Information design in multi-stage games

Miltiadis Makris & Ludovic Renou

Working Paper No. 861 June 2018 ISSN 1473-0278

School of Economics and Finance

Queen Mary
University of London

Information design in multi-stage games*

Miltiadis Makris† & Ludovic Renou‡

June 25, 2018

Abstract

We consider multi-stage games, where at each stage, players receive private signals about past and current states, past actions and past signals, and choose actions. We fully characterize the distributions over actions, states, and signals that obtain in any (sequential) communication equilibrium of any expansion of multi-stage games, i.e., when players can receive additional signals about past and current states, past actions, and past and current signals (including the additional past signals). We interpret our results as revelation principles for information design problems. We apply our characterization to bilateral bargaining problems.

Keywords: multi-stage games, information design, communication equilibrium, sequential communication equilibrium, information structures, Bayes correlated equilibrium, revelation principle.

JEL Classification Numbers: C73, D82.

---

*Ludovic Renou acknowledges the financial support of the Agence Nationale pour la Recherche, ANR Cigne (ANR-15-CE38-0007-01). We thank Laura Doval, Stephen Morris, Sujoy Mukerji, Alessandro Pavan, and Alex Wolitzky for insightful comments and the audiences at the many seminars we have given. We are particularly indebted to Tristan Tomala for his generosity with time, pointed discussions and perspective comments.

†Department of Economics, University of Southampton, `mmakris.econ(at)gmail.com`.

‡Queen Mary University of London and University of Adelaide, `lrenou.econ(at)gmail.com`

1 Introduction

This paper studies information design problems in dynamic games. Dynamic games offer novel and unique opportunities for information design: information about the past, the present, and even the future can be disclosed. As an example, consider the refinancing operations of central banks. Typically, central banks organize weekly tender auctions to provide short-term liquidities. Besides the choice of the auction format, another important design question is how much information to release from one auction to the next. Should central banks disclose all past bids? Should central banks disclose summary statistics of past bids and nothing else? Should central banks disclose their internal forecasts of liquidity needs? As another example, committees, boards of shareholders, members of parliaments and international organizations repeatedly vote on issues ranging from the setting of interest rates to CEO remunerations through new legislations. Again, besides the choice of the voting rules, another important design question is how much information to release from one vote to the next. Should all past votes made public? Should only the fraction of votes in favor of a reform made public? These questions echo the main theme of the paper: what are the implications of providing economic agents with additional information in dynamic problems?

To address this question, we consider multi-stage games, as in Myerson (1986) and Forges (1986). In a multi-stage game, a set of players interact over several stages and, at each stage, players receive private signals – referred to as base signal – about past and current (payoff-relevant) states, past actions and past signals, and choose actions. Repeated games and, more generally, stochastic games are examples of multi-stage games. The approach we follow is to fix a multi-stage game, which we call the *base game*, and to consider expansions of the base game, i.e., multi-stage games where players receive *additional* signals about past and current states, past actions, and past signals, including the additional ones, at each stage. We view any expansion of the base game as the result of a choice of an information structure by an information designer.

Our main contribution is to characterize the distributions over states, base signals and actions induced by *all* equilibria of *all* expansions of the base game. Our first characterization theorem (Theorem 1) states an equivalence between (i) the set of all distributions over states, signals and actions induced by all *communication equilibria* of all expansions of the base game, (ii) the set of all distributions over states, signals and actions induced by all *Bayes-Nash equilibria* of all expansions of the base game, and (iii) the set of all distributions over states, base signals and actions induced by all *Bayes correlated equilibria* of the base game. At a Bayes correlated equilibrium of the base game, at each stage, an "omniscient" mediator makes private recommendations of actions to players, conditional on past and current states and signals, past actions and past recommendations. In other words, the mediator makes recommendations at each history of the base game. Moreover, at each stage, players have an incentive to be obedient, if they have never disobeyed in the past and expect others to have been obedient in the past and to continue to be in the future.

We provide an additional characterization theorem, which extends Theorem 1 to other equilibrium concepts, all imposing *sequential rationality*. Our second characterization theorem (Theorem 2) states an equivalence between (i) the set of all distributions over actions, signals and states induced by all *sequential communication equilibria* of all expansions of the base game, (ii) the set of all distributions over states, signals and actions induced by all *conditional probability perfect Bayesian equilibria* of all expansions of the base game, and (iii) the set of all distributions over actions, base signals and states induced by all *sequential Bayes correlated equilibria* of the base game.[1] The definition of a sequential Bayes correlated equilibrium is intricate (as is the definition of a sequential communication equilibrium and of a conditional probability perfect Bayesian equilibrium) and better left to the main body of the paper. It suffices to say that it requires a player to be obedient, even at histories where some players have disobeyed in the past.

We interpret our characterization theorems as *revelation principles for information design*.[2] Indeed, let us say that a distribution over states, base signals and actions is implementable if there exist an expansion of the base game (corresponding to the choice of an information structure by

[1] Sugaya and Wolitzky (2017) formally introduce the concept of conditional probability perfect Bayesian equilibrium, as a strengthening of weak perfect Bayesian equilibrium.

[2] We view information design as the *unrestricted* choice of information structures. In mechanism design, the choice of a mechanism also entails some choice of information structures. E.g., choosing an English auction to allocate an asset entails disclosing information about bidders' valuation through the observation of their bids. However, unlike information design, the information disclosure is restricted; agent must be incentivized. In other words, in mechanism design, only incentive-compatible information structures are feasible.

the information designer) and an equilibrium of the expansion, which induces that distribution. Theorem 1 (resp., Theorem 2) then states that a distribution over actions, signals and states is implementable if and only if it is a Bayes correlated equilibrium distribution (resp., a sequential Bayes correlated equilibrium distribution). As the revelation principle for mechanism design, our characterization theorems allow us to focus on the distributions to be implemented, rather than the information structures implementing them.

We like to stress two noteworthy features of our analysis. First, any solution concept, which nests the concept of Bayes-Nash equilibrium and is nested by the concept of communication equilibrium, e.g., normal-form correlated equilibrium (Aumann, 1974) or extensive-form correlated equilibrium (Forges, 1986), also leads to the same characterization as in Theorem 1. A similar remark applies to Theorem 2. Second, we need a restriction on the set of possible expansions, which we call *admissibility*, for our results to be valid. Admissibility subsumes two important properties of an expansion. First, admissibility states that the additional signals a player receive at a stage only depend on past actions, past and current signals (including the past additional signals), and past and current states. In other words, admissibility guarantees measurability with respect to the natural filtration on the histories induced by the multi-stage game. Second, it states that the additional signals have *no causal effects* on the realizations of the states and base signals. Put it differently, it guarantees that the additional signals are just signals, i.e., additional pieces of information. Admissibility thus explains why our "omniscient" mediator only needs to know what has been realized up to a stage before making recommendations at that stage. While this may seem an unnecessary condition, it is actually essential. Without assuming admissibility, our characterizations fail. (See Section 3 for an in-depth discussion.)

The paper contributes to the literature on information design, pioneered by Kamenica and Gentzkow (2011), and recently surveyed by Bergemann and Morris (2018). The closest paper to ours is Bergemann and Morris (2016). These authors characterize the set of distributions over actions, signals and states induced by all *Bayes-Nash* equilibria of all expansions of *static* base games, and show the equivalence with the distributions induced by the Bayes correlated equilibria of the static base games. The present paper generalizes their work to dynamic problems, a non-trivial generalization. Indeed, in dynamic games, there are many possible generalizations of the concept of *Bayes correlated equilibrium*, as introduced by Bergemann and Morris (2016). The generalizations differ according to when the mediator is active, what it knows when it is active, and what it recommends when it is active. For instance, a possible generalization was to have the mediator recommend *strategies* at the first stage and to be inactive at all other stages. Another possible generalization was to have the mediator randomly select a profile of strategies at the first stage and to recommend actions to the players at each of their information sets, according to the selected profile of strategies. The main challenge was to obtain the correct generalization for the revelation principles to hold. We genuinely need the mediator to be active at *all* histories. In particular, even in dynamic games where we can assume that all the states and signals about the states are drawn ex-ante, we still need the mediator to make recommendations at each history. It would not be enough to have the mediator recommend strategies as a function of the realized states and signals at the first stage only.

In addition, the present paper generalizes the analysis of Bergemann and Morris to other solution concepts, such as sequential communication equilibrium, which are perhaps more appropriate in dynamic games. This generalization is particularly important for many economic applications. Bargaining problems (e.g., Bergemann, Brooks and Morris, 2016), allocation problems with aftermarkets (e.g., Calzolari and Pavan, 2006, Giavannoni and Makris, 2014, and Dworczak, 2017), dynamic persuasion problems (Ely, 2017 and Renault, Solan and Vieille, 2017) are all instances of dynamic problems, where sequential rationality is a natural requirement. For example, it is very natural to assume that a buyer accepts any offer above his valuation in bargaining problems, as Bergemann, Brooks, and Morris do, even when the offers are off the equilibrium path. In all these instances, we need to invoke Theorem 2.

Doval and Ely (2017) is another generalization of the work of Bergemann and Morris (2016) and nicely complements our own generalization. These authors take as given states, consequences and state-contingent payoffs over the consequences, and characterize all the distribution over states and consequences consistent with the players playing according to *some* extensive-form game. Our work differs from theirs in two important dimensions. First, we take as given the base game (and, thus, the order of moves). In some economic applications, it is a reasonable assumption. For instance, if we think about the refinancing operations of central banks, the auction format and their frequencies define the base game. If a first-price auction is used to allocate liquidities, it would not make sense to consider games, where another auction format is used. In other applications, this is a more problematic assumption. For instance, if we think about Brexit and the ongoing negotiations between the European Union and the United Kingdom, it is difficult to have a well-defined base game in mind. Second, unlike Doval and Ely, we are able to accommodate dynamic problems, where the evolution of states and signals is controlled by the players through their actions. This is a natural assumption in many economic problems, such as mergers with ex-ante match-specific investments or inventory problems.

Finally, this paper contributes to the literature on correlated equilibrium and its generalizations, e.g., communication equilibrium (Myerson, 1986, Forges, 1986), extensive-form correlated equilibrium (von Stengel and Forges, 2015), or Bayesian solution (Forges, 1993, 2006).[3] The concept of Bayes correlated equilibrium is a generalization of all these notions.

2 Multi-stage games

The model follows closely Myerson (1986). There is a set I of n players, who interact over $T < +\infty$ stages, numbered 1 to T. (With a slight abuse of notation, we denote T the set of stages.) At each stage, a payoff-relevant state is drawn, players receive private signals about past and current states, past private signals and actions, and choose an action. We are interested in characterizing the joint distributions over profiles of states, actions and signals, which arise as equilibria of "expansions" of the game, i.e., games where players receive additional signals.

[3]The concept of extensive-form correlated equilibrium was first introduced in Forges (1986). The concept introduced in von Stengel and Forges (2015) differs from the one in Forges (1986).

2.1 The base game

We first define the base game Γ, which corresponds to the game being played if no additional signals are given to the players. At each stage t, a state $\omega_t \in \Omega_t$ is drawn, player $i \in I$ receives the private signal $s_{i,t} \in S_{i,t}$, which may depend probabilistically on the current and past states, past signals and actions, and then chooses an action $a_{i,t} \in A_{i,t}$. All sets are non-empty and finite.

We now introduce some notations. We write $A_t = \times_{i \in I} A_{i,t}$ for the set of actions at stage t and $A = \times_{t \in T} \times_{i \in I} A_{i,t}$ for the set of profiles of actions. We let $H_{i,t} = A_{i,t-1} \times S_{i,t}$ be the set of player i's new information at the beginning of stage $t \in \{2, \ldots, T\}$, $H_{i,1} = S_{i,1}$ the set of initial information, and $H_{i,T+1} = A_{i,T}$ the set of terminal information.

We denote $p_1(h_1, \omega_1)$ the joint probability of (h_1, ω_1) at the beginning of the first stage and $p_{t+1}(h_{t+1}, \omega_{t+1} | a_t, h^t, \omega^t)$ the joint probability of (h_{t+1}, ω_{t+1}) at stage $t+1$ given that a_t is the profile of actions played at stage t and (h^t, ω^t) is the history of actions played, signals received and states realized at the beginning of stage t. We assume perfect recall and, therefore, impose that $p_{t+1}((b_t, s_{t+1}), \omega_{t+1} | a_t, h^t, \omega^t) = 0$ if $b_t \neq a_t$.

We denote $H\Omega$ the subset of $\times_{t=1}^{T+1} (\times_{i \in I} H_{i,t} \times \Omega_t)$ that consists of all terminal histories of the game, with generic element (h, ω).[4] The history (h, ω) is in $H\Omega$ if and only if there exists a profile of actions $a \in A$ such that

$$p^a(h, \omega) := p_1(h_1, \omega_1) \cdot \prod_{t \in T} p_{t+1}(h_{t+1}, \omega_{t+1} | a_t, h^t, \omega^t) > 0.$$

For any vector (h, ω), we can denote various sub-vectors: $h_i = (h_{i,1}, \ldots, h_{i,t}, \ldots, h_{i,T+1})$ the private (terminal) history of player i, $h_i^t = (h_{i,1}, \ldots, h_{i,t})$ the private history of player i at stage t, $h_t = (h_{1,t}, \ldots, h_{n,t})$ the profile of actions played at stage $t-1$ and signals received at stage t, $h^t = (h_1, \ldots, h_t)$ the history of signals and actions at stage t, $\omega = (\omega_1, \ldots, \omega_T)$ the profile of realized states, and $\omega^t = (\omega_1, \ldots, \omega_t)$ the profile of states realized up to stage t, with corresponding sets $H_i = \{h_i : (h, \omega) \in H\Omega \text{ for some } \omega\}$, $H_i^t = \{h_i^t : (h, \omega) \in H\Omega \text{ for some } \omega\}$, $H_t = \{h_t : (h, \omega) \in H\Omega \text{ for some } \omega\}$, $H^t = \{h^t : (h, \omega) \in H\Omega \text{ for some } \omega\}$, $\Omega = \{\omega : (h, \omega) \in$

[4] The set Ω_{T+1} is defined to be a singleton.

$H\Omega$ for some $h\}$, $\Omega^t = \{\omega^t : (h, \omega) \in H\Omega$ for some $h\}$. We write $H^t\Omega^t$ for the restriction of $H\Omega$ to the first t stages. We let $\widehat{H} := \times_{i \in I} H_i$ and $\widehat{H}^t := \times_{i \in I} H_i^t$. Similar notations will apply to other sets. If there is no risk of confusion, we will not formally define these additional notations.

The payoff to player i is $u_i(h, \omega)$ when the terminal history is $(h, \omega) \in H\Omega$. We assume that payoffs do not depend on the signal realizations, i.e., for any two histories $h = (a, s)$ and $h' = (a', s')$ such that $a = a'$, $u_i(h, \omega) = u_i(h', \omega)$ for all ω, for all i.[5] Throughout, we refer to the signals in S as the *base* signals.

2.2 Expansions

In an *expansion* of the base game, at each stage, players receive additional signals, which may depend probabilistically on past and current states, past and current signals (including the past additional ones), and past actions. Thus, players can receive additional information not only about the realization of past (payoff-relevant) states (such as the past valuations for objects in auction problems) but also about the past realization of actions (as in repeated games with imperfect monitoring). Formally, at each stage t, player i receives the additional private signal $m_{i,t} \in M_{i,t}$. All sets of additional signals are non-empty and finite. We define the set $M_{i,T+1}$ as a singleton.

We let $\pi_1(h_1, m_1, \omega_1)$ be the probability of (h_1, m_1, ω_1) at the first stage and

$$\pi_{t+1}(h_{t+1}, m_{t+1}, \omega_{t+1} | a_t, h^t, m^t, \omega^t)$$

the probability of $(h_{t+1}, m_{t+1}, \omega_{t+1})$, when a_t is the profile of actions played at stage t and (h^t, m^t, ω^t) is the history of actions, signals and states at the beginning of stage t.

We denote $HM\Omega$ the set of all terminal histories with $(h, m, \omega) \in HM\Omega$ if and only if there exists a profile of actions $a \in A$ such that

$$\pi^a(h, m, \omega) := \pi_1(h_1, m_1, \omega_1) \cdot \prod_{t \in T} \pi_{t+1}(h_{t+1}, m_{t+1}, \omega_{t+1} | a_t, h^t, m^t, \omega^t) > 0.$$

We use the same notations as in the base game to denote relevant sub-vectors and their corresponding sets. With a slight abuse of language, we use the word "expansion" to refer to the

[5] This is without loss of generality as we can always redefine the states to include the signals.

collection of additional signals $M := (M_{i,t})_{i,t}$ and kernels $\pi := (\pi_t)_t$, as well as to the multi-stage game $\Gamma_{M,\pi}$. For simplicity, we omit the reference to the sets of additional signals and write Γ_π for the expansion of Γ.

2.3 Admissibility

We introduce a condition, called *admissibility*, which links the distributions over profiles of actions, signals and states of the base game and its expansions.

Definition 1 *An expansion is admissible if there exist kernels $(\xi_t)_t$ such that*

$$\pi_1(h_1, m_1, \omega_1) = \xi_1(m_1|h_1, \omega_1)p_1(h_1, \omega_1),$$

for all (h_1, m_1, ω_1), and

$$\pi_{t+1}(h_{t+1}, m_{t+1}, \omega_{t+1}|a_t, h^t, m^t, \omega^t) = \xi_{t+1}(m_{t+1}|h^{t+1}, m^t, \omega^{t+1})p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t), \quad (\star)$$

for all $(a_t, h^t, m^t, \omega^t, h_{t+1}, m_{t+1}, \omega_{t+1})$, for all t.

Admissibility guarantees that the additional signals players receive at stage t depend only on states and *base* signals realized up to (including) stage t and actions and *additional* signals realized up to (excluding) stage t. In other words, admissibility guarantees measurability with respect to the natural filtration induced by the multi-stage game.[6] In addition, it states that the probability of (h_{t+1}, ω_{t+1}) at stage $t + 1$ is independent of m^{t+1}, conditionally on (a_t, h^t, ω^t), and given by $p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t)$. Thus, conditional on (a_t, h^t, ω^t), the likelihood of (h_{t+1}, ω_{t+1}) does not vary with the realization of m^t, i.e., the realized additional signals have *no causal effects* on the realization of the states and base signals. To put it simply, admissibility guarantees that additional signals are just signals, i.e., additional pieces of information. It is easy to verify that an expansion is admissible if and only if

$$\text{marg}_{H_{t+1}\times\Omega_{t+1}}\pi_{t+1}(\cdot|a_t, h^t, m^t, \omega^t) = p_{t+1}(\cdot|a_t, h^t, \omega^t),$$

[6] I.e., at each stage, states and base signals are realized, then additional signals are realized and then actions are chosen.

for all $(a_t, h^t, m^t, \omega^t)$.

In closing, it is worth noting that admissibility is stronger than *consistency*. An expansion is consistent if $\text{marg}_{H\Omega}\pi^a = p^a$ for all $a \in A$, i.e., for all $a \in A$,

$$\sum_{(m_1,\ldots,m_T)} \left(\pi_1(h_1, m_1, \omega_1) \cdot \prod_{t\in T} \pi_{t+1}(h_{t+1}, m_{t+1}, \omega_{t+1}|a_t, h^t, m^t, \omega^t)\right) = p_1(h_1, \omega_1) \cdot \prod_{t\in T} p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t). \qquad (\dagger)$$

In static problems, admissibility and consistency coincide. To see this, observe that consistency requires that $\sum_m \pi(h, m, \omega) = p(h, \omega)$ and, therefore, $\pi(h, m, \omega) = \pi(m|h, \omega)p(h, \omega)$, i.e., π is admissible. In multi-stage games, however, admissibility imposes additional restrictions on the kernels $(\pi_t)_t$ than consistency, which are fully characterized in Proposition 2 in Appendix A. The following is an example of an expansion, which is consistent but not admissible.

