}(\mathring{s}_j(1),\varepsilon^T_j,\mathring{m}^T_j)\equiv\Big(\eta_{\mathring{\sigma}^T_j}(t)\big(\mathring{s}_j(1),\varepsilon^T_j,(\mathring{m}_j(\tau))^{t-1}_{\tau=1}\big)\Big)^T_{t=2} \tag{S.41}$$

the sequence $(\mathring{s}_j(t))^T_{t=2}$ of principal j's signals past round 1, and, given $\mathring{\gamma}^T_j\equiv((\mathring{\sigma}_j(1),\mathring{\sigma}^T_j),\mathring{\phi}^T_j)$, we denote by

$$\mathring{\chi}_{\mathring{\gamma}^T_j}(\mathring{s}_j(1),\varepsilon^T_j,\mathring{m}^T_j)\equiv\mathring{\phi}^T_j\Big(\Big(\mathring{s}_j(1),\eta^T_{\mathring{\sigma}^T_j}\big(\mathring{s}_j(1),\varepsilon^T_j,\mathring{m}^T_j\big)\Big),\mathring{m}^T_j\Big)\in\Delta(X_j) \tag{S.42}$$

the (possibly random) decision taken by principal j given her round-1 signals $\mathring{s}_j(1)$, her sequence of signals $\eta^T_{\mathring{\sigma}^T_j}(\mathring{s}_j(1),\varepsilon^T_j,\mathring{m}^T_j)$ past round 1, and the agents' reports $\mathring{m}^T_j$.

Step 3: Reduction of Dimensionality We now show that all the correlations between the principals' decisions and the agents' private information that may be generated by the agents using the principals' signals past round 1 can be captured by a single random variable, uniformly distributed over $[0,1]$. To see this, let us define a *message plan* for agent i in $G^{\mathring{S}\mathring{M}T}$ as a sequence $\mathring{\pi}^{iT}\equiv(\mathring{\pi}^i(t))^T_{t=1}$ of Borel-measurable functions $\mathring{\pi}^i(t):\Xi^i\times\mathring{H}^i(t)\times\Omega^i\to\mathring{M}^i(t)$, where $\mathring{H}^i(t) \equiv \times_{\tau=1}^{t} \times_{j=1}^{J} \mathring{S}_j^i(\tau)$ and $\mathring{M}^i(t) \equiv \times_{j=1}^{J} \mathring{M}_j^i(t)$ for all $t < \infty$, $1 \leq t \leq T$. Notice that, unlike a strategy, a message plan for an agent only depends on the signals he receives and on his type, and not on the mechanisms offered by the principals nor on his own past messages. The allocation induced in $G^{\mathring{S}\mathring{M}T}$ by a profile $\mathring{\gamma}^T \equiv (\mathring{\gamma}_j^T)_{j=1}^J$ of mechanisms and a profile $\mathring{\pi}^T \equiv (\mathring{\pi}^{iT})_{i=1}^I$ of message plans is given by

$$
\begin{aligned}
&\mathring{z}_{\mathring{\gamma}^T, \mathring{\pi}^T}(x \,|\, \omega) \\
&\equiv \int_{\times_{i=1}^I \Xi^i} \int_{\times_{j=1}^J \mathring{S}_j(1)} \int_{[0,1]^{I \times J \times (T-1)}} \prod_{j=1}^J \mathring{\chi}_{\mathring{\gamma}_j^T}\Big(\mathring{s}_j(1), \varepsilon_j^T, \big((\mathring{\pi}_j^{i,\xi^i}(t)(\mathring{h}^i(t), \omega^i))_{i=1}^I\big)_{t=1}^T\Big)(x_j) \\
&\qquad \bigotimes_{i=1}^I \bigotimes_{j=1}^J \bigotimes_{t=2}^T \mathrm{d}\varepsilon_j^i(t) \bigotimes_{j=1}^J \mathring{\sigma}_j(1)(\mathrm{d}\mathring{s}_j(1)) \bigotimes_{i=1}^I \mathrm{d}\xi^i \qquad \text{(S.43)}
\end{aligned}
$$

for all $(\omega, x) \in \Omega \times X$, where the private histories of signals $(\mathring{h}^i(t))_{t=1}^T$ for all agents i are recursively defined, under $(\mathring{\gamma}^T, \mathring{\pi}^T)$, by

$$
\begin{aligned}
\mathring{h}^i(1) &\equiv (\mathring{s}_j^i(1))_{j=1}^J, \\
\mathring{h}^i(t+1) &\equiv \bigg(\mathring{h}^i(t), \Big(\eta^i_{\mathring{\sigma}_j^T}(t+1)\big(\mathring{s}_j(1), \varepsilon_j^T, \big((\mathring{\pi}_j^{k,\xi^k}(\tau)(\mathring{h}^k(\tau), \omega^k))_{k=1}^I\big)_{\tau=1}^t\big)\Big)_{j=1}^J\bigg), \\
&\quad t < \infty,\ 1 \leq t \leq T-1, \qquad \text{(S.44)}
\end{aligned}
$$

given the round-1 signals $\mathring{s}_j^i(1)$ from every principal j to every agent i, the subsequent independent draws ε_j^T of every principal j, and the type ω^i of every agent i. Denoting by $\mathring{\Gamma}_j^T$ and $\mathring{\Pi}^{iT}$ the spaces of principal j's mechanisms and of agent i's message plans in $G^{\mathring{S}\mathring{M}T}$, respectively, and letting $\Xi_0 \equiv [0,1]$, the following result then holds.

Lemma S.3 *For each j, there exists a function $\varphi_j : \times_{j=1}^J \mathring{\Gamma}_j^T \times \times_{i=1}^I \mathring{\Pi}^{iT} \times \times_{j=1}^J \mathring{S}_j(1) \times \times_{i=1}^I \Xi^i \times \Xi_0 \times \Omega \to \Delta(X_j)$ such that: (i) $\varphi_j(\mathring{\gamma}^T, \mathring{\pi}^T, \cdot, \cdot, \cdot, \cdot)$ is Borel-measurable in $((\mathring{s}_j(1))_{j=1}^J, (\xi^i)_{i=1}^I, \xi_0, \omega) \in \times_{j=1}^J \mathring{S}_j(1) \times \times_{i=1}^I \Xi^i \times \Xi_0 \times \Omega$ for all $(\mathring{\gamma}^T, \mathring{\pi}^T) \in \times_{j=1}^J \mathring{\Gamma}_j^T \times \times_{i=1}^I \mathring{\Pi}^{iT}$; (ii) φ_j depends on $\mathring{\gamma}_{-j}^T \equiv (\mathring{\gamma}_k)_{k \neq j}$ only through the sequence of transition probabilities $\mathring{\sigma}_{-j}^T \equiv (\mathring{\sigma}_k^T)_{k \neq j}$ past round 1; and (iii) φ_j satisfies*

$$
\begin{aligned}
\mathring{z}_{\mathring{\gamma}^T, \mathring{\pi}^T}(x \,|\, \omega) &= \int_{\times_{i=1}^I \Xi^i} \int_{\times_{j=1}^J \mathring{S}_j(1)} \int_{\Xi_0} \prod_{j=1}^J \varphi_j\big(\mathring{\gamma}^T, \mathring{\pi}^T, (\mathring{s}_k(1))_{k=1}^J, (\xi^i)_{i=1}^I, \xi_0, \omega\big)(x_j) \\
&\qquad \mathrm{d}\xi_0 \bigotimes_{j=1}^J \mathring{\sigma}_j(1)(\mathrm{d}\mathring{s}_j(1)) \bigotimes_{i=1}^I \mathrm{d}\xi^i \qquad \text{(S.45)}
\end{aligned}
$$

for all $(\mathring{\gamma}^T, \mathring{\pi}^T, \omega, x) \in \times_{j=1}^J \mathring{\Gamma}_j^T \times \times_{i=1}^I \mathring{\Pi}^{iT} \times \Omega \times X$.

Proof. In light of (S.43), we only need to show that the measure $\bigotimes_{i=1}^I \bigotimes_{j=1}^J \bigotimes_{t=2}^T \mathrm{d}\varepsilon_j^i(t)$ over $[0,1]^{I \times J \times (T-1)}$ is a push-forward of Lebesgue measure $\mathrm{d}\xi_0$ over $[0,1]$. For completeness, we provide a short proof of this more or less standard fact (Steinhaus (1930), Bogachev (2007, Exercise 9.12.50)). Let Y be the set of sequences $(y_n)_{n\geq 1} \in \{0,1\}^{\mathbb{N}}$ such that either $y_n = 1$ for all $n \geq 1$ or $y_n = 0$ for infinitely many $n \geq 1$. Then the restriction $f_{|Y}$ to Y of the mapping $f : \{0,1\}^{\mathbb{N}} \to \Xi_0 : (y_n)_{n\geq 1} \mapsto \sum_{n\geq 1} \frac{1}{2^n} y_n$ is a Borel isomorphism whose inverse maps any number in Ξ_0 to its binary expansion that does not terminate with an infinite sequence of 1's, except for the sequence $(1, 1, \ldots)$ (Dudley (2004, Lemma 13.1.2)). For all $\xi_0 \in \Xi_0$ and $n \geq 1$, denote by $f^{-1}_{|Yn}(\xi_0)$ the n^{th} element of the sequence $f^{-1}_{|Y}(\xi_0)$, so that $\xi_0 = \sum_{n\geq 1} \frac{1}{2^n} f^{-1}_{|Yn}(\xi_0)$. Now, consider a partition $\bigsqcup_{i=1}^{I} \bigsqcup_{j=1}^{J} \bigsqcup_{t=2}^{T} N_{i,j,t}$ of $\mathbb{N} \setminus \{0\}$ such that $N_{i,j,t}$ is countably infinite for all i, j, and t, with corresponding bijection $\nu_{i,j,t} : N_{i,j,t} \to \mathbb{N} \setminus \{0\}$; and, for each $t < \infty$, $1 < t \leq T$, define $e^i_j(t)(\xi_0) \equiv f\big((f^{-1}_{|Y\nu^{-1}_{i,j,t}(n)}(\xi_0))_{n\geq 1}\big)$. To conclude, observe that, if ξ_0 is uniformly distributed over $[0, 1]$, then the sequence $(f^{-1}_{|Yn}(\xi_0))_{n\geq 1}$ is iid, with each component uniformly distributed over $\{0, 1\}$, so that the sequence $(((e^i_j(t)(\xi_0))_{i=1}^{I})_{j=1}^{J})_{t=2}^{T}$ is iid, with each component uniformly distributed over $[0, 1]$. The mapping $\xi_0 \mapsto (((e^i_j(t)(\xi_0))_{i=1}^{I})_{j=1}^{J})_{t=2}^{T}$, which is Borel-measurable as f is continuous and $f^{-1}_{|Y}$ is Borel-measurable, then provides the required push-forward function. The result follows. ■

Intuitively, Lemma S.3 reflects that, using a single draw of a uniformly distributed random variable ξ_0 over $[0, 1]$, one can generate a sequence of iid random variables $(((\varepsilon^i_j(t))_{i=1}^{I})_{j=1}^{J})_{t=2}^{T}$, all uniformly distributed over $[0, 1]$; and, conversely, that, by interlacing the digits of the binary expansions of the numbers in any such sequence, one can generate a uniformly distributed random variable over $[0, 1]$. For our purposes, this shows that all correlations induced by the realizations of private signals past round 1 can be captured by a single random variable. The property that every function φ_j depends neither on $\mathring{\sigma}_{-j}(1) \equiv (\mathring{\sigma}_k(1))_{k\neq j}$ nor on $\mathring{\phi}^T_{-j} \equiv (\mathring{\phi}^T_k)_{k\neq j}$ plays a crucial role in Steps 6 and 9.

Step 4: Emulating ξ_0 in a Non-Manipulable Way To emulate a random variable such as ξ_0 in Lemma S.3 in the short-communication game $G^{\mathring{S}\mathring{M}}$ without creating manipulation opportunities by either the principals or the agents, we proceed as in the proof of Theorem 1. First, every principal j independently draws I auxiliary sampling variables ξ^i_j from the uniform distribution over $[0, 1]$, and then sums them and takes the fractional part of the sum. The resulting random variable $\xi_j \equiv \{\sum_{i=1}^{I} \xi^i_j\}$ is uniformly distributed over $[0, 1]$ and independent of the random variables $(\xi^i_j)_{i=1}^{I}$. Second, the random variables $(\xi_j)_{j=1}^{J}$ are independent, so that the random variable $\xi_0 \equiv \{\sum_{j=1}^{J} \xi_j\}$ is also uniformly distributed over $[0, 1]$ and independent of the random variables $(\xi_j)_{j=1}^{I}$ and $((\xi^i_j)_{i=1}^{I})_{j=1}^{J}$. Notice that, by construction, ξ_0 is not manipulable by any principal j, in the sense that, if she deviates and independently draws $\tilde{\xi}_j$ from an arbitrary distribution over $[0, 1]$, then the resulting random variable $\tilde{\xi}_0 \equiv \{\tilde{\xi}_j + \sum_{k\neq j} \xi_k\}$ is still uniformly distributed over $[0, 1]$ and independent of $\tilde{\xi}_j$.