Example 1. There are two stages, $\Omega_1 = \Omega_2 = \{0, 1\}$, and no private signals and actions, i.e., A_1, S_1, A_2 and S_2 are singletons (for simplicity, we omit them). The states are uniformly and independently distributed, that is, $p_1(\omega_1) = p_2(\omega_2|\omega_1) = 1/2$ for all (ω_1, ω_2). Consider now the following expansion: $M_1 = \{0, 1\}$, M_2 is a singleton, $\pi_1(m_1, \omega_1) = 1/4$ for all (m_1, ω_1) and $\pi_2(\omega_2|m_1, \omega_1) = 1$ if and only if $\omega_2 = (\omega_1 + m_1) \pmod 2$. We can think of the second-stage state as the first-stage state plus a shock.

We now verify consistency. We have

$$\sum_{m_1} \pi_1(m_1, \omega_1)\pi_2(\omega_2|m_1, \omega_1) = \pi_1((\omega_2 - \omega_1) \pmod 2, \omega_1)\pi_2(\omega_2|(\omega_2 - \omega_1) \pmod 2, \omega_1) = 1/4.$$

The expansion is therefore consistent. Yet, it is not admissible since $\pi_2(\omega_2|m_1, \omega_1) \neq p_2(\omega_2|\omega_1)$ for all m_1. In this example, m_1 is not just an additional signal; it actually causes the second-stage state. It is also worth noting that the expansion also violates the measurability requirement since $m_1 = (\omega_2 - \omega_1) \pmod 2$, i.e., the probability of a first-stage signal depends on the realization of the second-stage state. ■

3 Equivalence theorems

This section contains our main results. It provides characterization theorems, which differ by the solution concepts adopted. In section 3.1, we consider first the concepts of Bayes-Nash equilibrium and communication equilibrium. This allows us to present our first characterization theorem in the simplest possible terms, without cluttering the analysis with issues such as consistency of beliefs, sequential rationality, or truthfulness and obedience at off-equilibrium path histories. Along with a number of illustrative examples, this helps us to highlight the role of admissibility. In section 3.2, we then extend our analysis to other solution concepts, which all impose sequential rationality. A reader more interested in this aspect of the analysis can directly jump to section 3.2 (and return to section 3.1 for an in-depth discussion on the role of admissibility).

3.1 A first equivalence theorem

We now present three equilibrium concepts: Bayes-Nash equilibria, communication equilibria, and Bayes correlated equilibria. Throughout, we fix an expansion Γ_π of Γ.

Bayes-Nash equilibrium. A behavioral strategy σ_i is a collection of maps $(\sigma_{i,t})_{t\in T}$, with $\sigma_{i,t}$: $H_i^t M_i^t \to \Delta(A_{i,t})$. A profile σ of behavioral strategies is a *Bayes-Nash equilibrium* of Γ_π if

$$\sum_{\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega}} u_i(\boldsymbol{h},\boldsymbol{\omega})\mathbb{P}_{\sigma,\pi}(\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega}) \geq \sum_{\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega}} u_i(\boldsymbol{h},\boldsymbol{\omega})\mathbb{P}_{(\sigma_i',\sigma_{-i}),\pi}(\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega}),$$

for all σ_i', for all i, with $\mathbb{P}_{\tilde{\sigma},\pi} \in \Delta(HM\Omega)$ denoting the distribution over profiles of actions, signals and states induced by $\tilde{\sigma}$ and π. We let $\mathcal{BNE}(\Gamma_\pi)$ be the set of distributions over $H\Omega$ induced by the Bayes-Nash equilibria of Γ_π.

Communication equilibrium. Consider the canonical Myersonian extension of the multi-stage Γ_π, denoted $\mathcal{C}(\Gamma_\pi)$. At each period t, player i observes the private signal $(h_{i,t}, m_{i,t})$, reports privately a signal $(\hat{h}_{i,t}, \hat{m}_{i,t})$ to a mediator, receives a private recommendation $\hat{a}_{i,t}$ from the mediator and chooses an action $a_{i,t}$. We let $\gamma_{i,t} : H_i^t M_i^t \times H_i^{t-1} M_i^{t-1} \times A_i^{t-1} \to \Delta(H_{i,t} M_{i,t})$ be a reporting strategy at period t and $\tau_{i,t} : H_i^t M_i^t \times H_i^t M_i^t \times A_i^t \to \Delta(A_{i,t})$ be an action strategy at period t. We denote $\gamma_{i,t}^*$ the truthful strategy and $\tau_{i,t}^*$ the obedient strategy. A (canonical) communication equilibrium of Γ_π is a collection of kernels $\mu_t : \widehat{HM}^t \times A^{t-1} \to \Delta(A_t)$ such that (γ^*, τ^*) is an equilibrium of the communication game, that is,

$$\sum_{\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega}} u_i(\boldsymbol{h},\boldsymbol{\omega})\mathbb{P}_{\mu\circ(\gamma^*,\tau^*),\pi}(\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega}) \geq \sum_{\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega}} u_i(\boldsymbol{h},\boldsymbol{\omega})\mathbb{P}_{\mu\circ((\gamma_i,\tau_i),(\gamma^*_{-i},\tau^*_{-i})),\pi}(\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega}),$$

for all (γ_i, τ_i), for all i, with $\mathbb{P}_{\mu\circ(\tilde{\gamma},\tilde{\tau}),\pi} \in \Delta(HM\Omega)$ denoting the distribution over profiles of actions, signals and states induced by $\mu \circ (\tilde{\gamma}, \tilde{\tau})$ and π. Note that we do not require the reports to the mediator to be jointly consistent, i.e., we do not require that $(\hat{h}, \hat{m}) \in HM$. Indeed, the mediator cannot force the players to make jointly consistent reports.[7] We let $\mathcal{CE}(\Gamma_\pi)$ be the set of distributions over $H\Omega$ induced by the communication equilibria of Γ_π. It is well-known that $\mathcal{BNE}(\Gamma_\pi) \subseteq \mathcal{CE}(\Gamma_\pi)$. See Myerson (1986) and Forges (1986) for more details on communication equilibria.[8]

Bayes correlated equilibrium. Consider the following mediated extension of Γ, denoted $\mathcal{M}(\Gamma)$. At each period t, player i observes the private signal $h_{i,t}$, receives a private recommendation $\hat{a}_{i,t}$ from the mediator and chooses an action $a_{i,t}$. We let $\bar{\tau}_{i,t} : H_i^t \times A_i^t \to \Delta(A_{i,t})$ be an action strategy at period t and write $\bar{\tau}^*_{i,t}$ for the obedient strategy.

A Bayes correlated equilibrium is a collection of kernels $\bar{\mu}_t : H^t\Omega^t \times A^{t-1} \to \Delta(A_t)$ such that τ^* is an equilibrium of the mediated game, that is,

$$\sum_{\boldsymbol{h},\boldsymbol{\omega}} u_i(\boldsymbol{h},\boldsymbol{\omega})\mathbb{P}_{\bar{\mu}\circ\bar{\tau}^*,p}(\boldsymbol{h},\boldsymbol{\omega}) \geq \sum_{\boldsymbol{h},\boldsymbol{\omega}} u_i(\boldsymbol{h},\boldsymbol{\omega})\mathbb{P}_{\bar{\mu}\circ(\bar{\tau}_i,\bar{\tau}^*_{-i}),p}(\boldsymbol{h},\boldsymbol{\omega})$$

for all $\bar{\tau}_i$, for all i, with $\mathbb{P}_{\bar{\mu}\circ\bar{\tau},p} \in \Delta(H\Omega)$ denoting the distribution over profiles of actions, base signals and states induced by $\bar{\mu} \circ \bar{\tau}^*$ and p. We let $\mathcal{BCE}(\Gamma)$ be the set of distributions over $H\Omega$ induced by the Bayes correlated equilibria of Γ.

It is instructive to compare the definition of Bayes correlated equilibrium and communication equilibrium. In a communication equilibrium, the mediator relies on the information provided

[7] Remember that $\widehat{HM} \supset HM$ is the set of reports the mediator can receive from the players.

[8] It is without loss of generality to restrict attention to canonical Myersonian extensions. The revelation principle holds for the concept of Bayes-Nash equilibrium: any equilibrium distribution over actions, states and signals of any Myersonian extension of Γ_π, where the sets of messages players can send to (and receive from) the mediator are arbitrary, is in $\mathcal{CE}(\Gamma_\pi)$.

by the players to make recommendations, while in a Bayes correlated equilibrium it is as if the mediator knows the realized states, actions and base signals prior to making recommendations. In fact, we can interpret a Bayes correlated equilibrium as a communication equilibrium of a fictitious game, with an additional dummy player observing the realized states, actions and base signals.

Note that the sets $\mathcal{BNE}(\Gamma_\pi)$, $\mathcal{CE}(\Gamma_\pi)$ and $\mathcal{BCE}(\Gamma)$ are non-empty, as a consequence of the finiteness of the game. We are now ready to state our first equivalence result.

Theorem 1 *We have the following equivalence:*

$$\mathcal{BCE}(\Gamma) = \bigcup_{\Gamma_\pi \text{ an admissible expansion of } \Gamma} \mathcal{CE}(\Gamma_\pi) = \bigcup_{\Gamma_\pi \text{ an admissible expansion of } \Gamma} \mathcal{BNE}(\Gamma_\pi).$$

Theorem 1 states an equivalence between (i) the set of distributions over actions, base signals and states induced by all Bayes correlated equilibria of Γ, (ii) the set of distributions over actions, base signals and states we can obtain by considering *all* communication equilibria of *all* admissible expansions of Γ, and (iii) the set of distributions over actions, base signals and states we can obtain by considering *all* Bayes-Nash equilibria of *all* admissible expansions of Γ. It is a *revelation principle* for information design. Indeed, Theorem 1 states that any distribution over actions, base signals and states a designer can implement by committing to an admissible information structure is a Bayes correlated equilibrium distribution. We can therefore focus on the Bayes correlated equilibrium distributions and abstract from the particular information structures implementing them. It mirrors the focus on incentive compatible social choice functions in mechanism design theory.

We note here that we consider *all* (admissible) expansions, echoing the approach in the mechanism design literature to consider *all* mechanisms. In particular economic applications, however, it might be reasonable to impose additional constraints on the set of expansions, e.g., to require additional signals to be public. In those instances, we can still use our results to obtain "upper bounds," and to check whether these upper bounds are achievable within the set of constrained expansions. The very same issue applies to mechanism design, and we are no different in that respect.

Theorem 1 generalizes the work of Bergemann and Morris (2016), who show the equivalence between (i) and (iii) for static games. We generalize their work not only to multi-stage games, but also to communication equilibria. It follows that any solution concept, which nests the concept of

Bayes-Nash equilibrium and is nested by the concept of communication equilibrium, also induces the set of distributions $\mathcal{BCE}(\Gamma)$. It is also worth noting that our definition of a Bayes correlated equilibrium is weaker than applying the definition of Bergemann and Morris on the strategic form of the base game, which would amount to making recommendations of strategies at the first stage, as a function of the realized states and signals. See Example 4 for an illustration. Finally, the set $\mathcal{BCE}(\Gamma)$ is convex.

The intuition for Theorem 1 is simple. Consider an expansion Γ_π of Γ and a communication equilibrium μ of Γ_π. At truthful histories (i.e., on path), the Myersonian mediator of $\mathcal{C}(\Gamma_\pi)$ knows the true profile of signals (h^t, m^t) at stage t and makes recommendations conditional on this profile of signals. In addition, it may be that this profile of signals is perfectly informative about ω^t (henceforth, the action profile a_t at stage t may be perfectly correlated with ω^t). To replicate the distribution induced by the communication equilibrium μ, we therefore need the mediator of Γ to be omniscient, i.e., to know (h^t, ω^t) at stage t. With such an omniscient mediator, we can construct kernels $(\bar{\mu}_t)_t$, which replicate the communication equilibrium μ. As for obedience, it is enough to observe that a player has more strategies in Γ_π than Γ, since he knows the additional signals. Thus, if a player has a profitable deviation from obedience in $\mathcal{M}(\Gamma)$, then he also has a profitable deviation from obedience in $\mathcal{C}(\Gamma_\pi)$, a contradiction. The converse follows from standard arguments, where recommendations are thought as additional signals.

Before applying Theorem 1, it is worth pausing over the role of admissibility. Admissibility guarantees that the omniscient mediator only needs to know the realized states, base signals and actions to replicate the distributions induced by any communication equilibrium of any expansion. Dispensing with admissibility presents a host of difficulties, which are best illustrated with the help of two simple examples.

Example 2. This example is an elaboration on Example 1. We first define the base game. There are a single player, two stages, two states at each stage $\Omega_1 = \Omega_2 = \{0, 1\}$, two base signals and actions at the first stage $S_1 = \{0, 1\} = A_1$, and no base signals and actions at the second stage (i.e. A_2 and S_2 are singletons). We assume that the player is perfectly informed about the first-stage state ω_1 and believes that the second-stage state is uniformly distributed, independently of first-stage actions, signals and states. Formally, $p_1(s_1, \omega_1) = 1/2$ if and only if $s_1 = \omega_1$ and $p_2(h_2, \omega_2|a_1, s_1, \omega_1) = 1/2$ for all $(s_1, \omega_1, a_1, h_2, \omega_2)$ such that $h_2 = a_1$. Finally, the player's payoff is one if his first-stage action matches the second-stage state, and zero otherwise.

Consider now the following expansion: $M_1 = \{0, 1\}$, M_2 is a singleton, $\pi_1(s_1, m_1, \omega_1) = 1/4$ for all (s_1, m_1, ω_1) such that $s_1 = \omega_1$, and $\pi_2(h_2, \omega_2|a_1, s_1, m_1, \omega_1) = 1$ if and only if $\omega_2 = (\omega_1 + m_1) \pmod 2$ and $h_2 = a_1$. As in Example 1, the expansion is consistent, but not admissible.

In the expansion Γ_π, the player achieves a payoff of one by perfectly correlating his first-stage action with the second-stage state. Indeed, the signal s_1 is perfectly informative of ω_1 and, together with the additional signal m_1, fully reveals the second-stage state ω_2. The player's optimal strategy is therefore to play $a_1 = 1$ (resp., $a_1 = 0$) with probability one if $(s_1 + m_1) \pmod 2 = 1$ (resp., $(s_1 + m_1) \pmod 2 = 0$).

The induced distribution over states, actions and signals is not a Bayes correlated distribution.

■

It is worth stressing out that the admissibility of an expansion depends on the description of the base game. To see this, consider an alternative interpretation of Example 2, where both states ω_1 and ω_2 are drawn at the first stage. This alternative interpretation induces a *different* base game, where the set of states at the first stage is $\mathring{\Omega}_1 = \Omega_1 \times \Omega_2$ (and there are no states at the second stage). Moreover, if we let $\mathring{p}_1(s_1, (\omega_1, \omega_2)) = 1/4$ if and only if $s_1 = \omega_1$, then the player is perfectly informed of ω_1 and believes that ω_2 is uniformly distributed, independently of first-stage signals, actions, and states. Consider now the following expansion: $M_1 = \{0, 1\}$, $\mathring{\pi}_1(s_1, m_1, (\omega_1, \omega_2)) = 1/4$ for all $(s_1, m_1, (\omega_1, \omega_2))$ such that $s_1 = \omega_1$ and $\omega_2 = (\omega_1 + m_1) \pmod 2$, and there are no additional signals at the second stage. As in Example 2, s_1 and m_1 are perfectly informative about ω_1 and ω_2, so that the player achieves a payoff of one. However, unlike Example 2, the expansion of the re-interpreted base game is *admissible*.[9] Therefore, Theorem 1 applies and the induced distribution is a Bayes-correlated distribution of the new base game. We haste to stress, however, that it is not always possible to do such an re-interpretation, as the next example demonstrates.

[9]For completeness, $\mathring{p}_2(h_2|a_1, s_1, (\omega_1, \omega_2)) = \mathring{\pi}_2(h_2|a_1, s_1, m_1, (\omega_1, \omega_2)) = 1$ if and only if $h_2 = a_1$. If we let $\xi_1(m_1|s_1, (\omega_1, \omega_2)) = 1$ if and only if $m_1 = (\omega_2 - \omega_1) \pmod 2$, then $\mathring{\pi}_1(s_1, m_1, (\omega_1, \omega_2)) = \xi_1(m_1|s_1, (\omega_1, \omega_2))\mathring{p}_1(s_1, (\omega_1, \omega_2))$, as required for admissibility.

Example 3. This example differs from Example 2 in that the second-stage state ω_2 is partially controlled by the player through his first-stage action a_1. This explains why we cannot reinterpret the model as we did above.

We first define the base game. There are a single player, two stages, two actions $A_1 = \{0, 1\}$ at the first stage, two states $\Omega_2 = \{0, 1\}$ at the second stage, and all other sets are singletons. The probabilities are: $p_2(\omega_2 = 1|a_1 = 1) = 5/6$ and $p_2(\omega_2 = 1|a_1 = 0) = 1/2$. Thus, unlike Example 2, the player's first-stage action influences the likelihood of the second-stage state. The player's payoff is one (resp., zero) if the second-stage state is zero (resp., one), regardless of his action.

Consider now the following expansion: $M_1 = \{0, 1\}$, M_2 is a singleton, $\pi_1(m_1 = 1) = 1/2$, $\pi_2(\omega_2 = 1|a_1 = 1, m_1 = 1) = 2/3$, $\pi_2(\omega_2 = 1|a_1 = 0, m_1 = 1) = 1$, $\pi_2(\omega_2 = 1|a_1 = 1, m_1 = 0) = 1$, and $\pi_2(\omega_2 = 1|a_1 = 0, m_1 = 0) = 0$. This expansion is consistent, but not admissible.

Player's optimal payoff is $2/3$ in the game Γ_π: the optimal strategy consists in playing $a_1 = 1$ (resp., $a_1 = 0$) when $m_1 = 1$ (resp., $m_1 = 0$). The player's optimal strategy consists in choosing the action that maximizes the likelihood of the second-stage state being 0. The induced distribution μ over actions and states is $\mu(a_1 = 0, \omega_2 = 0) = 1/2$, $\mu(a_1 = 0, \omega_2 = 1) = 0$, $\mu(a_1 = 1, \omega_2 = 0) = 1/6$, $\mu(a_1 = 1, \omega_2 = 1) = 1/3$. This is not a Bayes correlated distribution. In any Bayes correlated equilibrium, the probability of (a_1, ω_2) is $\overline{\mu}_1(a_1)p_2(\omega_2|a_1)$ and there is no $\overline{\mu}_1$ that induces the distribution μ.

Clearly, it is not possible to re-interpret the base game as we did above: the realization of the second-stage state cannot precede the choice of the first-stage action. This would violate measurability with respect to the natural filtration on the histories of the game. ■

To sum up, in a model where the players do not control the realizations of the states and signals about the states, we can re-interpret the model as if all states and signals are drawn ex-ante and players gradually learn about them, as we did for Example 2. This would guarantee admissibility. If, in addition, there is perfect observation of past actions, then admissibility is equivalent to consistency. See Appendix C for details. We close this section with an example illustrating how we can apply Theorem 1.