Similarly, even if each ξ_j is ultimately determined by the auxiliary sampling variables $(\xi_j^i)_{i=1}^I$, each $(\xi_j^i)_{j=1}^J$ provides no information to agent i about ξ_j, nor, a fortiori, about ξ_0. Hence ξ_0 is not manipulable by any agent i either.

Step 5: Indexing Message Plans We now show how to leverage the richness of the agents' message spaces in $G^{\mathring{S}\mathring{M}}$ to enable them to communicate which message plans they would have chosen in $G^{\mathring{S}\mathring{M}T}$. The fact that all message plans for an agent in $G^{\mathring{S}\mathring{M}T}$ can be indexed by an element of $[0,1]$ seems a priori unproblematic: as a message plan consists of $T \leq \aleph_0$ Borel-measurable functions, each of which maps an uncountable standard Borel space into an uncountable standard Borel space, we have that card $\mathring{\Pi}^{iT} = (2^{\aleph_0})^T = 2^{\aleph_0}$ for all i (Dudley (2004, Problem 4.2.8)). However, for reasons that will become clear in Steps 6–7, we need this indexing to be itself performed in a measurable way—hence our focus on SPBE of $G^{\hat{S}\hat{M}T}$ in which, following any profile of mechanisms, all the agents' behavior strategies are of uniformly bounded Young class. We refer the reader to the reminder on Young classes before the proof of Theorem 2 for definitions and notations.

The following assumption, which we will maintain in the remainder of the proof, formalizes our main restriction on the complexity of the agents' best responses in the SPBE $(\hat{\mu}^{*T}, \hat{\lambda}^{*T})$ of $G^{\hat{S}\hat{M}T}$.

Assumption S.1 *There exists an ordinal* $0 < \hat{\alpha} < \omega_1$ *such that, for all* i, j, *and* $t < \infty$, $1 \leq t \leq T$, *the mappings* $(\xi^i, \hat{h}^i(t), \omega^i) \mapsto \hat{\lambda}_j^{*i,\xi^i}(t)(\hat{\gamma}^T, \hat{h}^i(t), \omega^i)$, $\hat{\gamma}^T \in \hat{\Gamma}^t$, *all belong to* $\overline{B}_{\hat{\alpha}}(\Xi^i \times \hat{H}^i(t) \times \Omega^i)$.

The following lemma is then a simple consequence of the fact that the various embeddings involved in the proof of Lemma S.2, which are analogous to the embeddings κ_j^i, ρ_j^i, and τ_j^i introduced in the proof of Theorem 1, are of Young class of at most the first infinite ordinal ω, and that a behavior strategy for each agent in $G^{\mathring{S}\mathring{M}T}$ induces a message plan in a straightforward way.

Lemma S.4 *Let* $(\mathring{\mu}^{*T}, \mathring{\lambda}^{*T})$ *be the p-truthful SPBE of* $G^{\mathring{S}\mathring{M}T}$ *corresponding to the SPBE* $(\hat{\mu}^{*T}, \hat{\lambda}^{*T})$ *of* $G^{\hat{S}\hat{M}T}$ *according to Lemma S.2. Then, under Assumption S.1, there exists an ordinal* $0 < \mathring{\alpha} < \omega_1$ *such that, for all* i, j, *and* $t < \infty$, $1 \leq t \leq T$, *the* (i,j,t)*-components* $(\xi^i, \mathring{h}^i(t), \omega^i) \mapsto \mathring{\pi}_j^{i,\xi^i}(t)(\mathring{h}^i(t), \omega^i)$ *of the message plans induced by the behavior strategies* $\mathring{\lambda}^{*iT}(\mathring{\gamma}^T)$ *in the various subgames* $\mathring{\gamma}^T \in \mathring{\Gamma}^T$ *of* $G^{\mathring{S}\mathring{M}T}$ *all belong to* $\overline{B}_{\mathring{\alpha}}(\Xi^i \times \mathring{H}^i(t) \times \Omega^i)$.

With a slight abuse of terminology, we will say that such message plans are of Young class $\mathring{\alpha}$. Let then $\mathring{\pi}_j^i(t) \in \overline{B}_{\mathring{\alpha}}(\Xi^i \times \mathring{H}^i(t) \times \Omega^i)$ be the (i,j,t)-component of a message plan induced by the strategy $\mathring{\lambda}^{*iT}$ of agent i in some subgame of $G^{\mathring{S}\mathring{M}T}$. By Lemma S.1, there exists a universal function $\overline{U}_{\mathring{\alpha}}^i(t)$ for $\overline{B}_{\mathring{\alpha}}(\Xi^i \times \mathring{H}^i(t) \times \Omega^i)$ such that, for any such $\mathring{\pi}_j^i(t)$, there exists a code $\overline{\theta}^i_{j,\mathring{\alpha}}(t)(\mathring{\pi}^i_j(t)) \in [0,1]$ such that

$$\begin{aligned}&\mathring{\pi}^{i,\xi^i}_j(t)\big(\mathring{h}^i(t),\omega^i\big)\\&\qquad=\overline{U}^i_{\mathring{\alpha}}(t)\big(\big(\xi^i,\mathring{h}^i(t),\omega^i\big),\overline{\theta}^i_{j,\mathring{\alpha}}(t)(\mathring{\pi}^i_j(t))\big),\quad\big(\xi^i,\mathring{h}^i(t),\omega^i\big)\in\Xi^i\times\mathring{H}^i(t)\times\Omega^i.\end{aligned}\tag{S.46}$$

Given a Borel isomorphism $\iota^T_J : [0,1]^{J\times T} \to [0,1]$, the message plan $\mathring{\pi}^{iT} \equiv ((\mathring{\pi}^i_j(t))^J_{j=1})^T_{t=1}$ may finally be encoded as

$$\overline{\theta}^{iT}_{\mathring{\alpha}}(\mathring{\pi}^{iT}) \equiv \iota^T_J\Big(\big(\big(\overline{\theta}^i_{j,\mathring{\alpha}}(t)(\mathring{\pi}^i_j(t))\big)^J_{j=1}\big)^T_{t=1}\Big) \in [0,1].$$

(The same message plan may be encoded in several ways.) Conversely, to every code $\vartheta^{iT}_{\mathring{\alpha}} \in \Theta^{iT}_{\mathring{\alpha}} \equiv [0,1]$ corresponds a message plan $\mathring{\varpi}^{iT}(\vartheta^{iT}_{\mathring{\alpha}})$ of Young class $\mathring{\alpha}$. It is clear from (S.46) that the action of a message plan on an agent's private sample and history of signals is a Borel-measurable function of that sample and that history and of a code for the message plan. In line with Rao (1971), the above construction using universal functions allows us in Steps 6–7 to eschew the problem of admissibility associated to the joint measurability of the evaluation of a message plan at a given sample and a given history.

With these preliminaries out of the way, we can return to the proof of Theorem 2. Hereafter, and in line with Lemma S.2, Assumption S.1, and Lemma S.4, we fix a p-truthful SPBE $(\mathring{\mu}^{*T}, \mathring{\lambda}^{*T})$ of $G^{\mathring{S}\mathring{M}T}$ in which the players' message plans induced by the behavior strategies $\mathring{\lambda}^{*iT}(\mathring{\gamma}^T)$ in the subgames $\mathring{\gamma}^T \in \mathring{\Gamma}^T$ of $G^{\mathring{S}\mathring{M}T}$ are all of Young class $\mathring{\alpha}$, and we let $\mathring{\gamma}^{*T}_j \equiv ((\mathring{\sigma}^*_j(t))^T_{t=1}, \mathring{\phi}^{*T}_j)$ be the mechanism posted by principal j in that SPBE.

We are now ready to specify a PBE $(\mathring{\mu}^{\#}, \mathring{\lambda}^{\#})$ of $G^{\mathring{S}\mathring{M}}$ corresponding to the SPBE $(\mathring{\mu}^{*T}, \mathring{\lambda}^{*T})$ of $G^{\mathring{S}\mathring{M}T}$. We first describe the principals' and the agents' strategies (Steps 6–7). We then argue that the allocation induced by $(\mathring{\mu}^{\#}, \mathring{\lambda}^{\#})$ in $G^{\mathring{S}\mathring{M}}$ is the same as the one induced by $(\mathring{\mu}^{*T}, \mathring{\lambda}^{*T})$ in $G^{\mathring{S}\mathring{M}T}$ (Step 8). Finally, we show that $(\mathring{\mu}^{\#}, \mathring{\lambda}^{\#})$ satisfies all the equilibrium requirements in $G^{\mathring{S}\mathring{M}}$ (Steps 9–10) and is outcome-equivalent to a p-truthful PBE of $G^{\mathring{S}\mathring{M}}$ (Step 11).

Step 6: Principals' Strategies We start by describing two embeddings that permit us to map information in $G^{\mathring{S}\mathring{M}T}$ into information in $G^{\mathring{S}\mathring{M}}$ and vice-versa.

Embeddings First, let $\mathring{a}^i_j : \mathring{S}^i_j(1) \times \Xi^i_j \to \mathring{S}^i_j$ be a Borel isomorphism mapping the round-1 signal $\mathring{s}^i_j(1) \in \mathring{S}^i_j(1) = [0,1]$ from principal j to agent i in $G^{\mathring{S}\mathring{M}T}$ and the auxiliary sampling variable $\xi^i_j \in \Xi^i_j \equiv [0,1]$ into a signal $\mathring{s}^i_j \in \mathring{S}^i_j = [0,1]$ from principal j to agent i in $G^{\mathring{S}\mathring{M}}$. We let $(\mathring{a}^{i-}_{j(s)}, \mathring{a}^{i-}_{j(\xi)}) \equiv (\mathring{a}^i_j)^{-1} : \mathring{S}^i_j \to \mathring{S}^i_j(1) \times \Xi^i_j$ be the inverse mapping of $\mathring{a}^i_j$.

Second, let $\mathring{b}^i_j : \Omega^i \times \mathring{S}^i_{-j}(1) \times \Xi^i \times \Xi^i_{-j} \times \Theta^{iT}_{\mathring{\alpha}} \to \mathring{M}^i_j$ be a Borel isomorphism mapping agent i's exogenous type $\omega^i \in \Omega^i$, the round-1 signals $\mathring{s}^i_{-j}(1) \in \mathring{S}^i_{-j}(1) = [0,1]^{J-1}$ from the principals other than j to agent i in $G^{\mathring{S}\mathring{M}T}$, the sampling variable $\xi^i \in \Xi^i$, the auxiliary sampling variables $\xi^i_{-j} \in \Xi^i_{-j} \equiv \times_{k\neq j} \Xi^i_k = [0,1]^{J-1}$, and a code $\vartheta^{iT}_{\mathring{\alpha}} \in \Theta^{iT}_{\mathring{\alpha}}$ for agent i's message plan in $G^{\mathring{S}\mathring{M}T}$ into a message $\mathring{m}^i_j \in \mathring{M}^i_j = \Omega^i \times [0,1]^{J-1}$ from agent i to principal j in $G^{\mathring{S}\mathring{M}}$. We let $(\mathring{b}^{i-}_{j(\omega)}, \mathring{b}^{i-}_{j(s_-)}, \mathring{b}^{i-}_{j(\xi)}, \mathring{b}^{i-}_{j(\xi_-)}, \mathring{b}^{i-}_{j(\vartheta)}) \equiv (\mathring{b}^i_j)^{-1} : \mathring{M}^i_j \to \Omega^i \times \mathring{S}^i_{-j}(1) \times \Xi^i \times \Xi^i_{-j} \times \Theta^{iT}_{\mathring{\alpha}}$ be the inverse mapping of $\mathring{b}^i_j$.

Let $\mathring{\mu}^{\#}_j$ be the strategy for principal j in $G^{\mathring{S}\mathring{M}}$ that consists in posting with probability 1 the mechanism $\mathring{\gamma}^{\#}_j = (\mathring{\sigma}^{\#}_j, \mathring{\phi}^{\#}_j)$ defined as follows.

We start with the distribution $\mathring{\sigma}^{\#}_j$. Principal j first draws a vector of signals $\mathring{s}_j(1)$ from the original distribution $\mathring{\sigma}^*_j(1)$, and then independently draws all the $(\xi^i_j)^I_{i=1}$ uniformly from $[0,1]$. Finally, she uses the Borel isomorphism $\mathring{a}^i_j$ described above to map each $(\mathring{s}^i_j(1), \xi^i_j)$ into the corresponding signal $\mathring{s}^i_j = \mathring{a}^i_j(\mathring{s}^i_j(1), \xi^i_j)$ to disclose to every agent i in $G^{\mathring{S}\mathring{M}}$. With a slight abuse of notation, the distribution $\mathring{\sigma}^{\#}_j$ of $\mathring{s}_j$ is thus the push-forward of the measure $\mathring{\sigma}^*_j(1) \otimes \bigotimes^I_{i=1} \mathrm{d}\xi^i_j$ by the mapping $(\mathring{a}^i_j)^I_{i=1} : \mathring{S}_j(1) \times \Xi_j \to \mathring{S}_j$, where $\Xi_j \equiv \times^I_{i=1} \Xi^i_j$; that is, for each $A \in \mathcal{B}([0,1])$,

$$\mathring{\sigma}^{\#}_j(A) = (\mathring{a}^i_j)^I_{i=1} \sharp\, \mathring{\sigma}^*_j(1) \otimes \bigotimes_{i=1}^{I} \mathrm{d}\xi^i_j(A) = \mathring{\sigma}^*_j(1) \otimes \bigotimes_{i=1}^{I} \mathrm{d}\xi^i_j\,(((\mathring{a}^i_j)^I_{i=1})^{-1}(A)). \tag{S.47}$$

Consider next the extended decision rule $\mathring{\phi}^{\#}_j$. Given the signals $\mathring{s}_j = (\mathring{s}^i_j)^I_{i=1}$ disclosed to the agents and the messages $\mathring{m}_j = (\mathring{m}^i_j)^I_{i=1}$ received from them, for every agent i, principal j's mechanism uncovers the information $(\mathring{s}^i_j(1), \xi^i_j) \equiv (\mathring{a}^i_j)^{-1}(\mathring{s}^i_j)$ from the signal $\mathring{s}^i_j$ disclosed to agent i, and the information $(\omega^i, \mathring{s}^i_{-j}(1), \xi^i, \xi^i_{-j}, \vartheta^{iT}_{\mathring{\alpha}}) = (\mathring{b}^i_j)^{-1}(\mathring{m}^i_j)$ from the message $\mathring{m}^i_j$ received from agent i. Using the procedure described in Step 4, principal j's mechanism then constructs the variable $\xi_0 = \{\sum^J_{j=1}\{\sum^I_{i=1} \xi^i_j\}\}$, and finally implements the decision $\varphi_j(\mathring{\gamma}^{*T}, (\mathring{\varpi}^{iT}(\vartheta^{iT}_{\mathring{\alpha}}))^I_{i=1}, (\mathring{s}^i(1))^I_{i=1}, (\omega^i)^I_{i=1}, \xi_0)$, which, as explained in Step 3, is invariant in $\mathring{\phi}^{*T}_{-j}$ and $\mathring{\sigma}^*_{-j}(1)$. Accordingly, we let

$$\begin{aligned}\mathring{\phi}^{\#}_j(\mathring{s}_j, \mathring{m}_j) \equiv \varphi_j\Bigg(&\mathring{\gamma}^{*T}, \left(\mathring{\varpi}^{iT} \circ \mathring{b}^{i-}_{j(\vartheta)}(\mathring{m}^i_j)\right)^I_{i=1}, \left(\mathring{a}^{i-}_{j(s)}(\mathring{s}^i_j), \mathring{b}^{i-}_{j(s_-)}(\mathring{m}^i_j)\right)^I_{i=1}, \left(\mathring{b}^{i-}_{j(\xi)}(\mathring{m}^i_j)\right)^I_{i=1},\\ &\left\{\left\{\sum_{i=1}^I \mathring{a}^{i-}_{j(\xi)}(\mathring{s}^i_j)\right\} + \sum_{k\neq j}\left\{\sum_{i=1}^I \mathring{b}^{i-}_{j,k(\xi_-)}(\mathring{m}^i_j)\right\}\right\}, \left(\mathring{b}^{i-}_{j(\omega)}(\mathring{m}^i_j)\right)^I_{i=1}\Bigg),\\ &(\mathring{s}_j, \mathring{m}_j) \in \mathring{S}_j \times \mathring{M}_j.\end{aligned} \tag{S.48}$$

Our next result shows that (S.48) yields a well-defined extended decision rule.

Lemma S.5 *For each j, $\mathring{\phi}^{\#}_j : \mathring{S}_j \times \mathring{M}_j \to \Delta(X_j)$ is Borel-measurable.*

Proof. Observe first that, by construction of the embeddings $\mathring{a}^i_j$ and $\mathring{b}^i_j$, the last four arguments of the function on the right-hand side of (S.48) are Borel-measurable in $(\mathring{s}_j, \mathring{m}_j)$.

Bearing this in mind, we shall thus simply refer to them as $(\mathring{s}_l(1))_{l=1}^J$, $(\xi^i)_{i=1}^I$, ξ_0, and ω. Now, from (S.43)–(S.45), (S.46), and (S.48), we have

$$\begin{aligned}&\mathring{\phi}_j^{\#}(\mathring{s}_j, \mathring{m}_j)\\ &= \mathring{\chi}_{\mathring{\gamma}_j^T}\Big(\mathring{s}_j(1), e_j(\xi_0), \big(\big(\overline{U}^i_{\mathring{\alpha}}(t)\big(\big(\xi^i, \mathring{h}^i(t), \omega^i\big), (\iota_J^T)^{-1}(j,t) \circ \overline{\theta}^{iT}_{\mathring{\alpha}} \circ \mathring{\varpi}^{iT} \circ \mathring{b}^{i-}_{j(\vartheta)}(\mathring{m}^i_j)\big)\big)_{i=1}^I\big)_{t=1}^T\Big),\\ &\qquad (\mathring{s}_j, \mathring{m}_j) \in \mathring{S}_j \times \mathring{M}_j, \end{aligned} \tag{S.49}$$

where e_j if the j-component of the push-forward function constructed in the proof of Lemma S.3, $(\iota_J^T)^{-1}(j,t)$ is the (j,t)-component of the inverse of the Borel isomorphism ι_J^T, and the histories $\mathring{h}^i(t)$ are recursively defined by (S.44), using $e_j(\xi_0)$ in lieu of ε_j^T and $\mathring{\varpi}^{iT} \circ \mathring{b}^{i-}_{j(\vartheta)}(\mathring{m}^i_j)$ in lieu of $\mathring{\pi}^{iT}$. By construction, we can choose $\overline{\theta}^{iT}_{\mathring{\alpha}}$ and $\mathring{\varpi}^{iT}$ so that $\overline{\theta}^{iT}_{\mathring{\alpha}} \circ \mathring{\varpi}^{iT} = \mathrm{Id}_{\Theta^{iT}_{\mathring{\alpha}}}$, which yields, by (S.49),

$$\begin{aligned}\mathring{\phi}_j^{\#}(\mathring{s}_j, \mathring{m}_j) &= \mathring{\chi}_{\mathring{\gamma}_j^T}\Big(\mathring{s}_j(1), e_j(\xi_0), \big(\big(\overline{U}^i_{\mathring{\alpha}}(t)\big(\big(\xi^i, \mathring{h}^i(t), \omega^i\big), (\iota_J^T)^{-1}(j,t) \circ \mathring{b}^{i-}_{j(\vartheta)}(\mathring{m}^i_j)\big)\big)_{i=1}^I\big)_{t=1}^T\Big),\\ &\qquad (\mathring{s}_j, \mathring{m}_j) \in \mathring{S}_j \times \mathring{M}_j. \end{aligned} \tag{S.50}$$

Because the functions $\mathring{\chi}_{\mathring{\gamma}_j^T}$, e_j, $\overline{U}^i_{\mathring{\alpha}}(t)$, $(\iota_J^T)^{-1}(j,t)$, and $\mathring{b}^{i-}_{j(\vartheta)}$ are all Borel-measurable, it follows from (S.50) and our initial observation that there only remains to check that, for each i, the history $\mathring{h}^i(t)$ is a Borel-measurable function of $((\mathring{s}_l(1))_{l=1}^J, (\xi^i)_{i=1}^I, \xi_0, \omega, \mathring{m}_j)$, and thus of $(\mathring{s}_j, \mathring{m}_j)$. We proceed by induction, using (S.44). The claim is obvious for $t = 1$. Suppose then the claim established for some $t \geq 1$. By (S.44) and (S.46), we have

$$\begin{aligned}\mathring{h}^i(t+1) \equiv \Bigg(&\mathring{h}^i(t), \Big(\eta^i_{\mathring{\sigma}_l^T}(t+1)\big(\mathring{s}_l(1), e_l(\xi_0),\\ &\big(\big(\overline{U}^k_{\mathring{\alpha}}(\tau)\big(\big(\xi^k, \mathring{h}^k(\tau), \omega^k\big), (\iota_J^T)^{-1}(l,\tau) \circ \mathring{b}^{k-}_{j(\vartheta)}(\mathring{m}^k_j)\big)\big)\big)_{k=1}^I\big)_{\tau=1}^t\Big)_{l=1}^J\Bigg),\end{aligned}$$

which, by the induction hypothesis, establishes the claim for $t+1$ as the functions $\eta^i_{\mathring{\sigma}_l^T}(t+1)$, e_l, $\overline{U}^k_{\mathring{\alpha}}(\tau)$, $(\iota_J^T)^{-1}(l,\tau)$, and $\mathring{b}^{k-}_{j(\vartheta)}$ are all Borel-measurable. The result follows. ■

Step 7: Agents' Strategies To construct every agent i's strategy $\mathring{\lambda}^{\# i}$ in $G^{\mathring{S}\mathring{M}}$, we distinguish three cases according to the profile of mechanisms $\mathring{\gamma} \equiv (\mathring{\gamma}_j)_{j=1}^J$ posted by the principals.

Case 1 If $\mathring{\gamma} = \mathring{\gamma}^{\#}$, that is, every principal j posts her candidate equilibrium mechanism $\mathring{\gamma}_j^{\#}$, then every agent i sends to every principal j the message

$$\mathring{m}^i_j = \mathring{b}^i_j(\omega^i, \mathring{s}^i_{-j}(1), \xi^i, \xi^i_{-j}, \vartheta^{iT}_{\mathring{\alpha}}), \tag{S.51}$$

where ω^i is agent i's true type, $\mathring{s}^i_{-j}(1) = (\mathring{a}^{i-}_{k(s)}(\mathring{s}^i_k))_{k \neq j}$ are round-1 signals in $G^{\mathring{S}\mathring{M}T}$, encoded into the signals $(\mathring{s}^i_k)_{k \neq j}$ agent i received from the principals other than j, ξ^i is agent i's sampling variable, $\xi^i_{-j} = (\mathring{a}^{i-}_{k(\xi)}(\mathring{s}^i_k))_{k\neq j}$ are the new auxiliary variables from the principals other than j, also encoded into the signals $(\mathring{s}^i_k)_{k\neq j}$, and $\vartheta^{iT}_{\mathring{\alpha}}$ is the code for agent i's message plan in $G^{\mathring{S}\mathring{M}T}$ that prescribes truthful reporting in all rounds no matter the realization of ξ^i; this message plan is continuous, hence certainly of Young class $\mathring{\alpha}$.

Case 2 If $\mathring{\gamma}_j = (\mathring{\sigma}_j, \mathring{\phi}_j) \neq \mathring{\gamma}^{\#}_j$ but $\mathring{\gamma}_{-j} = \mathring{\gamma}^{\#}_{-j}$, that is, principal j unilaterally deviates from $\mathring{\gamma}^{\#}$, then every agent i's behavior in $G^{\mathring{S}\mathring{M}}$ is predicated on the behavior that agent i would have followed in the subgame of $G^{\mathring{S}\mathring{M}T}$ in which every principal $k \neq j$ posts the mechanism $\mathring{\gamma}^{*T}_k$ and principal j posts the mechanism $L^T_j(\mathring{\gamma}_j)$ such that: (i) the distribution $\mathring{\sigma}_j(1)$ over round-1 signals is the same as the distribution $\mathring{\sigma}_j$ over signals in $\mathring{\gamma}_j$; (ii) for each $t < \infty$, $1 < t \leq T$, the transition probability $\mathring{\sigma}_j(t)$ is the same as the transition probability $\mathring{\sigma}^*_j(t)$ in $\mathring{\gamma}^{*T}_j$; (iii) the extended decision rule $\mathring{\phi}^T_j$ is invariant in signals sent and messages received past round 1, and implements the same decisions as $\mathring{\phi}_j$ when the round-1 signals and messages are the same under the two mechanisms, that is,

$$\mathring{\phi}^T_j\big((\mathring{s}_j(t), \mathring{m}_j(t))^T_{t=1}\big) \equiv \mathring{\phi}_j(\mathring{s}_j(1), \mathring{m}_j(1)), \quad (\mathring{s}_j(t), \mathring{m}_j(t))^T_{t=1} \in \mathring{H}^T_j,$$

where $\mathring{H}^T_j$ is defined analogously to $\hat{H}^T_j$. We accordingly postulate that every agent i of type ω^i draws ξ^i uniformly from $[0, 1]$ and then sends to principal j the message

$$\mathring{m}^i_j = \mathring{\lambda}^{*i,\xi^i}_j(1)\big(\big(L^T_j(\mathring{\gamma}_j), \mathring{\gamma}^{*T}_{-j}\big), \big(\mathring{s}^i_j, (\mathring{a}^{i-}_{k(s)}(\mathring{s}^i_k))_{k\neq j}\big), \omega^i\big) \tag{S.52}$$