Example 4. There are two players and two stages. Player 1 is active in the first stage and chooses an action $a_1 \in A_1$; player 2 is inactive. Player 2 is active in the second stage and chooses an action $a_2 \in A_2$; player 1 is inactive. There are no base signals and states, i.e., Ω_1, S_1, Ω_2 and S_2 are singletons. We are interested in characterizing the distributions $\mu \in \Delta(A_1 \times A_2)$ as we vary the information players have. In particular, this implies varying the information player 2 has about the action chosen by player 1 before choosing his own action. Formally, we consider expansions (π_1, π_2), where $\pi_1 \in \Delta(M_1)$ and $\pi_2 : M_1 \times A_1 \to \Delta(M_2)$. In words, player 1 receives the additional signal m_1 at the first period and player 2 receives the additional signal m_2 at the second stage, which may depend on the first-period signal and action (m_1, a_1). Clearly, these expansions are admissible.

From Theorem 1, we can restrict attention to the Bayes correlated equilibria of the game. By definition, (μ_1, μ_2) is a Bayes correlated equilibrium if:

$$\sum_{a_1,a_2,\hat{a}_1,\hat{a}_2} u_1(a_1,a_2)\mu_1(\hat{a}_1)\tau_1^*(a_1|\hat{a}_1)\mu_2(\hat{a}_2|a_1,\hat{a}_1)\tau_2^*(a_2|\hat{a}_2) \geq$$
$$\sum_{a_1,a_2,\hat{a}_1,\hat{a}_2} u_1(a_1,a_2)\mu_1(\hat{a}_1)\tau_1(a_1|\hat{a}_1)\mu_2(\hat{a}_2|a_1,\hat{a}_1)\tau_2^*(a_2|\hat{a}_2),$$

for all $\tau_1 : A_1 \to \Delta(A_1)$, and

$$\sum_{a_1,a_2,\hat{a}_1,\hat{a}_2} u_2(a_1,a_2)\mu_1(\hat{a}_1)\tau_1^*(a_1|\hat{a}_1)\mu_2(\hat{a}_2|a_1,\hat{a}_1)\tau_2^*(a_2|\hat{a}_2) \geq$$
$$\sum_{a_1,a_2,\hat{a}_1,\hat{a}_2} u_2(a_1,a_2)\mu_1(\hat{a}_1)\tau_1^*(a_1|\hat{a}_1)\mu_2(\hat{a}_2|a_1,\hat{a}_1)\tau_2(a_2|\hat{a}_2),$$

for all $\tau_2 : A_2 \to \Delta(A_2)$, with τ_1^* and τ_2^* the obedient strategies. Any Bayes correlated equilibrium (μ_1, μ_2) induces a distribution $\mu \in \Delta(A_1 \times A_2)$, given by $\mu(a_1, a_2) = \mu_1(a_1)\mu_2(a_2|a_1, a_1)$ for all (a_1, a_2). Moreover, it is easy to verify that a distribution $\mu \in \Delta(A_1 \times A_2)$ is a Bayes correlated distribution if and only if the following two constraints are satisfied:

(i) For all a_1 such that $\sum_{a_2} \mu(a_1, a_2) > 0$, we have

$$\sum_{a_2} u_1(a_1,a_2)\mu(a_2|a_1) \geq \max_{a_1 \in A_1} \min_{a_2 \in A_2} u_1(a_1,a_2).$$

(ii) For all a_2 such that $\sum_{a_1} \mu(a_1, a_2) > 0$, we have

$$\sum_{a_1} u_2(a_1, a_2)\mu(a_1|a_2) \geq \sum_{a_1} u_2(a_1, a_2')\mu(a_1|a_2),$$

for all a_2'.

Condition (i) states that if player 1 is recommended to play a_1, but plays $a_1' \neq a_1$ instead, the mediator may recommend player 2 to punish player 1, i.e., to play $a_2 \in \arg\min_{a_2' \in A_2} u_1(a_1', a_2')$. Consequently, any recommendation made to player 1, which gives player 1 a payoff higher than his (pure) maxmin payoff, can be sustained as a Bayes correlated equilibrium. Condition (ii) states that all recommendations the mediator makes to player 2 must be best responses to player 2's belief about player 1's action.

For a concrete example, consider the game below (player 1 is the row player):

	L	R
T	$2, 2$	$0, 1$
B	$3, 0$	$1, 1$

The set of Bayes correlated distributions is given by

$$\{\mu : \mu(T, L) \geq \mu(B, L), \mu(B, R) \geq \mu(T, R), \mu(T, L) \geq \mu(T, R)\}.$$

Indeed, if L (resp., R) is recommended to player 2, player 2 must conjecture that player 1 played T (resp., B) with probability at least $1/2$ for L (resp., R) to be a best response. We therefore need $\mu(T, L) \geq \mu(B, L)$ and $\mu(B, R) \geq \mu(T, R)$. Moreover, the maxmin payoff to player 1 is 1. Therefore, if action T is recommended to player 1, it must be that $2\mu(T, L)/(\mu(T, L)+\mu(T, R)) \geq 1$, i.e., $\mu(T, L) \geq \mu(T, R)$. The associated payoffs are depicted in the picture below (the dark gray triangle):

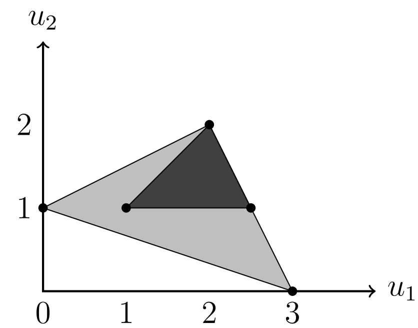

For instance, the payoff $(5/2, 1)$ corresponds to the following signalling structure and equilibrium strategies. There are two equally likely signals t and b at the first stage; player 1 is privately told the first-stage signal. There are two signals at the second-stage l and r; player 2 is privately told the second-stage signal. Player 2 receives l if and only if (T, t) and (B, b) are the first-stage profiles of signal and action. An equilibrium of that extended game consists in players playing according to their signals. This gives us the distribution $\mu(T, L) = \mu(B, L) = 1/2$ and its associated payoff $(5/2, 1)$, as required.

Finally, note that if we apply the definition of Bergemann and Morris to the strategic-form of the game, $\mu(B, R) = 1$ is the unique outcome. Indeed, if the mediator recommends a strategy to both players as a function of the realized signals and states, we simply obtain the correlated equilibria of the game, since there are no signals and states (and the strategies are the actions). This is also the unique distribution induced by the communication equilibria of the game.[10] To see this, note that it is never optimal for player 1 to play T when recommended to do so. Player 1 can disobey and play B, and report to have played T to the mediator. ■

3.2 A second equivalence theorem

The objective of this section is to enrich our analysis by requiring rational behavior both on and off the equilibrium path. The main message is that Theorem 1 generalizes to stronger solution concepts.

[10]This is also the unique distribution induced by the extensive-form correlated equilibria of the game, as defined by von Stengel and Forges (2015).

An important tool in modeling off-equilibrium path beliefs is the concept of conditional probability systems (henceforth, CPS). Fix a finite non-empty set $\mathcal{X}$. A conditional probability system $\tilde{\beta}$ on $\mathcal{X}$ is a function from $2^{\mathcal{X}} \times 2^{\mathcal{X}} \setminus \{\emptyset\}$ to $[0,1]$, which satisfies three properties: for all X, Y, Z with $X \subseteq \mathcal{X}$, $Y \subseteq \mathcal{X}$ and $\emptyset \neq Z \subseteq \mathcal{X}$,

(i) $\tilde{\beta}(Z|Z) = 1$ and $\tilde{\beta}(\mathcal{X}|Z) = 1$,

(ii) if $X \cap Y = \emptyset$, then $\tilde{\beta}(X \cup Y|Z) = \tilde{\beta}(X|Z) + \tilde{\beta}(Y|Z)$,

(iii) if $X \subseteq Y \subseteq Z$ and $Y \neq \emptyset$, then $\tilde{\beta}(X|Z) = \tilde{\beta}(X|Y)\tilde{\beta}(Y|Z)$.

Conditional probability systems capture the idea of "conditional beliefs" even after zero-probability events. In particular, if $\mathcal{X}$ is the set of terminal histories of a game, a conditional probability system induces a belief system, i.e., a belief over histories at each information set of a player. A conditional probability system also captures the beliefs players have about the strategies and beliefs of others. Finally, using a conditional probability system to represent the players' beliefs imposes that all differences in beliefs come from differences in information. We refer the reader to Myerson (1986) for more on conditional probability systems.[11]

We now present three equilibrium concepts: sequential equilibrium, sequential communication equilibrium and sequential Bayes correlated equilibrium. Throughout, we fix an admissible expansion Γ_π of Γ. Remember that there exist $(\xi_t)_t$ such that

$$\pi_t(h_t, m_t, \omega_t|a_{t-1}, h^{t-1}, m^{t-1}, \omega^{t-1}) = \xi_t(m_t|h^t, m^{t-1}, \omega^t)p_t(h_t, \omega_t|a_{t-1}, h^{t-1}, \omega^{t-1}).$$

Sequential equilibrium. Let $\overline{\mathbb{P}}_{\sigma,\pi}(\cdot|h^t, m^t, \omega^t)$ denote the distribution over $HM\Omega$ induced by the profile of behavioral strategies σ and the expansion π, given the history (h^t, m^t, ω^t). Formally,

[11] Myerson shows that for any conditional probability system $\tilde{\beta}$, there exists a sequence of probability measures $\mathbb{P}^n$ on $\mathcal{X}$ such that (i) $\mathbb{P}^n(\{x\}) > 0$ for all $x \in \mathcal{X}$ and (ii) $\tilde{\beta} = \lim_n \mathbb{P}^n$, that is, $\tilde{\beta}(X|Y) = \lim_n \frac{\mathbb{P}^n(X \cap Y)}{\mathbb{P}^n(Y)}$ for all X, for all $Y \neq \emptyset$.

we have that

$$\overline{\mathbb{P}}_{\sigma,\pi}(\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega}|h^t,m^t,\omega^t) := \mathbb{1}\{(\boldsymbol{h}^t,\boldsymbol{m}^t,\boldsymbol{\omega}^t)=(h^t,m^t,\omega^t)\}\sigma_t(\boldsymbol{a}_t|\boldsymbol{h}^t,\boldsymbol{m}^t)\times \prod_{t'=t}^{T}\pi_{t'+1}(\boldsymbol{h}_{t'+1},\boldsymbol{m}_{t'+1},\boldsymbol{\omega}_{t'+1}|\boldsymbol{a}_{t'},\boldsymbol{h}^{t'},\boldsymbol{m}^{t'},\boldsymbol{\omega}^{t'})\sigma_{t'+1}(\boldsymbol{a}_{t'+1}|\boldsymbol{h}^{t'+1},\boldsymbol{m}^{t'+1}).$$

Note that the probability $\overline{\mathbb{P}}_{\sigma,\pi}(\cdot|h^t,m^t,\omega^t)$ is well-defined even if (h^t,m^t,ω^t) has zero probability under $\mathbb{P}_{\sigma,\pi}$, and it is equal to $\mathbb{P}_{\sigma,\pi}(\cdot|h^t,m^t,\omega^t)$ when $\mathbb{P}_{\sigma,\pi}(h^t,m^t,\omega^t)>0$. Intuitively, this probability represents the beliefs an outside observer has at (h^t,m^t,ω^t) if it is conjectured that players continue to follow their equilibrium strategies even after deviations. We adopt the convention that $\overline{\mathbb{P}}_{\sigma,\pi}(\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega}) := \overline{\mathbb{P}}_{\sigma,\pi}(\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega}|h^0,m^0,\omega^0)$.

At any given history (h^t,m^t,ω^t), player i's expected payoff is

$$U_i(\sigma|h^t,m^t,\omega^t) := \sum_{\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega}} u_i(\boldsymbol{h},\boldsymbol{\omega})\overline{\mathbb{P}}_{\sigma,\pi}(\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega}|h^t,m^t,\omega^t).$$

Finally, at any private history (h_i^t,m_i^t), player i's expected payoff is

$$U_i(\sigma,\hat{\beta}|h_i^t,m_i^t) := \sum_{\boldsymbol{h^t},\boldsymbol{m^t},\boldsymbol{\omega^t}} U_i(\sigma|\boldsymbol{h^t},\boldsymbol{m^t},\boldsymbol{\omega^t})\hat{\beta}(\boldsymbol{h^t},\boldsymbol{m^t},\boldsymbol{\omega^t}|h_i^t,m_i^t),$$

where $\hat{\beta}$ is a CPS on $HM\Omega$.[12]

A *sequential equilibrium* of Γ_π is a profile σ of behavioral strategies and a CPS $\hat{\beta}$ on $HM\Omega$, which satisfy:

(i) *Sequential optimality:* For all t, for all i, for all (h_i^t,m_i^t),

$$U_i(\sigma,\hat{\beta}|h_i^t,m_i^t) \geq U_i((\sigma_i',\sigma_{-i}),\hat{\beta}|h_i^t,m_i^t),$$

for all σ_i'.

[12]Formally, $\hat{\beta}(\boldsymbol{h^t},\boldsymbol{m^t},\boldsymbol{\omega^t}|h_i^t,m_i^t) := \hat{\beta}(\{(\bar{\boldsymbol{h}},\bar{\boldsymbol{m}},\bar{\boldsymbol{\omega}})\in HM\Omega : (\bar{\boldsymbol{h}}^{\boldsymbol{t}},\bar{\boldsymbol{m}}^{\boldsymbol{t}},\bar{\boldsymbol{\omega}}^{\boldsymbol{t}}) = (\boldsymbol{h^t},\boldsymbol{m^t},\boldsymbol{\omega^t})\}|\{(\bar{\boldsymbol{h}},\bar{\boldsymbol{m}},\bar{\boldsymbol{\omega}})\in HM\Omega : (\bar{\boldsymbol{h}}^{\boldsymbol{t}}_{\boldsymbol{i}},\bar{\boldsymbol{m}}^{\boldsymbol{t}}_{\boldsymbol{i}}) = (h_i^t,m_i^t)\})$. We also write $\hat{\beta}(\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega})$ for $\hat{\beta}(\{(\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega})\}|HM\Omega)$. Similar notations apply to all other "conditional probabilities."

(ii) *CPS consistency:* The CPS $\hat{\beta}$ is consistent with (σ, π) and admissibility, that is, for all $(h, m, \omega) \in HM\Omega$, for all (i, t),

$$\hat{\beta}(a_{i,t}|h_i^t, m_i^t) = \sigma_{i,t}(a_{i,t}|h_i^t, m_i^t),$$
$$\hat{\beta}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t) = p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t),$$
$$\hat{\beta}(m_{t+1}|h^{t+1}, m^t, \omega^{t+1}) = \xi_{t+1}(m_{t+1}|h^{t+1}, m^t, \omega^{t+1}).$$

(iii) *Independence:* There exists a sequence of fully supported strategy profiles $(\sigma^n)_n$ converging to σ such that $\hat{\beta}(h^t, m^t, \omega^t|h_i^t, m_i^t) = \lim_{n\to+\infty} \mathbb{P}_{\sigma^n,\pi}(h^t, m^t, \omega^t|h_i^t, m_i^t)$ for all (h^t, m^t, ω^t), for all t, i.

Moreover, if we substitute condition (iii) with: for all t, i, for all (h^t, m^t, ω^t),

$$\hat{\beta}(h^t, m^t, \omega^t|h_i^t, m_i^t) = \frac{\mathbb{P}_{\sigma,\pi}(h^t, m^t, \omega^t)}{\mathbb{P}_{\sigma,\pi}(h_i^t, m_i^t)},$$

whenever $\mathbb{P}_{\sigma,\pi}(h_i^t, m_i^t) > 0$, then we obtain the concept of *conditional probability perfect Bayesian equilibrium*, as introduced by Sugaya and Wolitzky (2017). We let $\mathcal{SE}(\Gamma_\pi)$ (resp., $\mathcal{CPPBE}(\Gamma_\pi)$) be the set of distributions over $H\Omega$ induced by the sequential equilibria (resp., conditional probability perfect Bayesian equilibrium) of Γ_π.

Sequential communication equilibrium (Myerson, 1986). In what follows, we use notations, which parallel the ones used in previous definitions, and thus do not rehash formal definitions. We consider Myersonian extensions $\mathcal{C}(\Gamma_\pi)$ of the game Γ_π, where at each stage the set of recommendations made to a player may be a strict subset of the set of actions available to the player. Formally, for each history $((\hat{h}_i^t, \hat{m}_i^t), \hat{a}_i^{t-1})$ of past and current reports and past recommendations, $R_{i,t}((\hat{h}_i^t, \hat{m}_i^t), \hat{a}_i^{t-1}) \subseteq A_{i,t}$ is the set of possible recommendations to player i. We refer to the function $R_{i,t}$ as the *mediation range* of player i at stage t. We denote $\mathscr{H}(R)$ the set of all terminal histories consistent with the mediation ranges in $\mathcal{C}(\Gamma_\pi)$, i.e., $(h, m, \omega, \hat{h}, \hat{m}, \hat{a}) \in \mathscr{H}(R)$ if and only if $(h, m, \omega) \in HM\Omega$ and $\hat{a}_{i,t} \in R_{i,t}((\hat{h}_i^t, \hat{m}_i^t), \hat{a}_i^{t-1})$ for all i, for all t.

We denote $\overline{\mathbb{P}}_{\mu\circ(\gamma,\tau),\pi}(\cdot|h^t, m^t, \omega^t, \hat{h}^t, \hat{m}^t, \hat{a}^t)$ the distribution over $\mathscr{H}(R)$ induced by the profile of strategies (γ, τ), the recommendation kernels μ and the expansion π, given the history

$(h^t, m^t, \omega^t, \hat{h}^t, \hat{m}^t, \hat{a}^t)$. Note that $\overline{\mathbb{P}}_{\mu\circ(\gamma,\tau),\pi}(\cdot|h^t, m^t, \omega^t, \hat{h}^t, \hat{m}^t, \hat{a}^t)$ is equal to

$$\mathbb{P}_{\mu\circ(\gamma,\tau),\pi}(\cdot|h^t, m^t, \omega^t, \hat{h}^t, \hat{m}^t, \hat{a}^t)$$

at all histories $(h^t, m^t, \omega^t, \hat{h}^t, \hat{m}^t, \hat{a}^t)$ having positive probability under $\mathbb{P}_{\mu\circ(\gamma,\tau),\pi}$.

At any history $(h^t, m^t, \omega^t, \hat{h}^t, \hat{m}^t, \hat{a}^t)$, player i's expected payoff is

$$U_i(\mu\circ(\gamma,\tau)|h^t, m^t, \omega^t, \hat{h}^t, \hat{m}^t, \hat{a}^t) := \sum_{\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega},\hat{\boldsymbol{h}},\hat{\boldsymbol{m}},\hat{\boldsymbol{a}}} u_i(\boldsymbol{h},\boldsymbol{\omega})\overline{\mathbb{P}}_{\mu\circ(\gamma,\tau),\pi}(\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega},\hat{\boldsymbol{h}},\hat{\boldsymbol{m}},\hat{\boldsymbol{a}}|h^t, m^t, \omega^t, \hat{h}^t, \hat{m}^t, \hat{a}^t).$$

Finally, at any private history $(h_i^t, m_i^t, \hat{h}_i^t, \hat{m}_i^t, \hat{a}_i^t)$, player i's expected payoff is:

$$U_i(\mu \circ (\gamma,\tau), \beta|h_i^t, m_i^t, \hat{h}_i^t, \hat{m}_i^t, \hat{a}_i^t) :=$$
$$\sum_{\boldsymbol{h^t},\boldsymbol{m^t},\boldsymbol{\omega^t},\hat{\boldsymbol{h}}^{\boldsymbol{t}},\hat{\boldsymbol{m}}^{\boldsymbol{t}},\hat{\boldsymbol{a}}^{\boldsymbol{t}}} U_i(\mu \circ (\gamma,\tau)|\boldsymbol{h^t},\boldsymbol{m^t},\boldsymbol{\omega^t},\hat{\boldsymbol{h}}^{\boldsymbol{t}},\hat{\boldsymbol{m}}^{\boldsymbol{t}},\hat{\boldsymbol{a}}^{\boldsymbol{t}})\beta(\boldsymbol{h^t},\boldsymbol{m^t},\boldsymbol{\omega^t},\hat{\boldsymbol{h}}^{\boldsymbol{t}},\hat{\boldsymbol{m}}^{\boldsymbol{t}},\hat{\boldsymbol{a}}^{\boldsymbol{t}}|h_i^t, m_i^t, \hat{h}_i^t, \hat{m}_i^t, \hat{a}_i^t),$$

where β is a CPS on $\mathscr{H}(R)$. Expected payoffs are defined similarly at the private histories $(h_i^t, m_i^t, \hat{h}_i^{t-1}, \hat{m}_i^{t-1}, \hat{a}_i^{t-1})$.