he would send to her in round 1 of the subgame $(L^T_j(\mathring{\gamma}_j), \mathring{\gamma}^{*T}_{-j})$ of $G^{\mathring{S}\mathring{M}T}$ upon receiving the signals $(\mathring{s}^i_j, (\mathring{a}^{i-}_{k(s)}(\mathring{s}^i_k))_{k\neq j})$. In addition, agent i sends to every principal $k \neq j$ the message

$$\mathring{m}^i_k = \mathring{b}^i_k(\omega^i, \mathring{s}^i_{-k}(1), \xi^i, \xi^i_{-k}, \vartheta^{iT}_{\mathring{\alpha}}), \tag{S.53}$$

where ω^i is agent i's true type, $\mathring{s}^i_{-k}(1) = (\mathring{s}^i_j, (\mathring{a}^{i-}_{l(s)}(\mathring{s}^i_l))_{l\neq j,k})$ gathers the true signal $\mathring{s}^i_j$ agent i received from principal j and the round-1 signals in $G^{\mathring{S}\mathring{M}T}$ encoded into the signals $(\mathring{s}^i_l)_{l\neq j,k}$ agent i received from the principals other than j and k, ξ^i is agent i's sampling variable, $\xi^i_{-k} = (\xi^i_j, (\mathring{a}^{i-}_{l(\xi)}(\mathring{s}^i_l))_{l\neq j,k})$ gathers an arbitrary ξ^i_j for principal j that agent i draws himself uniformly from $[0, 1]$ and the new auxiliary variables from the principals other than j and k, also encoded into the signals $(\mathring{s}^i_l)_{l\neq j,k}$, and $\vartheta^{iT}_{\mathring{\alpha}}$ is the code for agent i's message plan in $G^{\mathring{S}\mathring{M}T}$ associated to agent i's behavior strategy $\mathring{\lambda}^{*iT}(L^T_j(\mathring{\gamma}_j), \mathring{\gamma}^{*T}_{-j})$ in the subgame $(L^T_j(\mathring{\gamma}_j), \mathring{\gamma}^{*T}_{-j})$; this message plan, by Assumption S.1 and Lemma S.4, is of Young class $\mathring{\alpha}$.

Case 3 Finally, if more than one principal deviate from $\mathring{\gamma}^{\#}$, then every agent i's behavior in $G^{\mathring{S}\mathring{M}}$ is predicated on the behavior that agent i would have followed in round 1 of the subgame of $G^{\hat{S}\hat{M}T}$ in which every principal j posts the mechanism $L^T_j(\mathring{\gamma}_j)$. That is, we postulate that every agent i of type ω^i draws ξ^i uniformly from $[0,1]$ and then sends the message

$$\mathring{m}_j^i = \mathring{\lambda}_j^{*i,\xi^i}(1)\Big(\big(L_j^T(\mathring{\gamma}_j)\big)_{j=1}^J, \mathring{s}^i, \omega^i\Big) \tag{S.54}$$

to every principal j.

This completes the description of every agent i's candidate equilibrium strategy $\mathring{\lambda}^{\#i}$ in $G^{\mathring{S}\mathring{M}}$. Notice that, because the functions $(a_j^i)^{-1}$ and b_j^i are Borel-measurable for all i and j, the measurability restrictions imposed on the functions $\mathring{\lambda}^{*i}$ imply that $\mathring{\lambda}^{\#i}$ is $(\mathcal{B}([0,1]) \otimes \hat{\boldsymbol{\Sigma}} \otimes \mathring{S}^i \otimes 2^{\Omega^i}, \mathring{M}^i)$-measurable, as requested.

We let $\mathring{\lambda}^{\#} \equiv (\mathring{\lambda}^{\#i})_{i=1}^I$ be the profile of agents' candidate equilibrium strategies in $G^{\mathring{S}\mathring{M}}$.

Step 8: Outcome Equivalence of $(\mathring{\mu}^{\#}, \mathring{\lambda}^{\#})$ and $(\mathring{\mu}^{*T}, \mathring{\lambda}^{*T})$ We now claim that the strategy profiles $(\mathring{\mu}^{\#}, \mathring{\lambda}^{\#})$ and $(\mathring{\mu}^{*T}, \mathring{\lambda}^{*T})$ are outcome-equivalent. Indeed, the allocation $z_{\mathring{\mu}^{\#},\mathring{\lambda}^{\#}}$ induced by $(\mathring{\mu}^{\#}, \mathring{\lambda}^{\#})$ in $G^{\mathring{S}\mathring{M}}$ satisfies

$$\begin{aligned}
& z_{\mathring{\mu}^{\#},\mathring{\lambda}^{\#}}(x\,|\,\omega) \\
&= \int_{\times_{i=1}^I \Xi^i} \int_{\times_{j=1}^J \mathring{S}_j} \prod_{j=1}^J \mathring{\phi}_j^{\#} \left(\mathring{s}_j, \big(\underbrace{\mathring{b}_j^i\big(\omega^i, (\mathring{a}_{k(s)}^{i-}(\mathring{s}_k^i))_{k\neq j}, \xi^i, (\mathring{a}_{k(\xi)}^{i-}(\mathring{s}_k^i))_{k\neq j}, \overline{\theta}_{\mathring{\alpha}}^{iT}(\mathring{\pi}_{tr}^{iT})\big)}_{\mathring{m}_j^i} \big)_{i=1}^I \right)(x_j) \\
&\qquad \bigotimes_{j=1}^J \left[(\mathring{a}_j^i)_{i=1}^I \sharp\, \mathring{\sigma}_j^*(1) \otimes \bigotimes_{i=1}^I \mathrm{d}\xi_j^i \right](\mathrm{d}\mathring{s}_j) \bigotimes_{i=1}^I \mathrm{d}\xi^i \\
&= \int_{\times_{i=1}^I \Xi^i} \int_{\times_{j=1}^J \mathring{S}_j(1)} \int_{\Xi_0} \prod_{j=1}^J \varphi_j\big(\mathring{\gamma}^{*T}, \mathring{\pi}_{tr}^T, (\mathring{s}_j(1))_{j=1}^J, (\xi^i)_{i=1}^I, \xi_0, \omega\big)\, \mathrm{d}\xi_0 \bigotimes_{j=1}^J \mathring{\sigma}_j^*(1)(\mathrm{d}\mathring{s}_j(1)) \bigotimes_{i=1}^I \mathrm{d}\xi^i \\
&= \mathring{z}_{\mathring{\gamma}^{*T},\mathring{\pi}_{tr}^T}(x\,|\,\omega) \\
&= z_{\mathring{\mu}^{*T},\mathring{\lambda}^{*T}}(x\,|\,\omega)
\end{aligned} \tag{S.55}$$

for all $(x,\omega) \in X \times \Omega$, where the first equality follows from (S.47) and (S.51), denoting by $\mathring{\pi}_{tr}^{iT}$ agent i's message plan in $G^{\mathring{S}\mathring{M}T}$ that prescribes truthful reporting in all rounds no matter the realization of ξ^i, the second equality follows from (S.48), the change-of-variable formula for push-forward measures, the equalities

$$\begin{aligned}
\mathring{\varpi}^{iT} \circ \mathring{b}_{j(\vartheta)}^{i-}(\mathring{m}_j^i) &= \mathring{\pi}_{tr}^{iT}, \\
\mathring{a}_{j(s)}^{i-}(\mathring{s}_j^i) &= \mathring{s}_j^i(1), \\
\mathring{b}_{j(s_-)}^{i-}(\mathring{m}_j^i) &= \mathring{s}_{-j}^i(1), \\
\mathring{b}_{j(\xi)}^{i-}(\mathring{m}_j^i) &= \xi^i, \\
\mathring{a}_{j(\xi)}^{i-}(\mathring{s}_j^i) &= \xi_j^i, \\
\mathring{b}_{j(\xi_-)}^{i-}(\mathring{m}_j^i) &= \xi_{-j}^i,
\end{aligned}$$

$$\mathring{b}^{i-}_{j(\omega)}(\mathring{m}^i_j) = \omega^i,$$

the definition of ξ_0, and the fact that ξ_0 is uniformly distributed over $[0,1]$, the third equality follows from (S.45), and the fourth equality follows from the fact that $(\mathring{\mu}^{*T}, \mathring{\lambda}^{*T})$ is a p-truthful SPBE of $G^{\mathring{S}\mathring{M}T}$ in which the principals post the mechanisms $\mathring{\gamma}^{*T}$. Hence $z_{\mathring{\mu}^{\#},\mathring{\lambda}^{\#}} = z_{\mathring{\mu}^{*T},\mathring{\lambda}^{*T}}$, as claimed.

Step 9: Equilibrium Properties of $\mathring{\lambda}^{\#}$ We distinguish three cases. In each case, we study the incentives of some agent i, assuming that the other agents stick to their candidate equilibrium strategies $\mathring{\lambda}^{\#-i}$.

Case 1 Suppose first that $\mathring{\gamma} = \mathring{\gamma}^{\#}$. If agent i does not deviate from $\mathring{\lambda}^{\#i}$, then the allocation implemented in $G^{\mathring{S}\mathring{M}}$ is given by (S.55). Now, observe first that, given the way the random variable ξ_0 is constructed in Step 4, there is no way for agent i to infer the value of ξ_0 from the auxiliary sampling variables $(\xi^i_j)^J_{j=1} = (\mathring{a}^{i-}_{j(\xi)}(\mathring{s}^i_j))^J_{j=1}$ encoded into the signals $(\mathring{s}^i_j)^J_{j=1}$ he receives, nor, a fortiori, to manipulate the distribution of ξ_0 through his messages to the principals. Second, recall that b^i_j is a Borel isomorphism between $\Omega^i \times \mathring{S}^i_{-j}(1) \times \Xi^i \times \Xi^i_{-j} \times \Theta^{iT}_{\mathring{\alpha}}$ and $\mathring{M}^i_j$. Thus, if agent i, in $G^{\mathring{S}\mathring{M}}$, unilaterally deviates from (S.51) for at least some j, then every mechanism $\mathring{\gamma}^{\#}_j$, given the signals disclosed to the agents and the messages received from them, implements the same decision

$$\varphi_j(\mathring{\gamma}^{*T}, (\hat{\mathring{\pi}}^{iT}, \mathring{\pi}^{-iT}_{tr}), (\mathring{s}^i_j(1), \hat{\mathring{s}}^i_{-j}(1), \mathring{s}^{-i}(1)), (\xi^i, \xi^{-i}), \xi_0, (\hat{\omega}^i, \omega^{-i}))$$

that would be implemented in the subgame $\mathring{\gamma}^{*T}$ of $G^{\mathring{S}\mathring{M}T}$ if agent i unilaterally deviated to some message plan $\hat{\mathring{\pi}}^{iT}$ different from $\mathring{\pi}^{iT}_{tr}$, or reported round-1 signals $\hat{\mathring{s}}^i_{-j}(1)$ received from all principals other than j different from their true signals $\mathring{s}^i_{-j}(1)$, or reported a type $\hat{\omega}^i$ different from his true type ω^i. Because $(\mathring{\mu}^{*T}, \mathring{\lambda}^{*T})$ is a p-truthful SPBE of $G^{\mathring{S}\mathring{M}T}$, we conclude that, when $\mathring{\gamma} = \mathring{\gamma}^{\#}$ and the other agents follow their candidate equilibrium strategies $\mathring{\lambda}^{\#-i}(\mathring{\gamma}^{\#})$ in the subgame $\mathring{\gamma}^{\#}$ of $G^{\mathring{S}\mathring{M}}$, agent i cannot profitably deviate from his candidate equilibrium strategy $\mathring{\lambda}^{\#i}(\mathring{\gamma}^{\#})$.