Recall that (γ^*, τ^*) denotes the profile of truthful and obedient strategies. We write $\Gamma_i^{*,t}$ for the subset of reporting strategies of player i, where player i is honest up to (including) stage t, i.e., an element of $\Gamma_i^{*,t}$ is of the form

$$(\gamma_{i,1}^* \dots \gamma_{i,t-1}^*, \gamma_{i,t}^*, \gamma_{i,t+1}, \dots, \gamma_{i,T}),$$

for some $(\gamma_{i,t+1}, \dots, \gamma_{i,T})$. Similarly, we write $\mathcal{T}_i^{*,t}$ for the subset of action strategies of player i, where player i is obedient up to (including) stage t.

A collection of recommendation kernels $\mu_t : \widehat{HM}^t \times A^{t-1} \to \Delta(A_t)$ is a *sequential communication equilibrium* of Γ_π if there exist mediation ranges $R_{i,t} : H_i^t \times M_i^t \times A_i^{t-1} \to 2^{A_{i,t}} \setminus \{\emptyset\}$ for all (i,t) and a CPS β on $\mathscr{H}(R)$ such that the following are satisfied:

(i) *Honesty:* For all t, for all i, for all private histories $(h_i^t, m_i^t, \hat{h}_i^{t-1}, \hat{m}_i^{t-1}, \hat{a}_i^{t-1})$ such that (a) $(h_i^{t-1}, m_i^{t-1}) = (\hat{h}_i^{t-1}, \hat{m}_i^{t-1})$ and (b) $\hat{a}_{i,t'} \in R_{i,t}((\hat{h}_i^{t'}, \hat{m}_i^{t'}), \hat{a}_i^{t'-1})$ for all $t' < t$,

$$U_i(\mu\circ(\gamma^*,\tau^*), \beta|h_i^t, m_i^t, \hat{h}_i^{t-1}, \hat{m}_i^{t-1}, \hat{a}_i^{t-1}) \geq U_i(\mu\circ((\gamma_i,\tau_i),(\gamma_{-i}^*,\tau_{-i}^*)), \beta|h_i^t, m_i^t, \hat{h}_i^{t-1}, \hat{m}_i^{t-1}, \hat{a}_i^{t-1}),$$

for all $(\gamma_i, \tau_i) \in \Gamma_i^{*,t-1} \times \mathcal{T}_i^{*,t-1}$.

(ii) *Obedience:* For all t, for all i, for all private histories $(h_i^t, m_i^t, \hat{h}_i^t, \hat{m}_i^t, \hat{a}_i^t)$ such that (a) $(h_i^t, m_i^t) = (\hat{h}_i^t, \hat{m}_i^t)$ and (b) $\hat{a}_{i,t'} \in R_{i,t}((\hat{h}_i^{t'}, \hat{m}_i^{t'}), \hat{a}_i^{t'-1})$ for all $t' \leq t$,

$$U_i(\mu \circ (\gamma^*, \tau^*), \beta | h_i^t, m_i^t, \hat{h}_i^t, \hat{m}_i^t, \hat{a}_i^t) \geq U_i(\mu \circ ((\gamma_i, \tau_i), (\gamma_{-i}^*, \tau_{-i}^*)), \beta | h_i^t, m_i^t, \hat{h}_i^t, \hat{m}_i^t, \hat{a}_i^t),$$

for all $(\gamma_i, \tau_i) \in \Gamma_i^{*,t} \times \mathcal{T}_i^{*,t-1}$.

(iii) *CPS consistency:* β is consistent with $((\gamma^*, \tau^*), \mu, \pi)$ and admissibility, that is, for all $(h, m, \omega, \hat{h}, \hat{m}, \hat{a}) \in \mathscr{H}(R)$, for all (i, t),

$$\begin{aligned}
&\beta(a_{i,t} | h_i^t, m_i^t, \hat{h}_i^t, \hat{m}_i^t, \hat{a}_i^t) = \tau_{i,t}^*(a_{i,t} | h_i^t, m_i^t, \hat{h}_i^t, \hat{m}_i^t, \hat{a}_i^t), \\
&\beta(h_{t+1}, \omega_{t+1} | a_t, h^t, \omega^t) = p_{t+1}(h_{t+1}, \omega_{t+1} | a_t, h^t, \omega^t), \\
&\beta(m_{t+1} | h^{t+1}, m^t, \omega^{t+1}) = \xi_{t+1}(m_{t+1} | h^{t+1}, m^t, \omega^{t+1}), \\
&\beta(\hat{a}_t | \hat{h}^t, \hat{m}^t, \hat{a}^{t-1}) = \mu_t(\hat{a}_t | \hat{h}^t, \hat{m}^t, \hat{a}^{t-1}),
\end{aligned}$$

and

$$\hat{\beta}(h^t, m^t, \omega^t, \hat{h}^t, \hat{m}^t, \hat{a}^t | h_i^t, m_i^t, \hat{h}_i^t, \hat{m}_i^t, \hat{a}_i^t) = \frac{\mathbb{P}_{\mu \circ (\gamma^*, \tau^*), \pi}(h^t, m^t, \omega^t, \hat{h}^t, \hat{m}^t, \hat{a}^t)}{\mathbb{P}_{\mu \circ (\gamma^*, \tau^*), \pi}(h_i^t, m_i^t, \hat{h}_i^t, \hat{m}_i^t, \hat{a}_i^t)},$$

whenever $\mathbb{P}_{\mu \circ (\gamma^*, \tau^*), \pi}(h_i^t, m_i^t, \hat{h}_i^t, \hat{m}_i^t, \hat{a}_i^t) > 0$.

In a sequential communication equilibrium, a player has an incentive to be honest and obedient at all histories consistent with the mediation ranges, at which the player has been truthful in the past. Put it differently, even at histories where some players have lied in the past or been disobedient, a player has an incentive to be honest and obedient, if he has never lied in the past (but may have been disobedient). It is worth stressing that the recommendation kernel μ_t maps to $\Delta(A_t)$ and does not implicitly take into account the mediation ranges. Yet, the consistency requirement implies that the probability of observing a history of recommendations and reports, inconsistent with the mediation ranges, is zero when the players are honest and obedient. Indeed, since the CPS is defined over $\mathscr{H}(R)$, consistency implies that

$$\sum_{\hat{a}_t \in \times_i R_{i,t}((\hat{h}_i^t, \hat{m}_i^t), \hat{a}_i^{t-1})} \mu_t(\hat{a}_t | \hat{h}^t, \hat{m}^t, \hat{a}^{t-1}) = \sum_{\hat{a}_t \in \times_i R_{i,t}((\hat{h}_i^t, \hat{m}_i^t), \hat{a}_i^{t-1})} \beta(\hat{a}_t | \hat{h}^t, \hat{m}^t, \hat{a}^{t-1}) = 1,$$

which gives us the desired result, that is,

$$\sum_{(\hat{h},\hat{m},\hat{a}):\, [\forall (i,t):\, \hat{a}_{i,t} \in R_{i,t}((\hat{h}_i^t,\hat{m}_i^t),\hat{a}_i^{t-1})]} \mathbb{P}_{\mu\circ(\gamma^*,\tau^*),\pi}(\hat{h},\hat{m},\hat{a}) = 1.$$

It is worth noting that a recommendation $\hat{a}_t$ can have probability zero under μ_t at some history of reports and recommendations and yet be in the mediation ranges.[13] We let $\mathcal{SCE}(\Gamma_\pi)$ be the set of distributions over $H\Omega$ induced by the sequential communication equilibria of Γ_π.

Few remarks are worth making. First, our definition of sequential communication equilibrium differs slightly from Myerson (1986) in that we assume that the mediator randomizes over recommendations at every stage (a behavioral strategy), while Myerson assumes that the mediator randomizes over *pure* recommendation kernels at the ex-ante stage (a mixed strategy). From Kuhn's theorem (viewing the mediator as a player), both definitions induce the same set $\mathcal{SCE}(\Gamma_\pi)$. Second, we have restricted attention to canonical extensions of the game Γ_π. In a recent work, Sugaya and Wolitzky (2017) prove that this is without loss of generality, that is, the revelation principle holds.[14] Third, in a sequential communication equilibrium, a player may believe that the mediator has send recommendation $\hat{a}$ at a previous stage, even though $\hat{a}$ is not in the support of the recommendation kernels. In other words, it is as if the mediator can "tremble" in a sequential

[13] We provide more details on the calculation. By definition, we have that $\beta(\hat{a}_t|\hat{h}^t,\hat{m}^t,\hat{a}^{t-1})$ is equal to

$$\beta(\{(\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega},\hat{\boldsymbol{h}},\hat{\boldsymbol{m}},\hat{\boldsymbol{a}}) \in \mathscr{H}(R) : (\hat{\boldsymbol{h}}^{\boldsymbol{t}},\hat{\boldsymbol{m}}^{\boldsymbol{t}},\hat{\boldsymbol{a}}^{\boldsymbol{t}}) = (\hat{h}^t,\hat{m}^t,\hat{a}^t)\}|\{(\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega},\hat{\boldsymbol{h}},\hat{\boldsymbol{m}},\hat{\boldsymbol{a}}) \in \mathscr{H}(R) : (\hat{\boldsymbol{h}}^{\boldsymbol{t}},\hat{\boldsymbol{m}}^{\boldsymbol{t}},\hat{\boldsymbol{a}}^{\boldsymbol{t-1}}) = (\hat{h}^t,\hat{m}^t,\hat{a}^{t-1})\}).$$

The result then follows from the observation that

$$\bigcup_{\hat{a}_t \in \times_i R_{i,t}((\hat{h}_i^t,\hat{m}_i^t),\hat{a}_i^{t-1})} \{(\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega},\hat{\boldsymbol{h}},\hat{\boldsymbol{m}},\hat{\boldsymbol{a}}) \in \mathscr{H}(R) : (\hat{\boldsymbol{h}}^{\boldsymbol{t}},\hat{\boldsymbol{m}}^{\boldsymbol{t}},\hat{\boldsymbol{a}}^{\boldsymbol{t}}) = (\hat{h}^t,\hat{m}^t,\hat{a}^t)\} =$$
$$\{(\boldsymbol{h},\boldsymbol{m},\boldsymbol{\omega},\hat{\boldsymbol{h}},\hat{\boldsymbol{m}},\hat{\boldsymbol{a}}) \in \mathscr{H}(R) : (\hat{\boldsymbol{h}}^{\boldsymbol{t}},\hat{\boldsymbol{m}}^{\boldsymbol{t}},\hat{\boldsymbol{a}}^{\boldsymbol{t-1}}) = (\hat{h}^t,\hat{m}^t,\hat{a}^{t-1})\}.$$

[14] More precisely, Sugaya and Wolitzky study general Myersonian extensions of multi-stage games, where at each stage players report arbitrary messages to the mediator and receive arbitrary messages from the mediator. For the concept of conditional probability perfect Bayesian equilibrium, they show that any equilibrium distribution over states, signals and actions is in $\mathcal{SCE}(\Gamma_\pi)$. Conversely, any distribution in $\mathcal{SCE}(\Gamma_\pi)$ is an equilibrium distribution of the canonical extension of Γ_π.

communication equilibrium. We will return to that point at the end of this section and in Section 4. Finally, conditional probability perfect Bayesian equilibria are sequential communication equilibria, which in turn are communication equilibria.

Sequential Bayes correlated equilibrium. We consider mediated extensions $\mathcal{M}(\Gamma)$ of the game Γ, where at each stage the set of recommendations made to a player may be a strict subset of the set of actions available to the player. Formally, for each private history $(h_i^t, \hat{a}_i^{t-1})$ of past and current base signals, past actions and past recommendations, $\overline{R}_{i,t}(h_i^t, \hat{a}_i^{t-1}) \subseteq A_{i,t}$ is the set of possible recommendations to player i. We refer to the function $\overline{R}_{i,t}$ as the *mediation range* of player i at stage t. We denote $\overline{\mathscr{H}}(\overline{R})$ the set of all terminal histories consistent with the mediation ranges in the mediated extension $\mathcal{M}(\Gamma)$, i.e., $(h, \omega, \hat{a}) \in \overline{\mathscr{H}}(\overline{R})$ if and only if $(h, \omega) \in H\Omega$ and $\hat{a}_{i,t} \in \overline{R}_{i,t}(h_i^t, \hat{a}_i^{t-1})$ for all i, for all t.

We denote $\overline{\mathbb{P}}_{\overline{\mu}\circ\overline{\tau},p}(\cdot|h^t, \omega^t, \hat{a}^t)$ the distribution over $\overline{\mathscr{H}}(\overline{R})$ induced by the profile of strategies $\overline{\tau}$, the recommendation kernels $\overline{\mu}$ and the kernels p, given the history $(h^t, \omega^t, \hat{a}^t)$. Note that $\overline{\mathbb{P}}_{\overline{\mu}\circ\overline{\tau},p}(\cdot|h^t, \omega^t, \hat{a}^t)$ is equal to $\mathbb{P}_{\overline{\mu}\circ\overline{\tau},p}(\cdot|h^t, \omega^t, \hat{a}^t)$ at all histories $(h^t, \omega^t, \hat{a}^t)$ having positive probability under $\mathbb{P}_{\overline{\mu}\circ\overline{\tau},p}$.

At any history $(h^t, \omega^t, \hat{a}^t)$, player i's expected payoff is

$$U_i(\overline{\mu}\circ\overline{\tau}|h^t, \omega^t, \hat{a}^t) := \sum_{\boldsymbol{h},\boldsymbol{\omega},\hat{\boldsymbol{a}}} u_i(\boldsymbol{h}, \boldsymbol{\omega})\overline{\mathbb{P}}_{\overline{\mu}\circ\overline{\tau},p}(\boldsymbol{h}, \boldsymbol{\omega}, \hat{\boldsymbol{a}}|h^t, \omega^t, \hat{a}^t).$$

Finally, at any private history $(h_i^t, \hat{a}_i^t)$, player i's expected payoff is:

$$U_i(\overline{\mu}\circ\overline{\tau}, \overline{\beta}|h_i^t, \hat{a}_i^t) :=$$
$$\sum_{\boldsymbol{h^t},\boldsymbol{\omega^t},\hat{\boldsymbol{a}}^{\boldsymbol{t}}} U_i(\overline{\mu}\circ\overline{\tau}|\boldsymbol{h^t}, \boldsymbol{\omega^t}, \hat{\boldsymbol{a}}^{\boldsymbol{t}})\overline{\beta}(\boldsymbol{h^t}, \boldsymbol{\omega^t}, \hat{\boldsymbol{a}}^{\boldsymbol{t}}|h_i^t, \hat{a}_i^t),$$

where $\overline{\beta}$ is a CPS on $\overline{\mathscr{H}}(\overline{R})$. We write $\overline{\mathcal{T}}_i^{*,t}$ for the subset of action strategies of player i, where player i is obedient up to (including) stage t. We are now ready to define the concept of sequential Bayes correlated equilibrium.

A collection of recommendation kernels $\overline{\mu}_t : H^t\Omega^t \times A^{t-1} \to \Delta(A_t)$ is a *sequential Bayes correlated equilibrium* of Γ if there exist mediation ranges $\overline{R}_{i,t} : H_i^t \times A_i^{t-1} \to 2^{A_{i,t}} \setminus \{\emptyset\}$ for all (i, t) and a CPS $\overline{\beta}$ on $\overline{\mathscr{H}}(\overline{R})$ such that the following are satisfied:

(i) *Obedience:* For all t, for all i, for all private histories $(h_i^t, \hat{a}_i^t)$ such that $\hat{a}_{i,t'} \in \overline{R}_{i,t'}(h_i^{t'}, \hat{a}_i^{t'-1})$ for all $t' \leq t$,

$$U_i(\overline{\mu} \circ \overline{\tau}^*, \overline{\beta} | h_i^t, \hat{a}_i^t) \geq U_i(\overline{\mu} \circ (\overline{\tau}_i, \overline{\tau}_{-i}^*), \overline{\beta} | h_i^t, \hat{a}_i^t),$$

for all $\overline{\tau}_i \in \overline{\mathcal{T}}_i^{*,t-1}$.

(ii) *CPS consistency:* $\overline{\beta}$ is consistent with $(\overline{\tau}^*, \overline{\mu}, p)$, that is, for all $(h, \omega, \hat{a}) \in \overline{\mathscr{H}}(\overline{R})$, for all (i, t),

$$\overline{\beta}(a_{i,t} | h_i^t, \hat{a}_i^t) = \overline{\tau}_{i,t}^*(a_{i,t} | h_i^t, \hat{a}_i^t),$$
$$\overline{\beta}(h_{t+1}, \omega_{t+1} | a_t, h^t, \omega^t) = p_{t+1}(h_{t+1}, \omega_{t+1} | a_t, h^t, \omega^t),$$
$$\overline{\beta}(\hat{a}_t | h^t, \omega^t, \hat{a}^{t-1}) = \overline{\mu}_t(\hat{a}_t | h^t, \omega^t, \hat{a}^{t-1}),$$

and

$$\overline{\beta}(h^t, \omega^t, \hat{a}^t | h_i^t, \hat{a}_i^t) = \frac{\mathbb{P}_{\overline{\mu} \circ \overline{\tau}^*, p}(h^t, \omega^t, \hat{a}^t)}{\mathbb{P}_{\overline{\mu} \circ \overline{\tau}^*, p}(h_i^t, \hat{a}_i^t)},$$

whenever $\mathbb{P}_{\overline{\mu} \circ \overline{\tau}^*, p}(h_i^t, \hat{a}_i^t) > 0$.

In a sequential Bayes correlated equilibrium, a player has an incentive to be obedient at all histories consistent with the mediation ranges. It is worth noting that the recommendation kernel $\overline{\mu}_t$ maps to $\Delta(A_t)$ and does not implicitly take into account the mediation ranges. Nevertheless, the consistency requirement implies that the probability of observing a history of recommendations, actions, base signals and states, inconsistent with the mediation ranges, is zero when the players are obedient. Indeed, we have

$$\sum_{(\hat{h}, \hat{\omega}, \hat{a}) : [\forall (i,t) : \, \hat{a}_{i,t} \in \overline{R}_{i,t}(\hat{h}_i^t, \hat{a}_i^{t-1})]} \mathbb{P}_{\overline{\mu} \circ \overline{\tau}^*, p}(\hat{h}, \hat{\omega}, \hat{a}) = 1,$$

since the CPS is defined over $\overline{\mathscr{H}}(\overline{R})$. (See above for more details.) We let $\mathcal{SBCE}(\Gamma)$ be the set of distributions over $H\Omega$ induced by the sequential correlated equilibria of Γ.

With all these preliminaries done, we can now state our second equivalence result.

Theorem 2 *We have the following equivalence:*

$$\mathcal{SBCE}(\Gamma) = \bigcup_{\Gamma_\pi \text{ an admissible expansion of } \Gamma} \mathcal{SCE}(\Gamma_\pi) = \bigcup_{\Gamma_\pi \text{ an admissible expansion of } \Gamma} \mathcal{CPPBE}(\Gamma_\pi).$$

Theorem 2 states an equivalence between (i) the set of distributions over actions, base signals and states induced by all sequential Bayes correlated equilibria of Γ, (ii) the set of distributions over actions, base signals and states we can obtain by considering *all* sequential communication equilibria of *all* admissible expansions of Γ, and (iii) the set of distributions over actions, base signals and states we can obtain by considering *all* conditional probability perfect Bayesian equilibria of *all* admissible expansions of Γ. This theorem states, in effect, that any distribution over actions, base signals and states a designer can implement by committing to an admissible information structure, where no player would expect to gain by disobeying his recommendations, even after zero-probability histories (but consistent with the mediation ranges), is a sequential Bayes correlated distribution.

We close this section with few remarks. First, the set $\mathcal{SBCE}(\Gamma)$ is convex. To see this, take two distributions ν and ν' in $\mathcal{SBCE}(\Gamma)$. It follows from Theorem 2 that there exist two expansions Γ_π and $\Gamma_{\pi'}$ and two associated sequential communication equilibria $(\mu_t, \beta, R_{i,t})_{i,t}$ and $(\mu'_t, \beta', R'_{i,t})_{i,t}$, which induce ν and ν', respectively. Take $\alpha \in [0, 1]$ and consider the expansion $\Gamma_{\alpha\pi+(1-\alpha)\pi'}$, where the information structure π (resp., π') is drawn with probability α (resp., $1 - \alpha$) and the players are informed about the draw. If players coordinate on $(\mu_t, \beta, R_{i,t})_{i,t}$ (resp., $(\mu'_t, \beta', R'_{i,t})_{i,t}$) when the drawn information structure is π (resp., π'), we obtain the distribution $\alpha\nu + (1 - \alpha)\nu'$. From Theorem 2, it is in $\mathcal{SBCE}(\Gamma)$. Second, the role of mediation ranges is to ensure the existence of an equilibrium, as our economic application in Section 4 highlights. Third, we do not have the equivalence with the set of distributions over actions, base signals and states we can obtain by considering *all* sequential equilibria of *all* admissible expansions of Γ, that is,

$$\bigcup_{\Gamma_\pi \text{ an admissible expansion of } \Gamma} \mathcal{SE}(\Gamma_\pi) \neq \mathcal{SBCE}(\Gamma),$$

as the following example demonstrates.

Example 5. This example is due to Laura Doval. The base game is as follows. There are two players, labelled 1 and 2, no states and no base signals. Player 1 has two actions, a_1 and a'_1, and player 2 has two actions, a_2 and a'_2. Player 1 moves first and player 2 moves second, without observing the move of player 1. The payoffs are in the table below, with player 1 the row player.

	a_2	a_2'
a_1	$0,0$	$1,0$
a_1'	$1,0$	$-1,-1$

We first argue that for all expansions, no sequential equilibria put probability one on (a_1, a_2). Consider any expansion, with (M_1, M_2) the set of messages, and $\pi_1 \in \Delta(M_1)$ and $\pi_2 : M_1 \times A_1 \to \Delta(M_2)$ the two probability kernels. Let (σ_1, σ_2) be a profile of strategies. The induced probability of (a_1, a_2) is

$$\sum_{m_1,m_2} \pi_1(m_1)\sigma_1(a_1|m_1)\pi_2(m_2|a_1,m_1)\sigma_2(a_2|m_2).$$

Suppose it is equal to 1. Thus, we need $\sigma_1(a_1|m_1) = 1$ for all m_1 such that $\pi_1(m_1) > 0$ and $\sigma_2(a_2|m_2) = 1$ for all m_2 such that $\sum_{m_1} \pi_1(m_1)\pi_2(m_2|a_1, m_1) > 0$, if we want to induce the distribution degenerated on (a_1, a_2).

For player 1 to have an incentive to play a_1, we need that

$$\sum_{m_2}(1 - \sigma_2(a_2|m_2))\pi_2(m_2|a_1, m_1) \geq \sum_{m_2}(2\sigma_2(a_2|m_2) - 1)\pi_2(m_2|a_1', m_1)$$

for all m_1 such that $\pi_1(m_1) > 0$. This implies that

$$\sum_{m_1,m_2} (1 - \sigma_2(a_2|m_2))\pi_2(m_2|a_1, m_1)\pi_1(m_1) \geq \sum_{m_1,m_2} (2\sigma_2(a_2|m_2) - 1)\pi_2(m_2|a_1', m_1)\pi_1(m_1).$$

This is equivalent to:

$$0 \geq \sum_{m_2 : \sum_{m_1} \pi_1(m_1)\pi_2(m_2|a_1,m_1)>0} \; \sum_{m_1} \pi_2(m_2|a_1', m_1)\pi_1(m_1) +$$
$$\sum_{m_2 : \sum_{m_1} \pi_1(m_1)\pi_2(m_2|a_1,m_1)=0} (2\sigma_2(a_2|m_2) - 1) \left(\sum_{m_1} \pi_2(m_2|a_1', m_1)\pi_1(m_1) \right).$$

In words, player 2 must be playing a_2' with high enough probability after receiving off-path messages m_2, i.e., such that $\sum_{m_1} \pi_1(m_1)\pi_2(m_2|a_1, m_1) = 0$. However, for player 2 to play a_2' with strictly positive probability after any off-path message m_2, we need player 2 to believe that a_1 was played with probability one (since a_2 is weakly dominant). In a sequential equilibrium, the probability that a_1' was played given the message m_2 is given by the limit of

$$\frac{\sum_{m_1} \pi_1(m_1)\sigma_1^\varepsilon(a_1'|m_1)\pi_2(m_2|a_1', m_1)}{\sum_{m_1} \pi_1(m_1)\sigma_1^\varepsilon(a_1'|m_1)\pi_2(m_2|a_1', m_1) + \sum_{m_1} \pi_1(m_1)\sigma_1^\varepsilon(a_1|m_1)\pi_2(m_2|a_1, m_1)}$$

as $\varepsilon \to 0$, for some fully mixed strategy σ_1^ε converging to σ_1. Since $\sum_{m_1} \pi_1(m_1)\pi_2(m_2|a_1, m_1) = 0$, $\pi_1(m_1)\pi_2(m_2|a_1, m_1) = 0$ for all m_1 and, therefore, the above limit is 1. Player 2 would never play a_2' with strictly positive probability in a sequential equilibrium, hence player 1 will never play a_1. We cannot implement the outcome (a_1, a_2), as a sequential equilibrium of some expansion of the base game.

We now argue that this distribution is implementable as a sequential Bayes correlated equilibrium. We simply need that $\overline{\mu}_1(a_1) = 1$, $\overline{\mu}_2(a_2|a_1, \hat{a}_1) = \overline{\mu}_2(a_2'|a_1', \hat{a}_1) = 1$ for all $\hat{a}_1 \in \{a_1, a_1'\}$, and $\overline{\beta}(a_1|a_2') = 1$. Intuitively, we need that player 2 assigns probability 1 to a_1 upon observing the recommendation a_2', so that he optimally plays a_2' (and, thus, is obedient). To see that this is possible, consider the sequence $(\overline{\tau}^\varepsilon, \overline{\mu}^\varepsilon)$ converging to $(\overline{\tau}^*, \overline{\mu})$ as $\varepsilon \to 0$, given by $\overline{\mu}_1^\varepsilon(a_1') = \varepsilon$, $\overline{\tau}_i^\varepsilon(a_i|a_i) = \overline{\tau}_i^\varepsilon(a_i'|a_i') = 1 - \varepsilon$ for all $i \in \{1, 2\}$, $\overline{\mu}_2^\varepsilon(a_2'|a_1', \hat{a}_1) = 1 - \varepsilon$, and $\overline{\mu}_2^\varepsilon(a_2'|a_1, \hat{a}_1) = \sqrt{\varepsilon}$ for all $\hat{a}_1 \in \{a_1, a_1'\}$. Conditional on a_2', the probability of a_1' is

$$\frac{2\varepsilon(1-\varepsilon)^2}{2\varepsilon(1-\varepsilon)^2 + (\varepsilon^2 + (1-\varepsilon)^2)\sqrt{\varepsilon}}.$$

The limit is zero as $\varepsilon \to 0$. Thus, if we consider the conditional probability system $\overline{\beta}$ given by $\lim_{\varepsilon \to 0} \overline{\mathbb{P}}_{(\overline{\tau}^\varepsilon, \overline{\mu}^\varepsilon)}$, we have a sequential Bayes correlated equilibrium. ■

Fourth, if we impose additional restrictions on the conditional probability systems in the definition of a sequential Bayes correlated equilibrium, then we can obtain a partial converse (see Appendix E for a formal statement). Indeed, it is easy to see that if we require the conditional probability systems to be consistent with trembles of the players only, then this stronger version of a sequential Bayes correlated equilibrium induces a *subset* of all the distributions over actions, states and base signals that we can obtain by considering all sequential equilibria of all admissible expansions.[15] Finally, we can obtain an analogous equivalence theorem for the concept of weak

[15]More formally, the restriction requires that there exists a sequence $(\tau^n)_n$ of completely mixed strategies converging to τ^* such that $\overline{\beta} = \lim_{n \to +\infty} \overline{\mathbb{P}}_{\overline{\mu} \circ \overline{\tau}^n, p}$. Sugaya and Wolitzky (2017) call such a concept a machine sequential equilibrium.

perfect Bayesian equilibrium.[16]

4 An economic application: bilateral bargaining

We consider a variation on the work of Bergemann, Brooks and Morris (2013). There are one buyer and one seller. The seller makes an offer $a_1 \in A_1 \subset \mathbb{R}_+$ to the buyer, who observes the offer and either accepts ($a_2 = 1$) or rejects ($a_2 = 0$) it. If the buyer accepts the offer a_1, the payoff to the buyer is $\omega - a_1$, while the payoff to the seller is a_1, with ω being the buyer's valuation (the payoff-relevant state). We assume that $\omega \in \Omega \subset \mathbb{R}_{++}$. If the buyer rejects the offer, the payoff to both the seller and the buyer is normalized to zero. The buyer and the seller are symmetrically informed and believe that the state is ω with probability $p(\omega) > 0$. We assume that the set of offers the seller can make is finite, but as fine as needed. For future reference, we write ω_L for the lowest state, ω_L^- for the largest offer a_1 strictly smaller than ω_L, and ω_H for the highest state.

This model differs from Bergemann, Brooks and Morris (2013) in one important aspect. In our model, both the seller and the buyer have no initial private information about the state, while Bergemann, Brooks and Morris assume that the buyer is privately informed of the state ω. The base game of Bergemann, Brooks and Morris thus corresponds to a particular expansion of our base game. Similarly, Roesler and Szentes (2017) consider all information structures, where the buyer has some signals about his own valuation (and the seller is uninformed). (See also Condorelli and Szentes, 2018.) Unlike these two papers, we consider *all* admissible information structures. In particular, the information the buyer receives may depend on the information the seller has received as well as the offer made. In addition, the seller can be better informed than the buyer in our model.

We first characterize the Bayes correlated equilibria. The profile of recommendation kernels

[16] It is available upon request from the authors.

(μ_1, μ_2) is a Bayes correlated equilibrium if the following constraints are satisfied:

$$\sum_{a_1,a_2,\hat{a}_1,\hat{a}_2,\omega} a_1 a_2 \mu_1(\hat{a}_1|\omega)\tau_1^*(a_1|\hat{a}_1)\mu_2(\hat{a}_2|a_1,\omega,\hat{a}_1)\tau_2^*(a_2|a_1,\hat{a}_2)p(\omega) \geq$$
$$\sum_{a_1,a_2,\hat{a}_1,\hat{a}_2,\omega} a_1 a_2 \mu_1(\hat{a}_1|\omega)\tau_1(a_1|\hat{a}_1)\mu_2(\hat{a}_2|a_1,\omega,\hat{a}_1)\tau_2^*(a_2|a_1,\hat{a}_2)p(\omega),$$

for all $\tau_1 : A_1 \to \Delta(A_1)$, and

$$\sum_{a_1,a_2,\hat{a}_1,\hat{a}_2,\omega} (\omega - a_1) a_2 \mu_1(\hat{a}_1|\omega)\tau_1^*(a_1|\hat{a}_1)\mu_2(\hat{a}_2|a_1,\omega,\hat{a}_1)\tau_2^*(a_2|a_1,\hat{a}_2)p(\omega) \geq$$
$$\sum_{a_1,a_2,\hat{a}_1,\hat{a}_2,\omega} (\omega - a_1) a_2 \mu_1(\hat{a}_1|\omega)\tau_1^*(a_1|\hat{a}_1)\mu_2(\hat{a}_2|a_1,\omega,\hat{a}_1)\tau_2(a_2|a_1,\hat{a}_2)p(\omega)$$

for all $\tau_2 : A_1 \times A_2 \to \Delta(A_2)$, where τ_1^* and τ_2^* are the obedient strategies.

For future reference, note that the sum of the payoffs (the surplus) is bounded from above by $\sum_{\omega_1} \omega_1 p_1(\omega_1) := \mathbb{E}(\omega)$. Moreover, the surplus is $\mathbb{E}(\omega)$ only if trade occurs with probability one in all states.

We now reformulate these constraints. Any Bayes correlated equilibrium induces a kernel $\mu : \Omega \to \Delta(A_1 \times A_2)$, given by $\mu(a_1, a_2|\omega) = \mu_1(a_1|\omega)\mu_2(a_2|a_1,\omega,a_1)$ for all (a_1, a_2). As in Example 4, it is easy to see that the incentive constraints are then equivalent to: for all a_1 such that $\sum_{\omega,a_2} \mu(a_1, a_2|\omega)p(\omega) > 0$,

$$\sum_{\omega,a_2} a_1 a_2 \mu(a_1, a_2|\omega)p(\omega) \geq 0,$$

and for all (a_1, a_2) such that $\sum_{\omega} \mu(a_1, a_2|\omega)p(\omega) > 0$,

$$\sum_{\omega} (\omega - a_1) a_2 \mu(a_1, a_2|\omega)p(\omega) \geq \sum_{\omega} (\omega - a_1)(1 - a_2)\mu(a_1, a_2|\omega)p(\omega).$$

The first inequality states that the seller's equilibrium payoff must be higher than his min-max payoff, which is zero. The second inequality states that the recommendations made to the buyer must be optimal on the equilibrium path.

It is then immediate to check that the following payoff distributions: (i) $\mu(\omega, 1|\omega) = 1$ for all ω, (ii) $\mu(0, 1|\omega) = 1$ for all ω, and (iii) $\mu(\omega_H, 0|\omega) = 1$ for all ω, are equilibrium distributions, with corresponding payoff profiles $(\mathbb{E}(\omega), 0)$, $(0, \mathbb{E}(\omega))$, and $(0, 0)$. Moreover, in any equilibrium, the payoff to the buyer and the seller is at least zero, and the sum of payoffs is at most $\mathbb{E}(\omega)$. Since the set of Bayes correlated payoffs is convex, the set of equilibrium payoffs is therefore

$$\text{co}\,\{(0,0), (\mathbb{E}(\omega), 0), (0, \mathbb{E}(\omega))\},$$

which is depicted in the picture below.

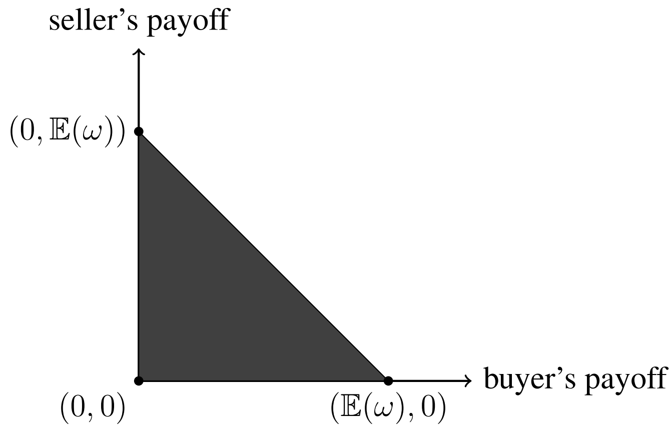

It is worth noting that the three payoff profiles $(0, 0)$, $(\mathbb{E}(\omega), 0)$ and $(0, \mathbb{E}(\omega))$ are Bayes-Nash equilibrium payoffs of the bargaining game, where the seller and the buyer are perfectly informed about the state. The set of Bayes correlated equilibrium payoffs thus corresponds to the correlated equilibria of the bargaining game with complete information. As a consequence, had we assumed that the buyer and seller had private signals about the state, the characterization of the Bayes correlated equilibrium payoffs would have been the same.

We now characterize the set of *sequential Bayes correlated equilibria*. With some abuse of notation, a sequential Bayes correlated equilibrium consists of recommendation kernels (μ_1, μ_2), a conditional probability system β, and mediation ranges (R_1, R_2), which jointly satisfy the following constraints. First, if the mediator recommends $\hat{a}_1$ to the seller, the seller must have an incentive to be obedient, i.e., for all $\hat{a}_1 \in R_1$,

$$\sum_{\omega, \hat{a}_2} \hat{a}_1 \hat{a}_2 \mu_2(\hat{a}_2 | \hat{a}_1, \omega, \hat{a}_1) \beta(\omega | \hat{a}_1) \geq \sum_{\omega, \hat{a}_2} a_1 \hat{a}_2 \mu_2(\hat{a}_2 | a_1, \omega, \hat{a}_1) \beta(\omega | \hat{a}_1)$$

for all a_1.[17] Second, if the offer made to the buyer is a_1 and the mediator recommends $\hat{a}_2 \in R_2(a_1)$ to the buyer, the buyer must have an incentive to be obedient, i.e.,

$$\sum_{\omega}(\omega - a_1)\hat{a}_2\beta(\omega|a_1, \hat{a}_2) \geq \sum_{\omega}(\omega - a_1)(1 - \hat{a}_2)\beta(\omega|a_1, \hat{a}_2).$$

Third, the conditional probability system must be consistent.

There are immediate bounds on the equilibrium payoffs: the sum of the buyer and seller's payoffs is bounded from above by $\mathbb{E}(\omega)$, the buyer's payoff is bounded from below by 0, and the seller's payoff is bounded from below by ω_L^-. The following proposition states that there are, in fact, no other restrictions on equilibrium payoffs.

Proposition 1 *The set of sequential Bayes correlated equilibrium payoffs is*

$$\mathrm{co}\,\{(\omega_L^-, 0), (\mathbb{E}(\omega), 0), (\omega_L^-, \mathbb{E}(\omega) - \omega_L^-)\}.$$

The set of equilibrium payoffs is depicted in the picture below.