Case 2 Suppose next that $\mathring{\gamma}_j \equiv (\mathring{\sigma}_j, \mathring{\phi}_j) \neq \mathring{\gamma}^{\#}_j$ but $\mathring{\gamma}_{-j} = \mathring{\gamma}^{\#}_{-j}$ for some j. We first claim that the strategy profiles $((\mathring{\gamma}_j, \mathring{\mu}^{\#}_{-j}), \mathring{\lambda}^{\#})$ in $G^{\mathring{S}\mathring{M}}$ and $((L^T_j(\mathring{\gamma}_j), \mathring{\mu}^{*T}_{-j}), \mathring{\lambda}^{*})$ in $G^{\mathring{S}\mathring{M}T}$ are outcome-equivalent. Indeed, the allocation $z_{(\mathring{\gamma}_j,\mathring{\mu}^{\#}_{-j}),\mathring{\lambda}^{\#}}$ induced by $((\mathring{\gamma}_j, \mathring{\mu}^{\#}_{-j}), \mathring{\lambda}^{\#})$ in $G^{\mathring{S}\mathring{M}}$ satisfies

$$z_{(\mathring{\gamma}_j,\mathring{\mu}^{\#}_{-j}),\mathring{\lambda}^{\#}}(x\,|\,\omega)$$

$$= \int_{\times_{i=1}^I \Xi^i} \int_{\times_{l=1}^J \mathring{S}_l} \mathring{\phi}_j \left(\mathring{s}_j, \Big(\underbrace{\mathring{\lambda}^{*i,\xi^i}_j(1)\big(\big(L^T_j(\mathring{\gamma}_j), \mathring{\gamma}^{*T}_{-j}\big), \big(\mathring{s}^i_j, (\mathring{a}^{i-}_{k(s)}(\mathring{s}^i_k))_{k\neq j}\big), \omega^i\big)}_{\mathring{m}^i_j} \Big)^I_{i=1} \right)(x_j)$$

$$\prod_{k\neq j} \mathring{\phi}_k^{\#}\left(\mathring{s}_k, \big(\underbrace{\mathring{b}_k^i\big(\omega^i, (\mathring{s}_j^i, (\mathring{a}_{l(s)}^{i-}(\mathring{s}_l^i))_{l\neq j,k}), \xi^i, (\xi_j^i, (\mathring{a}_{l(\xi)}^{i-}(\mathring{s}_l^i))_{l\neq j,k}), \overline{\theta}_{\mathring{\alpha}}^{iT}(\mathring{\pi}_{dev\, j}^{*iT})\big)}_{\mathring{m}_k^i}\big)_{i=1}^I \right)(x_k)$$

$$\bigotimes_{k\neq j}\left[(\mathring{a}_k^i)_{i=1}^I \sharp\, \mathring{\sigma}_k^*(1) \otimes \bigotimes_{i=1}^I \mathrm{d}\xi_k^i \right](\mathrm{d}\mathring{s}_k) \otimes \mathring{\sigma}_j(\mathrm{d}\mathring{s}_j) \bigotimes_{i=1}^I \mathrm{d}\xi^i$$

$$= \int_{\times_{i=1}^I \Xi^i} \int_{\times_{k=1}^J \mathring{S}_k(1)} \int_{\Xi_0} \prod_{k=1}^J \varphi_k\big((L_j^T(\mathring{\gamma}_j), \mathring{\gamma}_{-j}^{*T}), \mathring{\pi}_{dev\, j}^{*T}, (\mathring{s}_k(1))_{k=1}^J, (\xi^i)_{i=1}^I, \xi_0, \omega \big)$$

$$\mathrm{d}\xi_0 \bigotimes_{k\neq j} \mathring{\sigma}_k^*(1)(\mathrm{d}\mathring{s}_k(1)) \otimes \mathring{\sigma}_j(1)(\mathrm{d}\mathring{s}_j(1)) \bigotimes_{i=1}^I \mathrm{d}\xi^i$$

$$= \mathring{z}_{(L_j^T(\mathring{\gamma}_j), \mathring{\gamma}_{-j}^{*T}), \mathring{\pi}_{dev\, j}^{*T}}(x \,|\, \omega)$$

$$= z_{(L_j^T(\mathring{\gamma}_j), \mathring{\mu}_{-j}^{*T}), \mathring{\lambda}^{*T}}(x \,|\, \omega) \qquad \text{(S.56)}$$

for all $(x, \omega) \in X \times \Omega$, where the first equality follows from (S.47) and (S.52)–(S.53), denoting by $\mathring{\pi}_{dev\, j}^{*iT}$ agent i's message plan in $G^{\mathring{S}\mathring{M}T}$ associated to agent i's behavior strategy $\mathring{\lambda}^{*iT}(L_j^T(\mathring{\gamma}_j), \mathring{\gamma}_{-j}^{*T})$ in the subgame $(L_j^T(\mathring{\gamma}_j), \mathring{\gamma}_{-j}^{*T})$, the second equality follows from (S.48), the change-of-variable formula for push-forward measures, the equalities

$$\begin{aligned}
\mathring{\varpi}^{iT} \circ \mathring{b}_{k(\vartheta)}^{i-}(\mathring{m}_k^i) &= \mathring{\pi}_{dev\, j}^{*iT}, \\
\mathring{a}_{k(s)}^{i-}(\mathring{s}_k^i) &= \mathring{s}_k^i(1), \\
\mathring{b}_{k(s_-)}^{i-}(\mathring{m}_k^i) &= \mathring{s}_{-k}^i(1), \\
\mathring{b}_{k(\xi)}^{i-}(\mathring{m}_k^i) &= \xi^i, \\
\mathring{a}_{k(\xi)}^{i-}(\mathring{s}_k^i) &= \xi_k^i, \\
\mathring{b}_{k(\xi_-)}^{i-}(\mathring{m}_k^i) &= \xi_{-k}^i, \\
\mathring{b}_{k(\omega)}^{i-}(\mathring{m}_k^i) &= \omega^i,
\end{aligned}$$

the definition of ξ_0, the fact that ξ_0 is uniformly distributed over $[0,1]$, and the fact that, by Step 3, the functions φ_k depend neither on $\mathring{\sigma}_{-k}^*(1)$ nor on $\mathring{\phi}_{-k}^{*T}$ and that the mechanism $L_j^T(\mathring{\gamma}_j)$ involves, for each $t < \infty$, $1 < t \leq T$, the same transition probabilities $\mathring{\sigma}_j^*(t)$ as in $\mathring{\gamma}_j^{*T}$, the third equality follows from (S.45), and the fourth equality follows from the definition of $\mathring{\pi}_{dev\, j}^{*iT}$. Hence $z_{(\mathring{\gamma}_j, \mathring{\mu}_{-j}^{\#}), \mathring{\lambda}^{\#}} = z_{(L_j^T(\mathring{\gamma}_j), \mathring{\mu}_{-j}^{*T}), \mathring{\lambda}^{*T}}$, as claimed.

If agent i does not deviate from $\mathring{\lambda}^{\#i}$ following principal j's unilateral deviation to $\mathring{\gamma}_j$, then the allocation implemented in $G^{\mathring{S}\mathring{M}}$ is given by (S.56). The proof that agent i cannot be better off deviating from (S.52) vis-à-vis the deviating principal j, or from (S.53) vis-à-vis one or several of the non-deviating principals $k \neq j$ then proceeds as in Case 1. Again, agent i can neither infer ξ_0 nor influence its distribution, notably through his draw of the sampling variable ξ_j^i. Moreover, if agent i, in $G^{\mathring{S}\mathring{M}}$, unilaterally deviates from (S.52) vis-à-vis principal

j, or from (S.53) for at least one $k \neq j$, then the mechanisms $\mathring{\gamma}_j$ and $\mathring{\gamma}_k^{\#}$, given the signals disclosed to the agents and the messages received from them, implement the same decisions

$$\varphi_j\big(\big(L_j^T(\mathring{\gamma}_j), \mathring{\gamma}_{-j}^{*T}\big), (\hat{\mathring{\pi}}^{iT}, \mathring{\pi}_{dev\,j}^{-iT}), (\mathring{s}_j^i(1), \hat{\mathring{s}}_{-j}^i(1), \mathring{s}^{-i}(1)), (\xi^i, \xi^{-i}), \xi_0, (\hat{\omega}^i, \omega^{-i})\big),$$
$$\varphi_k\big(\big(L_j^T(\mathring{\gamma}_j), \mathring{\gamma}_{-j}^{*T}\big), (\hat{\mathring{\pi}}^{iT}, \mathring{\pi}_{dev\,j}^{-iT}), (\mathring{s}_k^i(1), \hat{\mathring{s}}_{-k}^i(1), \mathring{s}^{-i}(1)), (\xi^i, \xi^{-i}), \xi_0, (\hat{\omega}^i, \omega^{-i})\big)$$

that would be implemented in the subgame $(L_j^T(\mathring{\gamma}_j), \mathring{\gamma}_{-j}^{\#T})$ of $G^{\mathring{S}\mathring{M}T}$ if agent i unilaterally deviated to some message plan $\hat{\mathring{\pi}}^{iT}$ different from $\mathring{\pi}_{dev\,j}^{*iT}$, or reported round-1 signals $\hat{\mathring{s}}_{-j}^i(1)$ or $\hat{\mathring{s}}_{-k}^i(1)$ received from all principals other than j or k different from the signals he was supposed to report to principals j and k according to $\mathring{\lambda}^{*iT}(L_j^T(\mathring{\gamma}_j), \mathring{\gamma}_{-j}^{\#T})$, or reported a type $\hat{\omega}^i$ different from the type he was supposed to report to principals j and k, again according to $\mathring{\lambda}^{*iT}(L_j^T(\mathring{\gamma}_j), \mathring{\gamma}_{-j}^{\#T})$. Because $\mathring{\lambda}^{*T}(L_j^T(\mathring{\gamma}_j), \mathring{\gamma}_{-j}^{\#T})$ is a BNE of the subgame $(L_j^T(\mathring{\gamma}_j), \mathring{\gamma}_{-j}^{\#T})$, we conclude that, when $\mathring{\gamma}_j \neq \mathring{\gamma}_j^{\#}$ but $\mathring{\gamma}_{-j} = \mathring{\gamma}_{-j}^{\#}$ and the other agents follow their candidate equilibrium strategies $\mathring{\lambda}^{\#-i}(\mathring{\gamma}_j, \mathring{\gamma}_j^{\#})$ in the subgame $(\mathring{\gamma}_j, \mathring{\gamma}_j^{\#})$ of $G^{\mathring{S}\mathring{M}}$, agent i cannot profitably deviate from his candidate equilibrium strategy $\mathring{\lambda}^{\#i}(\mathring{\gamma}_j, \mathring{\gamma}_j^{\#})$.

Case 3 Suppose finally that $\mathring{\gamma}_j \equiv (\mathring{\sigma}_j, \mathring{\phi}_j) \neq \mathring{\gamma}_j^{\#}$ for at least two principals j. Then, the subgames $\mathring{\gamma}$ and $(L_j^T(\mathring{\gamma}_j))_{j=1}^J$ of $G^{\mathring{S}\mathring{M}}$ and $G^{\mathring{S}\mathring{M}T}$ are strategically equivalent. It follows that letting the agents send, in $\mathring{\gamma}$, messages according to (S.54) forms a BNE of $\mathring{\gamma}$.

Step 10: Equilibrium Properties of $\mathring{\mu}^{\#}$ We now check that, given the agents' strategy profile $\mathring{\lambda}^{\#}$, the strategy profile $\mathring{\mu}^{\#}$ is a NE in the principals' game. By Step 8, the allocation induced by $\mathring{\mu}^{\#}$ and $\mathring{\lambda}^{\#}$ in $G^{\mathring{S}\mathring{M}}$ coincides with the allocation induced by $\mathring{\mu}^{*T}$ and $\mathring{\lambda}^{*T}$ in $G^{\mathring{S}\mathring{M}T}$. Moreover, by Case 2 of Step 9, if some principal j unilaterally deviates from $\mathring{\mu}^{\#}$ by posting a mechanism $\mathring{\gamma}_j$, the allocation induced by $(\mathring{\gamma}_j, \mathring{\mu}_{-j}^{\#})$ and $\mathring{\lambda}^{\#}$ in $G^{\mathring{S}\mathring{M}}$ coincides with the allocation induced by $(L_j^T(\mathring{\gamma}_j), \mathring{\mu}_{-j}^{*T})$ and $\mathring{\lambda}^{*T}$ in $G^{\mathring{S}\mathring{M}T}$. Because $(\mathring{\mu}^{*T}, \mathring{\lambda}^{*T})$ is an SPBE of $G^{\mathring{S}\mathring{M}T}$, it follows that no principal j can profitably deviate from $\mathring{\mu}_j^{\#}$ in $G^{\mathring{S}\mathring{M}}$ given the other principals' strategy profile $\mathring{\mu}_{-j}^{\#}$ and the agents' strategy profile $\mathring{\lambda}^{\#}$. Thus $(\mathring{\mu}^{\#}, \mathring{\lambda}^{\#})$ is a PBE of $G^{\mathring{S}\mathring{M}}$ that is outcome-equivalent to $(\mathring{\mu}^{*T}, \mathring{\lambda}^{*T})$.

Step 11: Towards a P-Truthful PBE of $G^{\mathring{S}\mathring{M}}$ The PBE $(\mathring{\mu}^{\#}, \mathring{\lambda}^{\#})$ of $G^{\mathring{S}\mathring{M}}$ is not necessarily p-truthful as the agents need not report truthfully on path, that is, in the subgame $\mathring{\gamma}^{\#}$. In fact, every agent i, under $\mathring{\lambda}^{i\#}$, make reports according to (S.51) and, whereas the message plan $\mathring{\pi}_{tr}^{*iT}$ prescribes truthful reporting in all rounds no matter the realization of ξ^i, there is no reason why $\mathring{m}_j^i$ as given by (S.51) should be equal to $q_j^i \equiv (\omega^i, \mathring{s}_{-j}^i)$. Applying Theorem 1, however, enables us to construct a p-truthful PBE $(\mathring{\mu}^*, \mathring{\lambda}^*)$ of $G^{\mathring{S}\mathring{M}}$ that is outcome-equivalent to $(\mathring{\mu}^{\#}, \mathring{\lambda}^{\#})$.

This concludes the proof of Theorem 2(i), establishing universality.

Part (ii) (Robustness) By Theorem 1, it is sufficient to establish the result for a p-truthful PBE $(\mathring{\gamma}^*, \mathring{\lambda}^*)$ of $G^{\mathring{S}\mathring{M}}$ in which every principal j posts a mechanism $\mathring{\gamma}_j^* \equiv (\mathring{\sigma}_j^*, \mathring{\phi}_j^*)$. As $G^{\hat{S}\hat{M}T}$ admits an SPBE by assumption, Lemma S.2 ensures that $G^{\mathring{S}\mathring{M}T}$ admits a p-truthful SPBE. The bulk of the argument below consists in constructing a p-truthful SPBE of $G^{\mathring{S}\mathring{M}T}$ that is outcome-equivalent to $(\mathring{\gamma}^*, \mathring{\lambda}^*)$. The proof consists of six steps.

Step 1: Principals' Strategies Every principal j posts with probability 1 a mechanism $\mathring{\gamma}_j^{*T} \equiv (\mathring{\sigma}_j^{*T}, \mathring{\phi}_j^{*T})$ defined as follows. First, the distribution $\mathring{\sigma}_j^*(1)$ over round-1 signals in $\mathring{\gamma}_j^{*T}$ is the same as the distribution $\mathring{\sigma}_j^*$ over signals in $\mathring{\gamma}_j^*$ and, for each $t < \infty$, $1 < t \leq T$, every transition probability $\mathring{\sigma}_j^*(t)$ is degenerate, sending the same signal to every agent i no matter principal j's private history of signals and messages. Second, $\mathring{\phi}_j^{*T}$ is invariant in signals sent and messages received past round 1, and implements the same decisions as $\mathring{\phi}_j^*$ when the round-1 signals and messages are the same under the two mechanisms, that is,

$$\mathring{\phi}_j^{*T}\big((\mathring{s}_j(t), \mathring{m}_j(t))_{t=1}^T\big) \equiv \mathring{\phi}_j^*(\mathring{s}_j(1), \mathring{m}_j(1)), \quad (\mathring{s}_j(t), \mathring{m}_j(t))_{t=1}^T \in \mathring{H}_j^T.$$

Step 2: Agents' Strategies To construct every agent i's strategy $\mathring{\lambda}^{*iT}$ in $G^{\mathring{S}\mathring{M}T}$, we distinguish three cases according to the profile of mechanisms $\mathring{\gamma}^T \equiv (\mathring{\gamma}_j^T)_{j=1}^J$ posted by the principals.

Case 1 In any subgame $\mathring{\gamma}^T$ of $G^{\mathring{S}\mathring{M}T}$ in which the extended decision rule $\mathring{\phi}_j^T$ in the mechanism $\mathring{\gamma}_j^T$ posted by every principal j is invariant in signals sent and messages received past round 1, every agent i's behavior in $G^{\mathring{S}\mathring{M}T}$ is predicated on the behavior that agent i would have followed in the subgame of $G^{\mathring{S}\mathring{M}}$ in which every principal j posts the mechanism $C_j(\mathring{\gamma}_j^T)$ such that: (i) the distribution $\mathring{\sigma}_j$ over signals is the same as the distribution $\mathring{\sigma}_j(1)$ over round-1 signals in $\mathring{\gamma}_j^T$; (ii) the extended decision rule $\mathring{\phi}_j$ implements the same decisions as $\mathring{\phi}_j^T$ when the round-1 signals and messages are the same under the two mechanisms, that is, we have

$$\mathring{\phi}_j(\mathring{s}_j, \mathring{m}_j) = \mathring{\phi}_j^T\big(\big(\mathring{s}_j, (\mathring{s}_j(t))_{t=2}^T\big), \big(\mathring{m}_j, (\mathring{m}_j(t))_{t=2}^T\big)\big), \quad \big(\big(\mathring{s}_j, (\mathring{s}_j(t))_{t=2}^T\big), \big(\mathring{m}_j, (\mathring{m}_j(t))_{t=2}^T\big)\big) \in \mathring{H}_j^T.$$

We accordingly postulate that, in round 1 of the subgame $\mathring{\gamma}^T$, every agent i of type ω^i draws ξ^i uniformly from $[0, 1]$ and then sends to principal j the message

$$\mathring{m}_j^i = \mathring{\lambda}_j^{*i,\xi^i}\big((C_j(\mathring{\gamma}_j^T))_{j=1}^J, (\mathring{s}_j^i)_{j=1}^J, \omega^i\big)$$

he would send to her in the subgame $(C_j(\mathring{\gamma}_j^T))_{j=1}^J$ of $G^{\mathring{S}\mathring{M}}$ upon receiving the signals $(\mathring{s}_j^i)_{j=1}^J$, and then sends arbitrary messages past round 1. In particular, if every principal j posts her candidate equilibrium mechanism $\mathring{\gamma}_j^{*T}$, then every agent i truthfully reports $q_j^i(1) \equiv (\omega^i, \mathring{s}_{-j}^i(1))$ to every principal j in round 1.

Case 2 In any subgame $\mathring{\gamma}^T$ of $G^{\mathring{S}\mathring{M}T}$ in which there exists some j such that the extended decision rule $\mathring{\phi}_j^T$ in the mechanism $\mathring{\gamma}_j^T$ posted by principal j is not invariant in signals sent or messages received past round 1, whereas the extended decision rule $\mathring{\phi}_k^T$ in the mechanism $\mathring{\gamma}_k^T$ posted by every principal $k \neq j$ is invariant in signals sent and messages received past round 1, every agent i's behavior in $G^{\mathring{S}\mathring{M}T}$ is predicated on the behavior that agent i would have followed in the subgame of $G^{\mathring{S}\mathring{M}}$ in which every principal $k \neq j$ posts the mechanism $C_k(\mathring{\gamma}_k^T)$ constructed as in Case 1, and principal j posts a mechanism $D_j(\mathring{\gamma}_j^T)$ that we shall now construct.

To this end, let us first define a *simple j-message plan* for agent i in $G^{\mathring{S}\mathring{M}T}$ as a sequence $\tilde{\pi}_j^{iT} \equiv (\tilde{\pi}_j^i(t))_{t=1}^T$, where $\tilde{\pi}_j^i(1) \in \mathring{M}_j^i(1)$ and, for each $t < \infty$, $1 < t \leq T$, $\tilde{\pi}_j^i(t) : \times_{\tau=2}^t \mathring{S}_j^i(\tau) \to \mathring{M}_j^i(t)$ is a Borel-measurable function. That is, a simple j-message plan specifies a message from agent i to principal j in round 1, and messages from agent i to principal j at rounds $t < \infty$, $1 < t \leq T$, as functions of the signals sent by principal j to agent i from round 2 to round t. Such a j-message plan is simple in that it does not explicitly condition on $(\xi^i, \mathring{s}^i(1), \omega^i)$; furthermore, contrary to the message plans defined in the proof of Theorem 2(i), simple j-message plans specify messages at all rounds, including round 1, and only to principal j. Now, using a sequence of independent draws $\varepsilon_j^T \equiv ((\varepsilon_j^i(t))_{i=1}^I)_{t=2}^T$ from the uniform distribution over $[0,1]$, we may, exactly as in (S.40)–(S.41), define the functions $\eta_{\mathring{\sigma}_j^T}(t)$ and $\eta_{\mathring{\sigma}_j^T}^T$, which provide an alternative representation of the sequence of principal j's signals past round 1. As in (S.42), we denote by $\mathring{\chi}_{\mathring{\gamma}_j^T}(\mathring{s}_j(1), \varepsilon_j^T, \mathring{m}_j^T)$ the (possibly random) decision taken by principal j given her round-1 signals $\mathring{s}_j(1)$, her sequence of signals $\eta_{\mathring{\sigma}_j^T}^T(\mathring{s}_j(1), \varepsilon_j^T, \mathring{m}_j^T)$ past round 1, and the agents' reports $\mathring{m}_j^T$. We can now represent by

$$\psi_{\mathring{\gamma}_j^T}\big(\mathring{s}_j(1), \varepsilon_j^T, \tilde{\pi}_j^T\big) \equiv \mathring{\chi}_{\mathring{\gamma}_j^T}\big(\mathring{s}_j(1), \varepsilon_j^T, \big(\big(\tilde{\pi}_j^i(t)\big(\tilde{h}_j^i(t)\big)\big)_{i=1}^I\big)_{t=1}^T\big) \in \Delta(X_j) \tag{S.57}$$

the decision of principal j induced by her mechanism $\mathring{\gamma}_j^T$ and the profile $\tilde{\pi}_j^T \equiv (\tilde{\pi}_j^{iT})_{i=1}^I$ of agents' simple j-message plans given $(\mathring{s}_j(1), \varepsilon_j^T)$, where the j-private histories of signals $(\tilde{h}_j^i(t))_{t=1}^T$ for all agents i are recursively defined, under $(\mathring{\gamma}^T, \tilde{\pi}_j^T)$, by

$$\tilde{h}_j^i(1) \equiv \{\emptyset\}, \tag{S.58}$$

$$\tilde{h}_j^i(t+1) \equiv \Big(\mathring{h}^i(t), \eta_{\mathring{\sigma}_j^T}^i(t+1)\big(\mathring{s}_j(1), \varepsilon_j^T, \big(\big(\tilde{\pi}_j^k(\tau)\big(\tilde{h}_j^k(\tau)\big)\big)_{k=1}^I\big)_{\tau=1}^t\big)\Big), \quad t < \infty, 1 \leq t \leq T-1.$$

(Recall that, according to a simple j-message plan, the messages of agent i at rounds $t < \infty$, $1 \leq t \leq T$, do not depend on principal j's round-1 signal to agent i, but only on her signals past round 1.) Under the assumptions of the theorem, we can with no loss of generality assume that there exists an ordinal $0 < \tilde{\alpha} < \omega_1$ such that $\tilde{\pi}_j^i(t) \in \overline{B}_{\tilde{\alpha}}(\times_{\tau=2}^t \mathring{S}_j^i(\tau))$ for all i and $t < \infty$, $1 < t \leq T$. Then, by (S.57) and Lemma S.1(i), for any such i and t, there exists a universal function $\overline{U}_{j,\tilde{\alpha}}^i(t)$ for $\overline{B}_{\tilde{\alpha}}(\times_{\tau=2}^t \mathring{S}_j^i(\tau))$ such that, for any such $\tilde{\pi}_j^i(t)$, there exists a code $\overline{\theta}^i_{j,\tilde{\alpha}}(t)(\tilde{\pi}^i_j(t)) \in [0,1]$ such that

$$\begin{aligned}\tilde{\pi}^i_j(t)\big(\tilde{h}^i_j(t)\big) &= \overline{U}^i_{j,\tilde{\alpha}}(t)\big(\tilde{h}^i_j(t), \overline{\theta}^i_{j,\tilde{\alpha}}(t)\big(\tilde{\pi}^i_j(t)\big)\big) \\ &= \overline{U}^i_{j,\tilde{\alpha}}(t)\big(\tilde{h}^i_j(t), (\iota^T_j)^{-1}(t) \circ \overline{\theta}^{iT}_{j,\tilde{\alpha}}\big(\tilde{\pi}^{iT}_j\big)\big), \quad \tilde{h}^i(t) \in \tilde{h}^i_j(t) \in \tilde{H}^i_j(t), \end{aligned} \tag{S.59}$$

where $\tilde{H}^i_j(t)$ is the set of agent i's j-private histories of signals at round t, $\iota^T_j : [0,1]^T \to [0,1]$ is a Borel isomorphism and $\overline{\theta}^{iT}_{j,\tilde{\alpha}}(\tilde{\pi}^{iT}_j) \equiv \iota^T_j((\overline{\theta}^i_{j,\tilde{\alpha}}(t)(\tilde{\pi}^i_j(t)))^T_{t=1}) \in [0,1]$ is the code for the simple j-message plan $\tilde{\pi}^{iT}_j$. For all i and j, let $\Theta^{iT}_{j,\tilde{\alpha}} \equiv [0,1]$ be the set of such codes and $\Theta^T_{j,\tilde{\alpha}} \equiv \times^I_{i=1} \Theta^{iT}_{j,\tilde{\alpha}}$. In light of (S.57)–(S.59), and taking advantage of the fact that the function $\chi_{\mathring{\gamma}^T_j}$ and the functions $\overline{U}^i_{j,\tilde{\alpha}}(t)$ and $\eta^i_{\mathring{\sigma}^T_j}(t+1)$, $t < \infty$, $1 \leq t \leq T$, are Borel-measurable, we obtain that there exists a Borel-measurable function $\overline{\psi}_{\mathring{\gamma}^T_j} : \mathring{S}_j(1) \times [0,1]^{I\times(T-1)} \times \Theta^T_{j,\tilde{\alpha}} \to \Delta(X_j)$ such that

$$\begin{aligned}\psi_{\mathring{\gamma}^T_j}\big(\mathring{s}_j(1), \varepsilon^T_j, \tilde{\pi}^T_j\big) &= \overline{\psi}_{\mathring{\gamma}^T_j}\big(\mathring{s}_j(1), \varepsilon^T_j, \big(\overline{\theta}^{iT}_{j,\tilde{\alpha}}\big(\tilde{\pi}^{iT}_j\big)\big)^I_{i=1}\big), \\ &\qquad \big(\mathring{s}_j(1), \varepsilon^T_j, \tilde{\pi}^T_j\big) \in \mathring{S}_j(1) \times [0,1]^{I\times(T-1)} \times \bigtimes^I_{i=1} \overline{\Pi}^{iT}_{j,\tilde{\alpha}}, \end{aligned} \tag{S.60}$$

where $\overline{\Pi}^{iT}_{j,\tilde{\alpha}}$ is the set of simple j-message plans $\tilde{\pi}^{iT}_j \equiv (\tilde{\pi}^i_j(t))^T_{t=1}$ of agent i such that $\tilde{\pi}^i_j(t) \in \overline{B}_{\tilde{\alpha}}(\times^t_{\tau=2} \mathring{S}^i_j(\tau))$ for all i and $t < \infty$, $1 < t \leq T$.