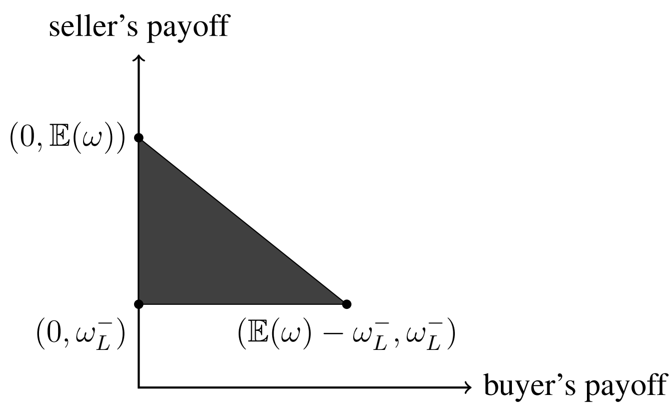

We prove this proposition in what follows. As a preliminary observation, note that the conditional probability system puts no restriction on the buyer's beliefs after observing an off-path offer a_1, i.e., an offer such that $\sum_\omega \mu_1(a_1|\omega)p(\omega) = 0$. To see this, note that $\beta(\omega, a_1) = \beta(\omega, a_1|a_1)\beta(a_1)$ for any conditional probability system β. Moreover, from the consistency of β with (p, τ^*, μ),

[17] The conditional probability system β is defined on (subsets of) $\Omega \times A_1 \times A_1 \times A_2 \times A_2$, where the first copy of A_1 and A_2 corresponds to the recommendations. To ease the notation, we write $\beta(\omega|\hat{a}_1)$ for $\beta(\{\omega\} \times \{\hat{a}_1\} \times A_1 \times A_2 \times A_2|\Omega \times \{\hat{a}_1\} \times A_1 \times A_2 \times A_2)$. Similarly, for $\beta(\omega|a_1, \hat{a}_2)$, $\beta(\omega, a_1)$, $\beta(a_1)$ and $\beta(\omega, a_1|a_1)$.

we have that $\beta(\omega, a_1) = \beta(a_1) = 0$ whenever $\sum_\omega \mu_1(a_1|\omega)p(\omega) = 0$. Therefore, $\beta(\omega, a_1|a_1)$ is arbitrary and, thus, is $\beta(\omega|a_1)$. In particular, we can assume that the buyer believes that the state is ω_L with probability one. We refer to those beliefs as the most pessimistic beliefs. Similarly, there are no restrictions on the buyer's beliefs after observing an off-path offer a_1 and a recommendation $\hat{a}_2$.

We are now ready to state how to obtain the payoff profile $(\omega_L^-, \mathbb{E}(\omega) - \omega_L^-)$. The mediator recommends the seller to offer ω_L^-, regardless of the state. If the offer ω_L^- is made, the mediator recommends the buyer to accept, regardless of the state and the recommendation made to the seller. If any offer $a_1 > \omega_L^-$ is made, the mediator recommends the buyer to reject the offer, regardless of the state and the recommendation made to the seller. Since any such offer is off-path, the buyer has an incentive to be obedient when he believes that the state is ω_L with probability one. As we have just argued, we can choose a well-defined conditional probability system capturing such beliefs. Finally, if any offer $a_1 < \omega_L^-$ is made, the mediator recommends the buyer to accept, regardless of the state and the recommendation made to the seller. Finally, the mediation ranges are $R_1 = \{\omega_L^-\}$, $R_2(a_1) = \{1\}$ if $a_1 < \omega_L$, $R_2(\omega_L) \subseteq \{0, 1\}$, and $R_2(a_1) = \{0\}$ if $a_1 > \omega_L$. It is worth pointing out that an equilibrium would fail to exist without restricting the mediation ranges. For instance, following an offer $a_1 < \omega_L$ and the recommendation to reject it, the buyer would not have an incentive to obey the recommendation. We thus need mediation ranges in order to guarantee that players indeed have an incentive to be obedient, even after off equilibrium path recommendations (that are consistent with the mediation ranges, though).

We now turn our attention to the two other payoff profiles $(\mathbb{E}(\omega), 0)$ and $(\omega_L^-, 0)$. The profile $(\mathbb{E}(\omega), 0)$ corresponds to full surplus extraction, which can be obtained with $\mu_1(\omega|\omega) = 1$ for all ω and $\mu_2(\hat{a}_2 = 1|a_1, \omega, \hat{a}_1) = 1$ for all $(a_1, \omega, \hat{a}_1)$ with $a_1 \leq \omega$ (and zero, otherwise). The mediation ranges are $R_1 = \Omega$, $R_2(a_1) = \{0\}$ if $a_1 > \omega_H$, $R_2(a_1) = \{1\}$ if $a_1 < \omega_L$, and $R_2(a_1) = \{0, 1\}$ if $a_1 \in \Omega$.

Lastly, the profile $(\omega_L^-, 0)$ is easily implementable when $\mathbb{E}(\omega) \in A_1$ (which we assume). The mediator recommends the seller to offer $\mathbb{E}(\omega)$, regardless of the state, and the buyer to accept that offer with probability $\omega_L^-/\mathbb{E}(\omega)$, on path. Off-path, we again use the most pessimistic beliefs to give the seller a payoff of zero, if he deviates.[18] The mediation ranges are $R_1 = \{\mathbb{E}(\omega)\}$, $R_2(a_1) = \{1\}$ if $a_1 < \omega_L$, $R_2(a_1) = \{0, 1\}$ if $a_1 = \mathbb{E}(\omega)$, and $R_2(a_1) = \{0\}$, otherwise. To complete the proof of Proposition 1, it is enough to invoke the bounds on the payoff profiles and the convexity of the set of sequential Bayes correlated equilibrium payoffs.

Discussion. This economic application sheds light on an important feature of our construction (and Myerson's definition of sequential communication equilibrium). As shown by Myerson, any conditional probability system is the limit of fully supported probabilities. Accordingly, for any conditional probability system β, there must exist a sequence $(\phi^n)_n$ of probabilities, fully supported on $\Omega \times A_1 \times A_1 \times A_2 \times A_2$, such that

$$\beta(\omega|a_1) = \lim_{n \to +\infty} \frac{\sum_{\hat{a}_1, \hat{a}_2, a_2} \phi^n(\omega, \hat{a}_1, a_1, \hat{a}_2, a_2)}{\sum_{\omega, \hat{a}_1, \hat{a}_2, a_2} \phi^n(\omega, \hat{a}_1, a_1, \hat{a}_2, a_2)},$$

for all (ω, a_1), where $\beta(\omega|a_1)$ is the buyer's belief about the state ω after observing the offer a_1. Thus, if we consider sequences such that

$$\lim_{n \to +\infty} \frac{\sum_{\hat{a}_1, \hat{a}_2, a_2} \phi^n(\omega, \hat{a}_1, a_1, \hat{a}_2, a_2)}{\sum_{\hat{a}_1, \hat{a}_2, a_2} \phi^n(\omega_L, \hat{a}_1, a_1, \hat{a}_2, a_2)} = 0,$$

for all $\omega \neq \omega_L$, then $\beta(\omega_L|a_1) = 1$.[19] For instance, if the buyer believes that the mediator is infinitely more likely to "tremble to recommending a_1" when the state is ω_L than in any other state, then an off-path offer of a_1 is an overwhelming signal that the state is ω_L.[20] In other words,

[18] If $\mathbb{E}(\omega) \notin A_1$, the analysis is slightly more complicated, but follows the same logic. The mediator recommends the seller to offer prices a_1 such that the expectation of ω given that a_1 is offered is exactly a_1, and recommends the buyer to randomize between accepting and rejecting with the appropriate probability.

[19] We can prove the existence of such sequences by construction. For all ω, let $A_\omega := \{\hat{a}_1 \in A_1 : \mu_1(\hat{a}_1|\omega) > 0\}$. By definition, it is non-empty. Moreover, its complement is also non-empty when there are off-path offers (since $p(\omega) > 0$ for all ω). For each n, choose $\varepsilon_n > 0$ sufficiently small. For all $\omega \neq \omega_L$, let $\mu_1^n(\hat{a}_1|\omega) = \mu_1(\hat{a}_1|\omega) - (\varepsilon_n^2/|A_\omega|)$ for all $\hat{a}_1 \in A_\omega$, $\mu_1^n(\hat{a}_1|\omega) = \varepsilon_n^2/|A_1 \backslash A_\omega|$, for all $\hat{a}_1 \in A_1 \backslash A_\omega$. Let $\mu_1^n(\hat{a}_1|\omega_L) = \mu_1(\hat{a}_1|\omega_L) - (\varepsilon_n/|A_{\omega_L}|)$ for all $\hat{a}_1 \in A_{\omega_L}$, $\mu_1^n(\hat{a}_1|\omega_L) = \varepsilon_n/|A_1 \setminus A_{\omega_L}|$, for all $\hat{a}_1 \in A_1 \setminus A_{\omega_L}$. Finally, let $\tau_1^n(\hat{a}_1|\hat{a}_1) = 1 - \varepsilon_n$ for all $\hat{a}_1$ and $\tau_1^n(a_1|\hat{a}_1) = \varepsilon_n/(|A_1| - 1)$. Choose any $\phi^n(\omega, \hat{a}_1, a_1, \hat{a}_2, a_2)$ such that $\sum_{a_2, \hat{a}_2} \phi^n(\omega, \hat{a}_1, a_1, \hat{a}_2, a_2) = p(\omega)\mu_1^n(\hat{a}_1|\omega)\tau_1^n(a_1|\hat{a}_1)$. It is then easy to verify that the ratios converge to zero as $(\varepsilon_n)_n \to 0$.

[20] Another instance is for the buyer to believe that the seller conditions his offer on both the recommendation and the state, even though the seller does not know the state, and to interpret an off-path offer of a_1 as an overwhelming signal that the state is ω_L. These beliefs may not be plausible, but they are entirely consistent with a conditional probability system.

the conditional probability system we have constructed is the limit of probabilities $\mathbb{P}_{\mu^n,\tau^n,p}$, where (μ^n, τ^n) is a fully supported strategy profile at each n.

If, however, we were to impose that the mediator does *not* tremble (and, thus, require that the conditional probability system is the limit of the sequence $(\mathbb{P}_{\mu,\tau^n,p})_n$ for some fully supported strategy τ^n), then the above construction would break down. In fact, we would lose the equivalence between the set of distributions induced by all sequential communication equilibria of all (admissible) expansions of the game and the set of distributions induced by all "trembling-free" sequential Bayes correlated equilibria of the game. To see this, we first construct an expansion of the base game with $(\omega_L^-, \mathbb{E}(\omega) - \omega_L^-)$ as a sequential equilibrium payoff profile (hence, a sequential communication equilibrium). To do so, assume that the seller is perfectly informed of the state, offers ω_L^- regardless of the state, that the buyer believes that the state is ω_L if an offer other than ω_L^- is made, and that the buyer accepts an offer if and only if it is strictly below his expectation of ω.[21]

We now argue that we *cannot* replicate this payoff profile, if our omniscient mediator does not tremble. To replicate this payoff profile, we need trade to occur with probability one in all states, since the sum of payoffs is $\mathbb{E}(\omega)$. Consequently, the only offer recommended to (and made by) the seller must be ω_L^-, i.e., we must have $\mu_1(\omega_L^-|\omega) = 1$ for all ω. If the mediator does not tremble, then the only consistent posterior is the prior. To satisfy the obedience constraint, however, we need the seller's payoff to be lower than ω_L^- after any deviation to $a_1 > \omega_L^-$. If the seller deviates and offers $a_1 \geq \mathbb{E}(\omega)$, then the seller's worst payoff is 0, since the mediator can, regardless of the state, recommend the buyer to not accept the offer (and the buyer would have an incentive to be obedient). If the seller deviates and offers $a_1 \in [\omega_L, \mathbb{E}(\omega))$, the lowest payoff the mediator can give to the seller is solution to the minization problem:

$$\min_{(\lambda_\ell, p_\ell)_{\ell=0,1}} a_1 \Big[\sum_{\ell=0,1} \lambda_\ell \mathbb{1}\big\{ \sum_{\omega} p_\ell(\omega)\omega > a_1 \big\} \Big],$$

with $(\lambda_\ell, p_\ell)_{\ell=0,1}$ a splitting of p. In words, the omniscient mediator makes recommendations to

[21] Let σ be the seller's strategy such that $\sigma(\omega_L^-|\omega) = 1$ for all ω. For all $\varepsilon > 0$ sufficiently small, let $\sigma^\varepsilon(a_1|\omega) = \frac{\varepsilon^2 p(\omega_L)}{1-p(\omega_L)}$ for all $a_1 \neq \omega_L^-$, $\omega \neq \omega_L$, $\sigma^\varepsilon(a_1|\omega_L) = \varepsilon$ for all $a_1 \neq \omega_L^-$, and $\sigma^\varepsilon(\omega_L^-|\omega) = 1 - \sum_{a_1 \neq \omega_L^-} \sigma^\varepsilon(a_1|\omega)$. By construction, $\sigma^\varepsilon \to \sigma$ as $\varepsilon \to 0$. Moreover, the posterior of ω_L given $a_1 \neq \omega_L^-$ goes to 1 as $\varepsilon \to 0$.

the buyer so as to minimize the probability of trading at price a_1; the mediator does so by choosing an optimal splitting of p (or, equivalently, an optimal kernel μ_2). It is not hard to see that a lower bound is

$$\min_{(\lambda_\ell,\mathbb{E}_\ell)_{\ell=0,1}} a_1\Big[\sum_{\ell=0,1}\lambda_\ell \mathbb{1}\big\{\mathbb{E}_\ell > a_1\big\}\Big],$$

subject to $\sum_{\ell=0,1}\lambda_\ell\mathbb{E}_\ell = \mathbb{E}(\omega)$, $\sum_{\ell=0,1}\lambda_\ell = 1$, $\lambda_\ell \in [0,1]$, and $\mathbb{E}_\ell \in [\omega_L,\omega_H]$ for all ℓ. We can think of $\mathbb{E}_\ell$ as the expectation of ω, given the recommendation ℓ. The solution to this problem is:

$$a_1\frac{\mathbb{E}(\omega)-a_1}{\omega_H - a_1}.$$

Therefore, whenever

$$\max_{a_1\in[\omega_L,\mathbb{E}(\omega))} a_1\frac{\mathbb{E}(\omega)-a_1}{\omega_H - a_1} > \omega_L^-,$$

the seller cannot be incentivized to offer ω_L^-.[22] In other words, even the harshest off-path punishment would not deter the seller from offering another price than ω_L^-. In those instances, we would not be able to replicate the payoff profile $(\omega_L^-, \mathbb{E}(\omega)-\omega_L^-)$, even though it is an equilibrium payoff of some expansion of the base game.

Appendices

A Admissibility

Proposition 2 *An expansion (M,π) is admissible if and only if*

$$\sum_{m^t,m_{t+1}} \pi_{t+1}(h_{t+1},m_{t+1},\omega_{t+1}|a_t,h^t,m^t,\omega^t)Q_t(a_t,h^t,m^t,\omega^t) = \\ p_{t+1}(h_{t+1},\omega_{t+1}|a_t,h^t,\omega^t)\left(\sum_{m^t} Q_t(a_t,h^t,m^t,\omega^t)\right),$$

for all $(a_t,h^t,\omega^t,h_{t+1},\omega_{t+1})$, for all sub-probability Q_t, for all t.

[22] Simple, albeit tedious, algebra shows that this inequality holds when $\omega_L^- < \left(\sqrt{\omega_H} - \sqrt{\omega_H - \mathbb{E}(\omega)}\right)^2$.

Proof $(\Leftarrow)$. Note that

$$\sum_{m^t, m_{t+1}} \pi_{t+1}(h_{t+1}, m_{t+1}, \omega_{t+1} | a_t, h^t, m^t, \omega^t) Q_t(a_t, h^t, m^t, \omega^t) =$$

$$\sum_{m^t, m_{t+1}} \pi_{t+1}(m_{t+1} | a_t, h^{t+1}, m^t, \omega^{t+1}) \pi_{t+1}(h_{t+1}, \omega_{t+1} | a_t, h^t, m^t, \omega^t) Q_t(a_t, h^t, m^t, \omega^t) =$$

$$\sum_{m^t} \pi_{t+1}(h_{t+1}, \omega_{t+1} | a_t, h^t, m^t, \omega^t) Q_t(a_t, h^t, m^t, \omega^t).$$

Consequently, if we choose two sub-probabilities $\overline{Q}_t$ and $\underline{Q}_t$, degenerated on arbitrary $\overline{m}^t$ and $\underline{m}^t$ respectively, we have that

$$\pi_{t+1}(h_{t+1}, \omega_{t+1} | a_t, h^t, \overline{m}^t, \omega^t) = \pi_{t+1}(h_{t+1}, \omega_{t+1} | a_t, h^t, \underline{m}^t, \omega^t) = p_{t+1}(h_{t+1}, \omega_{t+1} | a_t, h^t, \omega^t),$$

i.e., Equation $(\star)$ is satisfied, with $\xi_{t+1}(m_{t+1} | h^{t+1}, m^t, \omega^{t+1}) = \pi_{t+1}(m_{t+1} | a_t, h^{t+1}, m^t, \omega^{t+1})$ (we drop the dependence on a_t in the definition of ξ_{t+1} since conditioning on $(h^{t+1}, m^t, \omega^{t+1})$ is equivalent to conditioning on $(a_t, h^{t+1}, m^t, \omega^{t+1})$).

$(\Rightarrow)$. Immediate. $\square$

B Proof of Theorem 1

The proof is purely constructive.

$(\Rightarrow)$. We first prove that

$$\bigcup_{\Gamma_\pi \text{ an admissible expansion of } \Gamma} \mathcal{CE}(\Gamma_\pi) \subseteq \mathcal{BCE}(\Gamma).$$

Let Γ_π be an expansion of Γ and μ a communication equilibrium of Γ_π. Let $\mathbb{P}_{\mu \circ (\gamma^*, \tau^*), \pi}$ be the distribution over signals (including the additional ones), states, reports, and recommendations induced by $\mu \circ (\gamma^*, \tau^*)$ and π. Throughout, variables with a hat $\hat{}$ on top are either reports or recommendations.

The proof consists in constructing recommendation kernels $(\bar{\mu}_t : H^t \Omega^t \times A^{t-1} \to \Delta(A_t))_t$ such that the following two properties.

(i) Denoting $\mathbb{P}_{\bar{\mu}\circ\bar{\tau}^*,p}$ the distributions over actions, signals, states and recommendations induced by $\bar{\mu}\circ\bar{\tau}^*$ and p, we have that

$$\mathrm{marg}_{H\Omega}\mathbb{P}_{\bar{\mu}\circ\bar{\tau}^*,p} = \mathrm{marg}_{H\Omega}\mathbb{P}_{\mu\circ(\gamma^*,\tau^*),\pi},$$

implying that

$$\sum_{h,\hat{a},\omega} u_i(h,\omega)\mathbb{P}_{\bar{\mu}\circ\bar{\tau}^*,p}(h,\omega,\hat{a}) = \sum_{h,m,\omega,\hat{h},\hat{m},\hat{a}} u_i(h,\omega)\mathbb{P}_{\mu\circ(\gamma^*,\tau^*),\pi}(h,m,\omega,\hat{h},\hat{m},\hat{a}).$$

(ii) If there exist player i and strategy $\bar{\tau}_i$ such that

$$\sum_{h,\hat{a},\omega} u_i(h,\omega)\mathbb{P}_{\bar{\mu}\circ\bar{\tau}^*,p}(h,\omega,\hat{a}) < \sum_{h,\hat{a},\omega} u_i(h,\omega)\mathbb{P}_{\bar{\mu}\circ(\bar{\tau}_i,\bar{\tau}_{-i}),p}(h,\omega,\hat{a}),$$

then there exists a strategy τ_i such that

$$\sum_{h,m,\omega,\hat{h},\hat{m},\hat{a}} u_i(h,\omega)\mathbb{P}_{\mu\circ(\gamma^*,\tau^*),\pi}(h,m,\omega,\hat{h},\hat{m},\hat{a}) < \sum_{h,m,\omega,\hat{h},\hat{m},\hat{a}} u_i(h,\omega)\mathbb{P}_{\mu\circ(\gamma^*,(\tau_i,\tau^*_{-i})),\pi}(h,m,\omega,\hat{h},\hat{m},\hat{a}),$$

a contradiction with μ being a communication equilibrium.