We are now ready to define the mechanism $D_j(\mathring{\gamma}^T_j)$ of principal j in $G^{\mathring{S}\mathring{M}}$ associated to her mechanism $\mathring{\gamma}^T_j$ in $G^{\mathring{S}\mathring{M}T}$. Recall from Step 6 of part (i) the embedding $a^i_j : \mathring{S}^i_j(1) \times \Xi^i_j \to \mathring{S}^i_j$, and the j-component e_j of the push-forward function constructed in the proof of Lemma S.3. First, in line with (S.47), the distribution $\mathring{\sigma}_j$ of $\mathring{s}_j$ in $D_j(\mathring{\gamma}^T_j)$ is the push-forward of the measure $\mathring{\sigma}_j(1) \otimes \bigotimes^I_{i=1} \mathrm{d}\xi^i_j$ by the mapping $(\mathring{a}^i_j)^I_{i=1} : \mathring{S}_j(1) \times \times^I_{i=1} \Xi^i_j \to \mathring{S}_j$. Second, the extended decision rule $\mathring{\phi}_j$ in $D_j(\mathring{\gamma}^T_j)$ is defined by

$$\begin{aligned}\mathring{\phi}_j(\mathring{s}_j, \mathring{m}_j) \equiv \overline{\psi}_{\mathring{\gamma}^T_j}\left(\big(\mathring{a}^{i-}_{j(s)}(\mathring{s}^i_j)\big)^I_{i=1}, e_j\left(\left\{\sum^I_{i=1} \mathring{a}^{i-}_{j(\xi)}(\mathring{s}^i_j)\right\}\right), (\iota^{iT}_j(\mathring{m}^i_j))^I_{i=1}\right), \\ (\mathring{s}_j, \mathring{m}_j) \in \mathring{S}_j \times \mathring{M}_j, \end{aligned} \tag{S.61}$$

where, for each i, $\iota^{iT}_j : \mathring{M}^i_j \to \Theta^{iT}_{j,\tilde{\alpha}}$ is a Borel isomorphism that associates a code for agent i's simple j-message plan in $G^{\mathring{S}\mathring{M}T}$ to any message $\mathring{m}^i_j = (\omega^i, \mathring{s}^i_{-j})$ that agent i can send in $G^{\mathring{S}\mathring{M}}$. (It is clear that $\mathring{\phi}_j$ thus defined is Borel-measurable.) In words, the mechanism $D_j(\mathring{\gamma}^T_j)$ first uncovers the information $(\mathring{s}^i_j(1), \xi^i_j) \equiv (\mathring{a}^i_j)^{-1}(\mathring{s}^i_j)$ from the signal $\mathring{s}^i_j$ disclosed to every agent i. By construction of $\mathring{\sigma}_j$, the random variables $(\xi^i_j)^I_{i=1}$ thus obtained are independent and uniformly distributed over $[0,1]$. Using the procedure described in Step 4, $D_j(\mathring{\gamma}^T_j)$ then constructs the variable $\xi_{0,j} \equiv \{\sum^I_{i=1} \xi^i_j\}$, which is uniformly distributed over $[0,1]$ and independent of the random variables $(\xi^i_j)^I_{i=1}$. In turn, the only information solicited by principal j from every agent i is the code of the simple j-message plan that agent i would use in the subgame $\mathring{\gamma}^T$ of $G^{\mathring{S}\mathring{M}T}$ under consideration.

Finally, we construct the strategy of every agent i in the subgame $\mathring{\gamma}^T$ of $G^{\mathring{S}\mathring{M}T}$. In the continuation equilibrium of the corresponding subgame $(D_j(\mathring{\gamma}_j^T), (C_k(\mathring{\gamma}_k^T))_{k\neq j})$ of $G^{\mathring{S}\mathring{M}}$, agent i's strategy is described by the mapping

$$(\xi^i, \mathring{s}^i, \omega^i) \mapsto \mathring{\lambda}^{*i,\xi^i}\big((D_j(\mathring{\gamma}_j^T), (C_k(\mathring{\gamma}_k^T))_{k\neq j}), (\mathring{s}_j^i, (\mathring{s}_k^i)_{k\neq j}), \omega^i\big),$$

or, equivalently, according to the above construction, by the mapping

$$(\xi^i, \mathring{s}_j^i(1), (\mathring{s}_k^i)_{k\neq j}, \xi_j^i, \omega^i) \mapsto \mathring{\lambda}^{*i,\xi^i}\big((D_j(\mathring{\gamma}_j^T), (C_k(\mathring{\gamma}_k^T))_{k\neq j}), (\mathring{a}_j^i(\mathring{s}_j^i(1), \xi_j^i), (\mathring{s}_k^i)_{k\neq j}), \omega^i\big).$$

For each i, the random variables ξ^i and ξ_j^i are independently and uniformly distributed over $[0,1]$, with ξ^i controlled by agent i and ξ_j^i controlled by principal j. We now show how to recover from ξ^i and ξ_j^i a sampling variable for agent i in the subgame $\mathring{\gamma}^T$ of $G^{\mathring{S}\mathring{M}T}$. As in the proof of Lemma S.3, we can interlace the digits in the binary expansions in Y of ξ^i and ξ_j^i to obtain a number $\iota(\xi^i, \xi_j^i) \equiv \sum_{n\geq 1} \frac{1}{2^{2n-1}} f_{|Y n}^{-1}(\xi^i) + \sum_{n\geq 1} \frac{1}{2^{2n}} f_{|Y n}^{-1}(\xi_j^i)$. Clearly, ι is a Borel isomorphism, with inverse $\iota^{-1} \equiv (\iota_1^-, \iota_2^-)$. Notice also that ξ^i and ξ_j^i are independent and uniformly distributed over $[0,1]$ if and only if $\iota(\xi^i, \xi_j^i)$ is uniformly distributed over $[0,1]$. We may thus take $\xi^{iT} \equiv \iota(\xi^i, \xi_j^i)$ as agent i's sampling variable in $\mathring{\gamma}^T$. Now, define

$$L^i(\xi^{iT}, \mathring{s}^i(1), \omega^i) \equiv \mathring{\lambda}^{*i,\iota_1^-(\xi^{iT})}\big((D_j(\mathring{\gamma}_j^T), (C_k(\mathring{\gamma}_k^T))_{k\neq j}), (\mathring{a}_j^i(\mathring{s}_j^i(1), \iota_2^-(\xi^{iT})), (\mathring{s}_k^i(1))_{k\neq j}), \omega^i\big),$$
$$(\xi^{iT}, \mathring{s}^i(1), \omega^i) \in \Xi^i \times \mathring{S}^i(1) \times \Omega^i. \tag{S.62}$$

Thus the function L^i associates, to all ξ^{iT}, $\mathring{s}^i(1)$, and ω^i, a vector of messages $(\mathring{m}_j^i, (\mathring{m}_k^i)_{k\neq j}) \in \mathring{M}^i$. As we now show, these messages allow us to construct a strategy for agent i in $G^{\mathring{S}\mathring{M}T}$.

(i) For $k \neq j$, we take the corresponding message to be the one sent by agent i to principal k in round 1 of the subgame $\mathring{\gamma}^T$ of $G^{\mathring{S}\mathring{M}T}$. That is, we define

$$\mathring{\lambda}_k^{*i,\xi^{iT}}(1)(\mathring{\gamma}^T, \mathring{s}^i(1), \omega^i) \equiv L_k^i(\xi^{iT}, \mathring{s}^i(1), \omega^i), \quad (\xi^{iT}, \mathring{s}^i(1), \omega^i) \in \Xi^i \times \mathring{S}^i(1) \times \Omega^i, \tag{S.63}$$

and we let agent i send arbitrary messages to principal k past round 1.

(ii) As for principal j, observe that $L_j^i(\xi^{iT}, \mathring{s}^i(1), \omega^i)$ identifies, up to the isomorphism ι_j^{iT}, a code in $\Theta_{j,\tilde{\alpha}}^{iT}$, which pins down a simple j-message plan $\tilde{\varpi}_j^{iT}(\iota_j^{iT}(L_j^i(\xi^{iT}, \mathring{s}^i(1), \omega^i))) \in \overline{\Pi}_{j,\tilde{\alpha}}^{iT}$. The behavior of agent i vis-à-vis principal j is then determined by this j-message plan. That is, for each $t < \infty$, $1 \leq t \leq T$, we let

$$\mathring{\lambda}_j^{*i,\xi^{iT}}(t)(\mathring{\gamma}^T, \mathring{h}^i(t), \omega^i) \equiv \iota_j^{iT}\big(L_j^i(\xi^{iT}, \mathring{s}^i(1), \omega^i)\big)(t)(\tilde{h}_j^i(t)),$$
$$(\xi^{iT}, \mathring{h}^i(t), \omega^i) \in \Xi^i \times \mathring{H}^i(t) \times \Omega^i, \tag{S.64}$$

where $\mathring{s}_j^i(1)$ together with $\tilde{h}_j^i(t)$ reflect the signals received by agent i from principal j along his private history $\mathring{h}^i(t)$.

Case 3 Finally, in any subgame $\mathring{\gamma}^T$ of $G^{\mathring{S}\mathring{M}T}$ in which at least two principals post mechanisms with decision rules responding to signals and messages past round 1, the agents' candidate equilibrium strategies prescribe the same behaviors as those in some fixed SPBE of $G^{\mathring{S}\mathring{M}T}$. (It should be noted that this is where we crucially use the assumption that $G^{\mathring{S}\mathring{M}T}$ admits an SPBE.)

This completes the description of every agent i's candidate equilibrium strategy $\mathring{\lambda}^{*iT}$ in $G^{\mathring{S}\mathring{M}T}$. Proceeding as in the proof of Lemma S.5, we obtain that, for each $t < \infty$, $1 \leq t \leq T$, $\mathring{\lambda}^{*i}(t)$ is $(\mathcal{B}([0,1]) \otimes \hat{\mathbf{\Sigma}}(1) \otimes \mathring{\mathcal{H}}^i(t) \otimes 2^{\Omega^i}, \hat{\mathcal{M}}^i(t))$-measurable, as requested.

We let $\mathring{\lambda}^{*T} \equiv (\mathring{\lambda}^{*iT})_{i=1}^I$ be the profile of agents' candidate equilibrium strategies in $G^{\mathring{S}\mathring{M}T}$.

Step 3: Outcome Equivalence of $(\mathring{\gamma}^{*T}, \mathring{\lambda}^{*T})$ **and** $(\mathring{\gamma}^*, \mathring{\lambda}^*)$ That the strategy profiles $(\mathring{\gamma}^{*T}, \mathring{\lambda}^{*T})$ and $(\mathring{\gamma}^*, \mathring{\lambda}^*)$ are outcome-equivalent follows from the description of the principals' strategies in Step 1 and the description of the agents' strategies in Case 1 of Step 2.

Step 4: Equilibrium Properties of $\mathring{\lambda}^{*T}$ We distinguish three cases. In each case, we study the incentives of some agent i, assuming that the other agents stick to their candidate equilibrium strategies $\mathring{\lambda}^{*-iT}$.

Case 1 Suppose first that $\mathring{\gamma}^T$ is a subgame of $G^{\mathring{S}\mathring{M}T}$ in which the extended decision rule $\mathring{\phi}_j^T$ in the mechanism $\mathring{\gamma}_j^T$ posted by any principal j is invariant in signals sent and messages received past round 1. If all agents but agent i ignore the signals sent by the principals past round 1, then it is optimal for agent i to do the same. Because the distribution of signals in every mechanism $C_j(\mathring{\gamma}_j^T)$ is the same as the distribution of round-1 signals in the mechanism $\mathring{\gamma}_j^T$, and the mechanism $C_j(\mathring{\gamma}_j^T)$ implements the same decisions conditional on signals and messages as the mechanism $\mathring{\gamma}_j^T$ conditional on round-1 signals and messages, the candidate equilibrium strategies for the agents described in Case 1 of Step 2 form a BNE of $\mathring{\gamma}^T$.