In the sequel, we write $\tau_t(a_t|(h^t,m^t),(\hat{h}^t,\hat{m}^t),\hat{a}^t)$ for $\times_{i\in I}\tau_{i,t}(a_{i,t}|(h^t_i,m^t_i),(\hat{h}^t_i,\hat{m_i}^t),\hat{a}^t_i)$ and $\tau^*_t(a_t|(h^t,m^t),(\hat{h}^t,\hat{m}^t),\hat{a}^t)$ for $\times_{i\in I}\tau^*_{i,t}(a_{i,t}|(h^t_i,m^t_i),(\hat{h}^t_i,\hat{m_i}^t),\hat{a}^t_i)$. With a slight abuse of notation, we also write

$$\tau_t(a_t|(h^t,m^t),(\hat{h}^t,\hat{m}^t),\hat{a}^t)$$

for

$$\tau_{i,t}(a_{i,t}|(h^t_i,m^t_i),(\hat{h}^t_i,\hat{m_i}^t),\hat{a}^t_i)\times(\times_{j\in I\setminus\{i\}}\tau^*_{j,t}(a_{j,t}|(h^t_j,m^t_j),(\hat{h}^t_j,\hat{m_j}^t),\hat{a}^t_j)).$$

Similar notations are used for $\bar{\tau}_{i,t}$. To ease notations, we write $\mathbb{P}_{\mu\circ(\gamma^*,\tau),\pi}(h,m,\omega,\hat{a})$ for

$$\mathbb{P}_{\mu\circ(\gamma^*,\tau),\pi}(h,m,\omega,h,m,\hat{a}).$$

Since we focus on truthful histories throughout, this should not create any confusion.

As a preliminary observation, observe that with any strategy profile $(\bar{\tau}_t)_{t\in T}$ of $\mathcal{M}(\Gamma)$, we can associate a strategy profile $(\gamma^*_t,\tau_t)_{t\in T}$ of $\mathcal{C}(\Gamma_\pi)$ such that τ_t coincides with $\bar{\tau}_t$ at truthful histories and is independent of m^t, i.e.,

$$\tau_t(a_t|(h^t,m^t),(\hat{h}^t,\hat{m}^t),\hat{a}^t) = \bar{\tau}_t(a_t|h^t,\hat{a}^t)$$

for all a_t, for all histories such that $(h^t, m^t) = (\hat{h}^t, \hat{m}^t)$, for all m^t, for all t. We write $\bar{\tau}_t \equiv (\gamma_t^*, \tau_t)$ for such an association. The association is possible because players have more information in $\mathcal{C}(\Gamma_\pi)$ than in $\mathcal{M}(\Gamma)$. We associate $\bar{\tau}_t^*$ with (γ_t^*, τ_t^*).

We now define inductively the recommendation kernels $(\bar{\mu}_t : H^t\Omega^t \times A^{t-1} \to \Delta(A_t))_t$. Fix any strategy profile $\bar{\tau} := (\bar{\tau}_t)_{t\in T}$ and choose any $\tau := (\tau_t)_{t\in T}$ such that $\bar{\tau}_t \equiv (\gamma_t^*, \tau_t)$ for all t.

First, we construct $\bar{\mu}_1$ such that $\mathbb{P}_{\bar{\mu}\circ\bar{\tau},p}((a_1, h_1), \omega_1, \hat{a}_1) = \sum_{m_1} \mathbb{P}_{\mu\circ(\gamma^*,\tau),\pi}((a_1, h_1), m_1, \omega_1, \hat{a}_1)$ for all $(a_1, h_1, \omega_1, \hat{a}_1)$, that is,

$$\begin{aligned} p_1(h_1, \omega_1)\bar{\mu}_1(\hat{a}_1|h_1, \omega_1)\bar{\tau}_1(a_1|h_1, \hat{a}_1) &= \\ \sum_{m_1} \pi_1(h_1, m_1, \omega_1)\gamma_1^*(\hat{h}_1, \hat{m}_1|h_1, m_1)\mu_1(\hat{a}_1|\hat{h}_1, \hat{m}_1)\tau_1(a_1|(h_1, m_1), (\hat{h}_1, \hat{m}_1), \hat{a}_1) &= \\ \sum_{m_1} \pi_1(h_1, m_1, \omega_1)\mu_1(\hat{a}_1|h_1, m_1)\bar{\tau}_1(a_1|h_1, \hat{a}_1)&, \end{aligned}$$

where the last equality follows from $\bar{\tau} \equiv (\gamma^*, \tau)$. We therefore have that

$$\bar{\mu}_1(\hat{a}_1|h_1, \omega_1) = \frac{\sum_{m_1} \pi_1(h_1, m_1, \omega_1)\mu_1(\hat{a}_1|h_1, m_1)}{p_1(h_1, \omega_1)},$$

where $p_1(h_1, \omega_1) > 0$ since we consider histories in $H\Omega$. It is immediate to verify that $\bar{\mu}_1(\hat{a}_1|h_1, \omega_1) \geq 0$ for all $\hat{a}_1$ and, due to admissibility, $\sum_{\hat{a}_1} \bar{\mu}_1(\hat{a}_1|h_1, \omega_1) = 1$, for all (h_1, ω_1). Thus, $\bar{\mu}_1$ is well-defined. Moreover, its definition is independent of the choice of $\bar{\tau}$ and τ.

By induction, assume that $(\bar{\mu}_1, \ldots, \bar{\mu}_t)$ has been defined. We now construct $\bar{\mu}_{t+1}$ such that

$$\mathbb{P}_{\bar{\mu}\circ\bar{\tau},p}((a_{t+1}, h^{t+1}), \omega^{t+1}, \hat{a}^{t+1}) = \sum_{m^{t+1}} \mathbb{P}_{\mu\circ(\gamma^*,\tau),\pi}((a_{t+1}, h^{t+1}), m^{t+1}, \omega^{t+1}, \hat{a}^{t+1}),$$

for all $(a_{t+1}, h^{t+1}, \omega^{t+1}, \hat{a}^{t+1})$, that is,

$$\begin{aligned} &\mathbb{P}_{\bar{\mu}\circ\bar{\tau},p}((a_t, h^t), \omega^t, \hat{a}^t) \times \\ p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t)&\bar{\mu}_{t+1}(\hat{a}_{t+1}|h^{t+1}, \omega^{t+1}, \hat{a}^t)\bar{\tau}_{t+1}(a_{t+1}|h^{t+1}, \hat{a}^{t+1}) = \\ &\sum_{m^{t+1}} \Big(\mathbb{P}_{\mu\circ(\gamma^*,\tau),\pi}((a_t, h^t), m^t, \omega^t, \hat{a}^t) \times \\ \pi_{t+1}(h_{t+1}, m_{t+1}, \omega_{t+1}|a_t, h^t, m^t, \omega^t)&\gamma_{t+1}^*(\hat{h}_{t+1}, \hat{m}_{t+1}|(h^{t+1}, m^{t+1}), (\hat{h}^t, \hat{m}^t), \hat{a}^t) \\ \mu_{t+1}(\hat{a}_{t+1}|h^{t+1}, m^{t+1}, \hat{a}^t)&\tau_{t+1}(a_{t+1}|(h^{t+1}, m^{t+1}), (\hat{h}^{t+1}, \hat{m}^{t+1}), \hat{a}^{t+1})\Big), \end{aligned}$$

which is equivalent to

$$\mathbb{P}_{\bar{\mu}\circ\bar{\tau},p}((a_t, h^t), \omega^t, \hat{a}^t) \times$$
$$p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t)\bar{\mu}_{t+1}(\hat{a}_{t+1}|h^{t+1}, \omega^{t+1}, \hat{a}^t) =$$
$$\sum_{m^{t+1}} \Big(\mathbb{P}_{\mu\circ(\gamma^*,\tau),\pi}((a_t, h^t), m^t, \omega^t, \hat{a}^t) \times$$
$$\pi_{t+1}(h_{t+1}, m_{t+1}, \omega_{t+1}|a_t, h^t, m^t, \omega^t)\mu_{t+1}(\hat{a}_{t+1}|h^{t+1}, m^{t+1}, \hat{a}^t)\Big),$$

if $\bar{\tau}_{t+1}(a_{t+1}|h^t, \hat{a}^t) > 0$ since $\bar{\tau} \equiv (\gamma^*, \tau)$. (If $\bar{\tau}_{t+1}(a_{t+1}|h^t, \hat{a}^t) = 0$, then the above inequality is trivially satisfied.)

From the induction step, this is equivalent to

$$\Big(\sum_{m^t} \mathbb{P}_{\mu\circ(\gamma^*,\tau),\pi}((a_t, h^t), m^t, \omega^t, \hat{a}^t)\Big)p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t)\bar{\mu}_{t+1}(\hat{a}_{t+1}|h^{t+1}, \omega^{t+1}, \hat{a}^t) =$$
$$\sum_{m^{t+1}} \Big(\mathbb{P}_{\mu\circ(\gamma^*,\tau),\pi}((a_t, h^t), m^t, \omega^t, \hat{a}^t)\pi_{t+1}(h_{t+1}, m_{t+1}, \omega_{t+1}|a_t, h^t, m^t, \omega^t)\mu_{t+1}(\hat{a}_{t+1}|h^{t+1}, m^{t+1}, \hat{a}^t)\Big). \tag{1}$$

It remains to verify we can indeed construct $\bar{\mu}_{t+1}$ such that Equation (1) is satisfied for all strategies $\tau \in \mathcal{T}^\perp$, where

$$\mathcal{T}^\perp := \{\tau \in \mathcal{T} : [\tau_{i,t}(\cdot|h_i^t, m_i^t, h_i^t, m_i^t, \hat{a}_i^t) = \tau_{i,t}(\cdot|h_i^t, \overline{m}_i^t, h_i^t, \overline{m}_i^t, \hat{a}_i^t)\ \forall\ m_i^t, \overline{m}_i^t, h_i^t, \hat{a}_i^t, i, t]\}$$

is the set of action strategies, which does not depend on the additional signals at truthful histories.

We first note that since we consider histories in $H\Omega$, $p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t) > 0$. Second, we claim that there exists a strategy $\tau^\dagger \in \mathcal{T}^\perp$ such that

$$\sum_{m^t} \mathbb{P}_{\mu\circ(\gamma^*,\tau^\dagger),\pi}((a_t, h^t), m^t, \omega^t, \hat{a}^t) > 0,$$

whenever

$$\sum_{m^t} \underbrace{\pi_1(h_1, m_1, \omega_1)\mu_1(\hat{a}_1|h^1, m^1, \hat{a}^0) \times \cdots\cdots \times \pi_t(h_t, m_t, \omega_t|a_{t-1}, h^{t-1}, m^{t-1}, \omega^{t-1})\mu_t(\hat{a}_t|h^t, m^t, \hat{a}^{t-1})}_{\gamma(h^t, m^t, \omega^t, \hat{a}^t)} > 0.$$

To prove the claim, simply let $\tau^\dagger_{t'}(a_{t'}|h^{t'}, m^{t'}, h^{t'}, m^{t'}, \hat{a}^{t'}) = 1$ if and only if $h^{t'+1} = (a_{t'}, s_{t'+1})$ for all $t' \leq t$, that is, players play with probability one the actions specified in the history $(h^t, m^t, \omega^t, \hat{a}^t)$.

At all histories such that $\sum_{m^t} \gamma(h^t, m^t, \omega^t, \hat{a}^t) > 0$, we let

$$\bar{\mu}_{t+1}(\hat{a}_{t+1}|h^{t+1}, \omega^{t+1}, \hat{a}^t) =$$
$$\sum_{m^{t+1}} \frac{\mathbb{P}_{\mu\circ(\gamma^*,\tau^\dagger),\pi}((a_t, h^t), m^t, \omega^t, \hat{a}^t)\pi_{t+1}(h_{t+1}, m_{t+1}, \omega_{t+1}|a_t, h^t, m^t, \omega^t)}{\Big(\sum_{m^t} \mathbb{P}_{\mu\circ(\gamma^*,\tau^\dagger),\pi}((a_t, h^t), m^t, \omega^t, \hat{a}^t)\Big)p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t)}\mu_{t+1}(\hat{a}_{t+1}|h^{t+1}, m^{t+1}, \hat{a}^t) =$$
$$\sum_{m^{t+1}} \frac{\gamma(h^t, m^t, \omega^t, \hat{a}^t)}{\sum_{m^t} \gamma(h^t, m^t, \omega^t, \hat{a}^t)} \frac{\pi_{t+1}(h_{t+1}, m_{t+1}, \omega_{t+1}|a_t, h^t, m^t, \omega^t)}{p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t)}\mu_{t+1}(\hat{a}_{t+1}|h^{t+1}, m^{t+1}, \hat{a}^t) =$$
$$\sum_{m^{t+1}} \frac{\gamma(h^t, m^t, \omega^t, \hat{a}^t)}{\sum_{m^t} \gamma(h^t, m^t, \omega^t, \hat{a}^t)}\xi_{t+1}(m_{t+1}|h^{t+1}, m^t, \omega^{t+1})\mu_{t+1}(\hat{a}_{t+1}|h^{t+1}, m^{t+1}, \hat{a}^t),$$

for some ξ_{t+1}, where the last equality follows from admissibility. We can easily check that $\bar{\mu}_{t+1}(\hat{a}_{t+1}|h^{t+1}, \omega^{t+1}, \hat{a}^t) > 0$, and

$$\sum_{\hat{a}_{t+1}} \bar{\mu}_{t+1}(\hat{a}_{t+1}|h^{t+1}, \omega^{t+1}, \hat{a}^t) = 1,$$

and that its definition is independent of τ. It is then routine to check that for any other strategy $\tau \in \mathcal{T}^\perp$, Equation (1) is satisfied.

Finally, at all histories such that $\sum_{m^t} \gamma(h^t, m^t, \omega^t, \hat{a}^t) = 0$, we have that

$$\sum_{m^t} \mathbb{P}_{\mu\circ(\gamma^*,\tau),\pi}((a_t, h^t), m^t, \omega^t, \hat{a}^t) = 0,$$

for all $\tau \in \mathcal{T}^\perp$. At these histories, we choose an arbitrary $\bar{\mu}_{t+1}(\cdot|h^{t+1}, \omega^{t+1}, \hat{a}^t) \in \Delta(A_{t+1})$.

We have thus defined recommendation kernels $(\bar{\mu}_t : H^t\Omega^t \times A^{t-1} \to \Delta(A_t))_t$ such that

$$\sum_{\hat{a}} \mathbb{P}_{\bar{\mu}\circ\bar{\tau},p}(h, \omega, \hat{a}) = \sum_{m,\hat{a}} \mathbb{P}_{\mu\circ(\gamma^*,\tau),\pi}(h, m, \omega, \hat{a}),$$

for all (h, ω). Properties (i) and (ii) then follow immediately, which proves that

$$\bigcup_{\Gamma_\pi \text{ an admissible expansion of } \Gamma} \mathcal{CE}(\Gamma_\pi) \subseteq \mathcal{BCE}(\Gamma).$$

This completes the first part of the proof.

($\Leftarrow$). We now prove that

$$\mathcal{BCE}(\Gamma) \subseteq \bigcup_{\Gamma_\pi \text{ an admissible expansion of } \Gamma} \mathcal{BNE}(\Gamma_\pi).$$

Let $\bar{\mu}$ be a Bayes correlated equilibrium with distribution $\mathbb{P}_{\bar{\mu}\circ\bar{\tau}^*,p}$. We now construct an expansion Γ_π and a Bayes-Nash equilibrium σ^* of Γ_π, with the property that $\text{marg}_{H\Omega}\mathbb{P}_{\sigma^*,\pi} = \text{marg}_{H\Omega}\mathbb{P}_{\bar{\mu}\circ\bar{\tau}^*,p}$.

The expansion is as follows. Let $M_{i,t} = A_{i,t}$ for all (i,t),

$$\pi_1(h_1, m_1, \omega_1) = p_1(h_1, \omega_1)\bar{\mu}_1(\hat{a}_1|h_1, \omega_1),$$

with $m_1 = \hat{a}_1$, for all (h_1, m_1, ω_1), and

$$\pi_{t+1}(h_{t+1}, m_{t+1}, \omega_{t+1}|a_t, h^t, m^t, \omega^t) = p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t)\bar{\mu}_{t+1}(\hat{a}_{t+1}|h^{t+1}, \omega^{t+1}, \hat{a}^t),$$

with $(m^t, m_{t+1}) = (\hat{a}^t, \hat{a}_{t+1})$, for all $(a_t, h^t, m^t, \omega^t, h_{t+1}, m_{t+1}, \omega_{t+1})$. Clearly, the expansion is admissible.

By construction, any strategy $\bar{\tau}_t : H^t \times A^t \to \Delta(A_t)$ of $\mathcal{M}(\Gamma)$ is equivalent to a strategy $\sigma_t : H^t \times M^t \to \Delta(A_t)$ of Γ_π, i.e., $\sigma_t(a_t|h^t, m^t) := \times_i \sigma_{i,t}(a_{i,t}|h_i^t, m_i^t) = \times_i \bar{\tau}_{i,t}(a_{i,t}|h_i^t, \hat{a}_i^t)$ with $m^t = \hat{a}^t$, with the property that $\mathbb{P}_{\sigma,\pi}(h^t, m^t, \omega^t) = \mathbb{P}_{\bar{\mu}\circ\bar{\tau},p}(h^t, \hat{a}^t, \omega^t)$ when $m^t = \hat{a}^t$, for all (h^t, m^t, ω^t), for all t.

To see this last point, note that the definition of π_1 is clearly equivalent to $\mathbb{P}_{\sigma,\pi}(h_1, m_1, \omega_1) = \mathbb{P}_{\bar{\mu}\circ\bar{\tau},p}(h_1, \omega_1, \hat{a}_1)$ with $m_1 = \hat{a}_1$, for all (h_1, m_1, ω_1). By induction, assume that $\mathbb{P}_{\sigma,\pi}(h^t, m^t, \omega^t) = \mathbb{P}_{\bar{\mu}\circ\tau,p}(h^t, \omega^t, \hat{a}^t)$ with $m^t = \hat{a}^t$, for all (h^t, m^t, ω^t). We now compute the probability of $(h^{t+1}, m^{t+1}, \omega^{t+1})$. We have that

$$\begin{aligned}
\mathbb{P}_{\sigma,\pi}(h^{t+1}, m^{t+1}, \omega^{t+1}) &= \mathbb{P}_{\sigma,\pi}(h_{t+1}, m_{t+1}, \omega_{t+1}|h^t, m^t, \omega^t)\mathbb{P}_{\sigma,\pi}(h^t, m^t, \omega^t) \\
&= \pi_{t+1}(h_{t+1}, m_{t+1}, \omega_{t+1}|a_t, h^t, m^t, \omega^t)\sigma_t(a_t|h^t, m^t)\mathbb{P}_{\sigma,\pi}(h^t, m^t, \omega^t) \\
&= p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t)\bar{\mu}_{t+1}(\hat{a}_{t+1}|h^{t+1}, \omega^{t+1}, \hat{a}^t)\bar{\tau}_t(a_t|h^t, \hat{a}^t)\mathbb{P}_{\bar{\mu}\circ\bar{\tau},p}(h^t, \hat{a}^t, \omega^t) \\
&= \mathbb{P}_{\bar{\mu}\circ\bar{\tau},p}(h^{t+1}, \omega^{t+1}, \hat{a}^{t+1}),
\end{aligned}$$

with $\hat{a}^{t+1} = m^{t+1}$. Finally, since $\bar{\mu}$ is a Bayes correlated equilibrium of $\mathcal{M}(\Gamma)$, the strategy $\sigma^* \equiv \bar{\tau}^*$ is a Bayes-Nash equilibrium of Γ_π and, thus,

$$\mathcal{BCE}(\Gamma) \subseteq \bigcup_{\Gamma_\pi \text{ an admissible expansion of } \Gamma} \mathcal{BNE}(\Gamma_\pi) \subseteq \bigcup_{\Gamma_\pi \text{ an admissible expansion of } \Gamma} \mathcal{CE}(\Gamma_\pi).$$

This completes the proof.