Case 2 Suppose next that $\mathring{\gamma}^T$ is a subgame of $G^{\mathring{S}\mathring{M}T}$ in which there exists some j such that the extended decision rule $\mathring{\phi}_j^T$ in the mechanism $\mathring{\gamma}_j^T$ posted by principal j is not invariant in signals sent or messages received past round 1, whereas the extended decision rule $\mathring{\phi}_k^T$ in the mechanism $\mathring{\gamma}_k^T$ posted by every principal $k \neq j$ is invariant in signals sent and messages received past round 1. We first claim that the strategy profiles $(\mathring{\gamma}^T, \mathring{\lambda}^{*T})$ in $G^{\mathring{S}\mathring{M}T}$ and $((D_j(\mathring{\gamma}_j^T), (C_k(\mathring{\gamma}_k^T))_{k \neq j}), \mathring{\lambda}^*)$ in $G^{\mathring{S}\mathring{M}}$ are outcome-equivalent. Indeed, the allocation $z_{\mathring{\gamma}^T, \mathring{\lambda}^{*T}}$ induced by $(\mathring{\gamma}^T, \mathring{\lambda}^{*T})$ in $G^{\mathring{S}\mathring{M}T}$ satisfies

$$\begin{aligned} & z_{\mathring{\gamma}^T, \mathring{\lambda}^{*T}}(x \,|\, \omega) \\ & = \int_{\times_{i=1}^I \Xi^i} \int_{\times_{l=1}^J \mathring{S}_l(1)} \int_{[0,1]^{I \times (T-1)}} \psi_{\mathring{\gamma}_j^T}\Big(\mathring{s}_j(1), \varepsilon_j^T, \big(\tilde{\varpi}_j^i\big(\iota_j^{iT}(L_j^i(\xi^{iT}, \mathring{s}^i(1), \omega^i))\big)\big)_{i=1}^I\Big)(x_j) \end{aligned}$$

$$\prod_{k\neq j}\mathring{\phi}_k\Big(\mathring{s}_k(1),\big(L^i_k(\xi^{iT},\mathring{s}^i(1),\omega^i)\big)_{i=1}^I\Big)(x_k)\bigotimes_{i=1}^I\bigotimes_{t=2}^T \mathrm{d}\varepsilon^i_j(t)\bigotimes_{l=1}^J\mathring{\sigma}_l(1)(\mathrm{d}\mathring{s}_l(1))\bigotimes_{i=1}^I\mathrm{d}\xi^{iT}$$

$$=\int_{\times_{i=1}^I\Xi^i}\int_{\times_{i=1}^I\Xi^i_j}\int_{\times_{l=1}^J\mathring{S}_l(1)}\int_{[0,1]^{I\times(T-1)}}$$

$$\overline{\psi}_{\mathring{\gamma}^T_j}\left(\mathring{s}_j(1),e_j\left(\left\{\sum_{i=1}^j\xi^i_j\right\}\right),\big(\iota^{iT}_j(L^i_j(\iota(\xi^i,\xi^i_j),\mathring{s}^i(1),\omega^i))\big)_{i=1}^I\right)(x_j)$$

$$\prod_{k\neq j}\mathring{\phi}_k\Big(\mathring{s}_k(1),\big(L^i_k(\iota(\xi^i,\xi^i_j),\mathring{s}^i(1),\omega^i)\big)_{i=1}^I\Big)(x_k)\bigotimes_{l=1}^J\mathring{\sigma}_l(1)(\mathrm{d}\mathring{s}_l(1))\bigotimes_{i=1}^I\mathrm{d}\xi^i_j\bigotimes_{i=1}^I\mathrm{d}\xi^i$$

$$=\int_{\times_{i=1}^I\Xi^i}\int_{\times_{l=1}^J\mathring{S}_l}$$

$$\overline{\psi}_{\mathring{\gamma}^T_j}\left(\big(\mathring{a}^{i-}_{j(s)}(\mathring{s}^i_j)\big)_{i=1}^I,e_j\left(\left\{\sum_{i=1}^I\mathring{a}^{i-}_{j(\xi)}(\mathring{s}^i_j)\right\}\right),\right.$$

$$\left.\big(\iota^{iT}_j\big(\mathring{\lambda}^{*i,\xi^i}_j\big((D_j(\mathring{\gamma}^T_j),(C_k(\mathring{\gamma}^T_k))_{k\neq j}),(\mathring{s}^i_l)_{l=1}^J,\omega^i\big)\big)\big)_{i=1}^I\right)(x_j)$$

$$\prod_{k\neq j}\mathring{\phi}_k\Big(\mathring{s}_k,\big(\mathring{\lambda}^{*i,\xi^i}_k\big((D_j(\mathring{\gamma}^T_j),(C_l(\mathring{\gamma}^T_l))_{l\neq j}),(\mathring{s}^i_l)_{l=1}^J,\omega^i\big)\big)_{i=1}^I\Big)(x_k)\bigotimes_{l=1}^J\mathring{\sigma}_l(\mathrm{d}\mathring{s}_l)\bigotimes_{i=1}^I\mathrm{d}\xi^i$$

$$=\int_{\times_{i=1}^I\Xi^i}\int_{\times_{l=1}^J\mathring{S}_l}\prod_{l=1}^J\mathring{\phi}_l\Big(\mathring{s}_l,\big(\mathring{\lambda}^{*i,\xi^i}_l\big((D_j(\mathring{\gamma}^T_j),(C_k(\mathring{\gamma}^T_k))_{k\neq j}),(\mathring{s}^i_l)_{l=1}^J,\omega^i\big)\big)_{i=1}^I\Big)(x_l)$$

$$\bigotimes_{l=1}^J\mathring{\sigma}_l(\mathrm{d}\mathring{s}_l)\bigotimes_{i=1}^I\mathrm{d}\xi^i$$

$$=z_{(D_j(\mathring{\gamma}^T_j),(C_k(\mathring{\gamma}^T_k))_{k\neq j}),\mathring{\lambda}^*}(x\,|\,\omega)\tag{S.65}$$

for all $(x,\omega)\in X\times\Omega$, where the first equality follows from (S.63)–(S.64) along with the invariance, for $k\neq j$, of $\mathring{\phi}^T_k$ in signals sent and messages received past round 1, the second equality follows from (S.60) along with the definitions of $\xi_{0,j}$, e_j, ι, and ξ^{iT}, and the fact that we can choose $\overline{\theta}^{iT}_{j,\tilde{\alpha}}$ and $\tilde{\varpi}^{iT}$ so that $\overline{\theta}^{iT}_{j,\tilde{\alpha}}\circ\tilde{\varpi}^{iT}_j=\mathrm{Id}_{\Theta^{iT}_{j,\tilde{\alpha}}}$, the third inequality follows from (S.62) and the change-of-variable formula for push-forward measures, the fourth equality follows from (S.61), and the last equality follows from (5) and the definitions of $D_j(\mathring{\gamma}^T_j)$ and $(C_k(\mathring{\gamma}^T_k))_{k\neq j}$. Thus $z_{\mathring{\gamma}^T,\mathring{\lambda}^{*T}}=z_{(D_j(\mathring{\gamma}^T_j),(C_k(\mathring{\gamma}^T_k))_{k\neq j}),\mathring{\lambda}^*}$, as claimed.

For each i, if agent i does not deviate from $\mathring{\lambda}^{*iT}$ following principal j's unilateral deviation to $\mathring{\gamma}^T_j$, then the allocation implemented in $G^{\mathring{S}\mathring{M}T}$ is given by (S.65). That agent i cannot be better off unilaterally deviating from (S.63) vis-à-vis one or several of the principals $k\neq j$, or from (S.64) vis-à-vis principal j then follows from the following observation. Let $(\mathring{m}^i_k(1))_{k\neq j}$ and $\tilde{\pi}^{iT}_j$ be the round-1 messages to the non-deviating principals and the simple j-message plan corresponding to agent i's deviation. (Recall that the messages to principals $k\neq j$ at rounds $t<\infty$, $1<t\leq T$, play no role; hence we do not describe them here.) By definition of

the mechanisms $(D_j(\mathring{\gamma}_j^T), (C_k(\mathring{\gamma}_k^T))_{k\neq j})$, the allocation that agent i induces in $G^{\mathring{S}\mathring{M}T}$ under the deviation corresponds to the one the agent can induce in $G^{\mathring{S}\mathring{M}}$ by sending to every principal $k \neq j$ the message $\mathring{m}_k^i = \mathring{m}_k^i(1)$ and to principal j the message $\mathring{m}_j^i = (\iota_j^{iT})^{-1}(\overline{\theta}_{j,\tilde{\alpha}}^{iT}(\tilde{\pi}_j^{iT}))$. Because sending these messages $(\mathring{m}_j^i, (\mathring{m}_k^i(1))_{k\neq j})$ is feasible in $G^{\mathring{S}\mathring{M}}$, we conclude from the optimality of agent i's equilibrium strategy $\mathring{\lambda}^{*i}$ in $G^{\mathring{S}\mathring{M}}$ that, when the other agents follow their equilibrium strategies $\mathring{\lambda}^{*-iT}$ in $G^{\mathring{S}\mathring{M}T}$, agent i can do no better than playing according to (S.63)–(S.64) vis-à-vis principals $k \neq j$ and j.

Case 3 Suppose finally that $\mathring{\gamma}^T$ is a subgame of $G^{\mathring{S}\mathring{M}T}$ in which at least two principals post mechanisms with decision rules responding to signals and messages past round 1. Then, because the agents' candidate equilibrium strategies prescribe the same behaviors as those in some fixed SPBE of $G^{\mathring{S}\mathring{M}T}$, no agent i has an incentive to deviate.

Step 5: Equilibrium Properties of $\mathring{\gamma}^{*T}$ We now check that, given the agents' strategy profile $\mathring{\lambda}^{*T}$, the strategy profile $\mathring{\gamma}^{*T}$ is a NE in the principals' game. By Step 3, the allocation induced by $\mathring{\gamma}^{*T}$ and $\mathring{\lambda}^{*T}$ in $G^{\mathring{S}\mathring{M}T}$ coincides with the allocation induced by $\mathring{\gamma}^*$ and $\mathring{\lambda}^*$ in $G^{\mathring{S}\mathring{M}}$. Moreover, by Case 2 of Step 4, if some principal j unilaterally deviates from $\mathring{\gamma}^{*T}$ by posting a mechanism $\mathring{\gamma}_j^T$, the allocation induced by $(\mathring{\gamma}_j^T, \mathring{\gamma}_{-j}^{*T})$ and $\mathring{\lambda}^{*T}$ in $G^{\mathring{S}\mathring{M}T}$ coincides with the allocation induced by $((D_j(\mathring{\gamma}_j^T), \gamma_{-j}^*)$ and $\mathring{\lambda}^*$ in $G^{\mathring{S}\mathring{M}}$. Because $(\mathring{\gamma}^*, \mathring{\lambda}^*)$ is a PBE of $G^{\mathring{S}\mathring{M}}$, it follows that no principal j can profitably deviate from $\mathring{\gamma}_j^{*T}$ in $G^{\mathring{S}\mathring{M}T}$ given the other principals' strategy profile $\mathring{\gamma}_{-j}^{*T}$ and the agents' strategy profile $\mathring{\lambda}^{*T}$. Thus $(\mathring{\gamma}^{*T}, \mathring{\lambda}^{*T})$ is an SPBE of $G^{\mathring{S}\mathring{M}T}$ that is outcome-equivalent to $(\mathring{\gamma}^*, \mathring{\lambda}^*)$; in particular, it is p-truthful.

Step 6: Towards an SPBE of $G^{\hat{S}\hat{M}T}$ Following arguments analogous to those in in the proof of Theorem 1(i), one can verify that, starting from the equilibrium $(\mathring{\mu}^{*T}, \mathring{\gamma}^{*T})$ of $G^{\mathring{S}\mathring{M}T}$, there exists an outcome-equivalent equilibrium in the game $G^{\hat{S}\hat{M}T}$. To see this, it suffices to notice that the signal spaces are the same in both mechanisms, and that the message spaces $\mathring{M}_j^i(t)$ and $\hat{M}_j^i(t)$ are Borel-isomorphic for all i, j, and $t < \infty$, $1 \leq t \leq T$.

This concludes the proof of Theorem 2(ii), establishing robustness. Hence the result. ■

References

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

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

[3] Bogachev, V.I. (2007): *Measure Theory*, Vol. 1 and 2. Berlin, Heidelberg, New York: Springer.

[4] Cichoń, J., and M. Morayne (1988): "Universal Functions and Generalized Classes of Functions," *Proceedings of the American Mathematical Society*, 102(1), 83–89.

[5] Dudley, R.M. (2004): *Real Analysis and Probability.* Cambridge, UK: Cambridge University Press.

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

[7] Kechris, A.S. (1995): *Classical Descriptive Set Theory.* New York: Springer Verlag.

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

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

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

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

[12] Steinhaus, H. (1930): "Sur la probabilité de la convergence de séries. Première communication," *Studia Mathematica*, 2(1), 21–39.

[13] Šupina, J., and D. Uhrik (2010): "On a Lindenbaum Composition Theorem," *Tatra Mountains Mathematical Publications*, 74, 145–158.

[14] Young, W.H. (1911): "On a New Method in the Theory of Integration," *Proceedings of the London Mathematical Society*, 2(1), 15–50.

[15] Young, W.H. (1913): "On Functions and their Associated Sets of Points," *Proceedings of the London Mathematical Society*, 2(1), 260–287.