C A model with exogenous evolutions of the states.

We discuss a particular case of our general model, where the information about states and actions is decoupled. Assume that at each stage, a player receives a pair of base signals $(s_{i,t}, \mathring{s}_{i,t})$. The signal $s_{i,t}$ is informative about the past and current states and past base signals about the states, while the second signal $\mathring{s}_{i,t}$ is informative about the past actions, past base signals and past states. Formally, we have two processes $(p_t, \mathring{p}_t)$, with $p_{t+1}(s_{t+1}, \omega_{t+1}|s^t, \omega^t)$ the probability of (s_{t+1}, ω_{t+1}) given (s^t, ω^t), and $\mathring{p}_{t+1}(\mathring{s}_{t+1}|a^t, \mathring{s}^t, s^t, \omega^t)$ the probability of $\mathring{s}_{t+1}$, given $(a^t, \mathring{s}^t, s^t, \omega^t)$. It is worth observing that the probability of (s, ω) is independent of any profile of actions a and given by $p(s, \omega)$, with

$$p(s, \omega) = p_1(s_1, \omega_1) \cdot \prod_{t \in T} p_{t+1}(s_{t+1}, \omega_{t+1}|s^t, \omega^t).$$

Crucially, the players do not control the evolution of the state. As before, we let $p^a(s, \mathring{s}, \omega)$ be the probability of $(s, \mathring{s}, \omega)$, given the action profile a.

An expansion consists of message spaces $(M_{i,t})$ and a process $(\pi_t)_t$, with

$$\pi_{t+1}(s_{t+1}, \mathring{s}_{t+1}, m_{t+1}, \omega_{t+1}|a^t, s^t, \mathring{s}^t, m^t, \omega^t)$$

the probability of $(s_{t+1}, \mathring{s}_{t+1}, m_{t+1}, \omega_{t+1})$ given $(a^t, m^t, \mathring{s}^t, s^t, \omega^t)$. As before, we let $\pi^a(s, \mathring{s}, m, \omega)$ be the probability of $(s, \mathring{s}, m, \omega)$, given the action profile a. Conditional on (s, ω), we define

$$\mathring{\pi}_{t+1}(\mathring{s}_{t+1}, m_{t+1}|a, \mathring{s}^t, m^t, s, \omega) := \frac{\sum_{(\mathring{\mathbf{s}}, \mathbf{m}):(\mathring{\mathbf{s}}^{t+1}, \mathbf{m}^{t+1}) = ((m^t, m_{t+1}), (\mathring{s}^t, \mathring{s}_{t+1}))} \pi^a(s, \mathring{\mathbf{s}}, \mathbf{m}, \omega)}{\sum_{(\mathring{\mathbf{s}}, \mathbf{m}):(\mathring{\mathbf{s}}^t, \mathbf{m}^t) = (m^t, \mathring{s}^t)} \pi^a(s, \mathring{\mathbf{s}}, \mathbf{m}, \omega)},$$

for all $(a, m^t, \mathring{s}^t, \mathring{s}_{t+1}, m_{t+1})$ with a positive denominator, and arbitrary otherwise. It is worth noting that $\mathring{\pi}_{t+1}$ only depends on a^t (and not $(a_{t+1}, \ldots, a_T)$) when the marginal of π^a over $S \times \Omega$ is p, as we shall assume soon.

In that class of multi-stage games, the expansion is said to be weakly admissible if it satisfies the following two properties:

(i) The marginal of π^a over $S \times \Omega$ is p, i.e., $\sum_{\mathring{s}, m} \pi^a(s, \mathring{s}, m, \omega) = p(s, \omega)$ for all a.

(ii) There exist probability kernels $(\xi)_t$ such that

$$\xi_{t+1}(m_{t+1}|a^t, \mathring{s}^{t+1}, m^t, s, \omega)\mathring{p}_{t+1}(\mathring{s}_{t+1}|a^t, \mathring{s}^t, s^t, \omega^t) = \mathring{\pi}_{t+1}(m_{t+1}, \mathring{s}_{t+1}|a^t, \mathring{s}^t, m^t, s, \omega),$$

for all $(a^t, \mathring{s}^{t+1}, m^t, s, \omega, m_{t+1})$.

Examples 1 and 2 are instances of that class of multi-stage games and satisfy weak admissibility. It is also worth pointing out that if, in addition, we assume perfect observation of past actions, then weak admissibility is equivalent to consistency, i.e., $\sum_m \pi^a(s, \mathring{s}, m, \omega) = p^a(s, \mathring{s}, \omega)$.[23]

It is then immediate to replicate our analysis with recommendation kernels conditioning on (s, ω), i.e., $\overline{\mu}_{t+1}(\hat{a}_{t+1}|a^t, \hat{a}^t, \mathring{s}^t, s, \omega)$ is the probability of recommending $\hat{a}_{t+1}$ at stage $t+1$, when the profile of states and base signals is (s, ω) and the profile of past actions, past recommendations and past base signals about actions is $(a^t, \hat{a}^t, \mathring{s}^t)$. In words, the omniscient mediator not only knows all realizations of past actions, past and current base signals and states, but also all future states and base signals about states.[24]

D Proof of Theorem 2

Preliminary observations on conditional probability systems. We start with some preliminary observations on conditional probability systems. Let $\mathcal{X}$ be a finite set and β a CPS on $\mathcal{X}$. The CPS

[23] We have perfect observation of past actions if for all a_t, there exists $\mathring{s}_{t+1}(a_t)$ such that $\mathring{p}_{t+1}(\mathring{s}_{t+1}(a_t)|a^t, \mathring{s}^t, s^t, \omega^t) = 1$. Consistency implies that $\pi^a(s, \mathring{s}, m, \omega) = \pi^a(m|s, \mathring{s}, \omega)p^a(s, \mathring{s}, \omega)$ and, therefore,

$$\sum_{m_{t+1}} \mathring{\pi}_{t+1}(m_{t+1}, \mathring{s}_{t+1}|a^t, m^t, \mathring{s}^t, s, \omega) = \frac{\sum_{\mathring{\mathbf{s}}:\mathring{\mathbf{s}}^{t+1}=(\mathring{s}^t, \mathring{s}_{t+1})} p^a(\mathring{\mathbf{s}}|s, \omega)\Big(\sum_{\mathbf{m}:\mathbf{m}^t=m^t} \pi^a(\mathbf{m}|s, \mathring{s}, \omega)\Big)}{\sum_{\mathring{\mathbf{s}}:\mathring{\mathbf{s}}^t=\mathring{s}^t} p^a(\mathring{\mathbf{s}}|s, \omega)\Big(\sum_{\mathbf{m}:\mathbf{m}^t=m^t} \pi^a(\mathbf{m}|s, \mathring{s}, \omega)\Big)}.$$

It is thus equal to one if $\mathring{s}_{t+1} = \mathring{s}_{t+1}(a_t)$, i.e., weak admissibility holds. The converse is clearly true.

[24] It is worth observing that weak admissibility is actually equivalent to admissibility in a modified base game, where $\Omega_1^* = S \times \Omega$, $S_t^* = S_t \times \mathring{S}_t$, Ω_t^* is a singleton for all $t > 1$, $p_1^*((s_1, \mathring{s}_1), (s, \omega)) = \mathring{p}_1(\mathring{s}_1)p(s, \omega)$ and $p_{t+1}^*(s_{t+1}, \mathring{s}_{t+1}|a^t, (s^t, \mathring{s}^t), (s, \omega)) = \mathring{p}_{t+1}(\mathring{s}_{t+1}|a^t, \mathring{s}^t, s^t, \omega^t)$. Note that this implies that $p_{t+1}^*(\tilde{s}_{t+1}, \mathring{s}_{t+1}|a^t, (s^t, \mathring{s}^t), (s, \omega)) = 0$ if $\tilde{s}_{t+1} \neq s_{t+1}$. In words, in the modified base game, the states and base signals (s, ω) are drawn at the beginning of the game, and players gradually learn about them.

β satisfies the following two properties: for any $X \neq \emptyset$,

(i) $\beta(X'|X) = 1$ if $X \subseteq X'$.

(ii) $\beta(X'|X) = 0$ if $X' \cap X = \emptyset$.

To prove (i), simply observe that $1 \geq \beta(X'|X) = \beta(X|X) + \beta(X' \setminus X|X) \geq 1$. To prove (ii), note that $1 = \beta(\mathcal{X}|X) = \beta(X'|X) + \beta(\mathcal{X} \setminus X'|X) = \beta(X'|X) + 1$ (since $X \subseteq \mathcal{X} \setminus X'$).

Next, let $\mathcal{Y}$ be another non-empty finite set and $\mathcal{XY}$ a non-empty finite subset of $\mathcal{X} \times \mathcal{Y}$. Let β^* be a CPS on $\mathcal{XY}$. For every $X' \subseteq \mathcal{X}$ and $\emptyset \neq X \subseteq \mathcal{X}$, define

$$\beta(X'|X) := \beta^*(\{(x,y) \in \mathcal{XY} : x \in X'\}|\{(x,y) \in \mathcal{XY} : x \in X\}).$$

We argue that β is a CPS on $\mathcal{X}$ if the set $\{(x,y) \in \mathcal{XY} : x \in X\}$ is non-empty for every non-empty $X \neq \emptyset$.

We clearly have that $\beta(X|X) = 1$ and from (i) above that $\beta(\mathcal{X}|X) = 1$. Consider any triple (X, X', X'') such that $X \subseteq X' \subseteq X''$ and $X'' \neq \emptyset$. If $X \cap X' = \emptyset$, we have that

$$\begin{aligned} \beta(X \cup X'|X'') = \beta^*(\{(x,y) \in \mathcal{XY} : x \in X \cup X'\}|\{(x,y) \in \mathcal{XY} : x \in X''\}) = \\ \beta^*(\{(x,y) \in \mathcal{XY} : x \in X\} \cup \{(x,y) \in \mathcal{XY} : x \in X'\}|\{(x,y) \in \mathcal{XY} : x \in X''\}) = \\ \beta(X|X'') + \beta(X'|X''). \end{aligned}$$

Finally, we have that

$$\begin{aligned} \beta(X|X'') = \beta^*(\{(x,y) \in \mathcal{XY} : x \in X\}|\{(x,y) \in \mathcal{XY} : x \in X''\}) = \\ \beta^*(\{(x,y) \in \mathcal{XY} : x \in X\}|\{(x,y) \in \mathcal{XY} : x \in X'\}) \times \\ \beta^*(\{(x,y) \in \mathcal{XY} : x \in X'\}|\{(x,y) \in \mathcal{XY} : x \in X''\}) = \\ \beta(X|X')\beta(X'|X''). \end{aligned}$$

Therefore, β is a well-defined CPS. These observations will prove to be useful in the proof below of Theorem 2

$(\Rightarrow)$. We first show that $\bigcup_{\Gamma_\pi \text{ an admissible expansion of } \Gamma} \mathcal{SCE}(\Gamma_\pi) \subseteq \mathcal{SBCE}(\Gamma)$. Throughout, we fix an admissible expansion Γ_π of Γ. Recall that since the expansion is admissible, there exist kernels

$(\xi_t)_t$ such that

$$\pi_{t+1}(h_{t+1}, m_{t+1}, \omega_{t+1} | a_t, h^t, m^t, \omega^t) = \xi_{t+1}(m_{t+1} | h^{t+1}, m^t, \omega^{t+1}) p_{t+1}(h_{t+1}, \omega_{t+1} | a_t, h^t, \omega^t),$$

for all $(h^{t+1}, m^{t+1}, \omega^{t+1})$, for all t.

Fix a distribution $\boldsymbol{\mu}^d \in \mathcal{SCE}(\Gamma_\pi)$. From the revelation principle of Sugaya and Wolitzky (2017, Proposition 3), there exist a Myersonian extension $\mathcal{M}(\Gamma_\pi)$ of Γ_π, where at each stage t, player i receives the signal $(h_{i,t}, m_{i,t}) \in H_{i,t} \times M_{i,t}$, reports the message $x_{i,t} \in X_{i,t}$ to the mediator, receives the message $y_{i,t} \in Y_{i,t}$ from the mediator, and takes the action $a_{i,t} \in A_{i,t}$, and a CPPBE $((\tau^*, \sigma^*), \mu^*, \beta^*)$ of $\mathcal{M}(\Gamma_\pi)$, such that the marginal of $\mathbb{P}_{\mu^* \circ (\tau^*, \sigma^*), \pi}$ over $H\Omega$ is $\boldsymbol{\mu}^d$.[25] (Here, (τ^*, σ^*) is the profile of strategies of the players, μ^* is the mediator strategy, and β^* is the conditional probability system.)

To prove that $\boldsymbol{\mu}^d \in \mathcal{SBCE}(\Gamma)$, we construct a "fictitious" mediated multi-stage game $\mathcal{M}(\Gamma^*)$, where (i) each stage t has two sub-stages, and (ii) the mediator receives the signal (h_t, ω_t) at the beginning of each stage. The mediated game $\mathcal{M}(\Gamma^*)$ is as follows: at each stage $t \in T$,

First sub-stage:

- player i receives the signal $h_{i,t}$; the mediator receives the signal (h_t, ω_t);
- player i makes the report $\underline{x}_{i,t} \in \{\underline{x}_{i,t}\}$ to the mediator;
- the mediator sends the message $m_{i,t} \in M_{i,t}$ to player i;
- player i takes the action $\underline{a}_{i,t} \in \{\underline{a}_{i,t}\}$;

Second sub-stage:

- player i reports $x_{i,t} \in X_{i,t}$ to the mediator;
- the mediator sends the message $y_{i,t} \in Y_{i,t}$ to player i;

[25]The converse also holds, that is, for every Myersonian extension $\mathcal{M}(\Gamma_\pi)$ of Γ_π and every CPPBE of $\mathcal{M}(\Gamma_\pi)$, the marginal of $\mathbb{P}_{\mu^* \circ (\tau^*, \sigma^*), \pi}$ over $H\Omega$ is in $\mathcal{SCE}(\Gamma_\pi)$.

- player i takes the action $a_{i,t} \in A_{i,t}$.

Thus, the mediator is perfectly informed of the realization of (h_t, ω_t), while player i is informed of $h_{i,t}$. At the history $(a_{t-1}, h^{t-1}, m^{t-1}, \omega^{t-1}, x^{t-1}, y^{t-1}, \underline{x}^{t-1}, \underline{a}^{t-1})$, the probability of (h_t, ω_t) is

$$p_t(h_t, \omega_t | a_{t-1}, h^{t-1}, m^{t-1}, \omega^{t-1}).$$

Finally, the payoffs of $\mathcal{M}(\Gamma^*)$ are as in Γ, that is, if the profile of actions and states is $((\underline{a}, a), \omega)$, the payoff to player i is $u_i(a, \omega)$. This completes the description of the "fictitious" multi-stage game.

Throughout, we slightly abuse notations and do not refer to the trivial sequence of reports $\underline{x}$ and sequences of actions $\underline{a}$ to describe the histories. E.g., we write $(h_i^t, m_i^t, x_i^t, y_i^t)$ for $(h_i^t, m_i^t, x_i^t, y_i^t, \underline{x}_i^t, \underline{a}_i^t,)$. This is inconsequential since players have a single possible report and possible action at each first sub-stage of a stage.

A non-trivial strategy for player i at stage t is a pair $(\tau_{i,t}, \sigma_{i,t})$, with $\tau_{i,t}(x_{i,t} | h_i^t, m_i^t, x_i^{t-1}, y_i^{t-1})$ the probability of reporting $x_{i,t}$ at the private history $(h_i^t, m_i^t, x_i^{t-1}, y_i^{t-1})$ and $\sigma_{i,t}(a_{i,t} | h_i^t, m_i^t, x_i^t, y_i^t)$ the probability of playing $a_{i,t}$ at the private history $(h_i^t, m_i^t, x_i^t, y_i^t)$. (Formally, we would also need to specify that player i reports $\underline{x}_{i,t}$ and takes the action $\underline{a}_{i,t}$ at the first sub-stage of each stage, but again this is irrelevant.) Thus, the strategies in $\mathcal{M}(\Gamma^*)$ are essentially the same as in $\mathcal{M}(\Gamma_\pi)$.

By construction, we have that $((\tau^*, \sigma^*), (\xi, \mu^*), \beta^*)$ is a CPPBE of $\mathcal{M}(\Gamma^*)$, where again we slightly abuse the notations to define the conditional probability system, that is, $\beta^*(h, m, \omega, x, y, \underline{x}, \underline{a}) := \beta^*(h, m, \omega, x, y)$ for all $(h, m, \omega, x, y, \underline{x}, \underline{a})$.

To conclude the proof, we invoke again the revelation principle of Sugaya and Wolitzky (2017). Note that it is possible to do so, because Sugaya and Wolitzky allow for the mediator to receive private signals. Therefore, $\boldsymbol{\mu}^d$ is also a sequential communication equilibrium distribution of the canonical mechanism:

First sub-stage:

- player i receives the signal $h_{i,t}$; the mediator receives the signal (h_t, ω_t);

- player i reports $\tilde{h}_{i,t}$ to the mediator;

- the mediator recommends the trivial action $\underline{a}_{i,t}$ to player i;

- player i takes the trivial action $\underline{a}_{i,t}$;

  Second sub-stage:

- player i reports $\hat{h}_{i,t}$ to the mediator;

- the mediator recommends $\hat{a}_{i,t}$ to player i; the mediator's recommendation is conditional on $(h^t, \omega^t, \hat{a}^{t-1})$ as well as the players' past and current reports.

- player i takes the action $a_{i,t}$.

By the revelation principle, the players have an incentive to be truthful and obedient at all their private histories consistent with the mediation ranges at which they have been truthful in the past. Clearly, at all truthful histories consistent with the mediation ranges, the players must have an incentive to obey the recommendation of the informed mediator; that is, we have a sequential Bayes correlated equilibrium. Hence, $\boldsymbol{\mu^d} \in \mathcal{SBCE}(\Gamma)$, as required.

($\Leftarrow$). We show that $\mathcal{SBCE}(\Gamma) \subseteq \bigcup_{\Gamma_\pi \text{ an admissible expansion of } \Gamma} \mathcal{CPPBE}(\Gamma_\pi)$.

Let $(\overline{\mu}, \overline{R}, \overline{\beta})$ be a sequential Bayes correlated equilibrium, with induced probability $\overline{\mathbb{P}}_{\overline{\mu}\circ\overline{\tau}^*,p}$. We now construct an expansion Γ_π and a conditional probability perfect Bayesian equilibrium $(\sigma^*, \hat{\beta})$ of Γ_π.

The expansion is as follows. For all i, for all t, let $M_{i,t} = A_{i,t}$,

$$\pi_1(h_1, m_1, \omega_1) = p_1(h_1, \omega_1)\overline{\mu}_1(\hat{a}_1|h^1, \omega^1)$$

with $m_1 = \hat{a}_1$, for all $(h_1, m_1, \omega_1, \hat{a}_1)$, and

$$\pi_{t+1}(h_{t+1}, m_{t+1}, \omega_{t+1}|a_t, h^t, m^t, \omega^t) = p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t)\overline{\mu}_{t+1}(\hat{a}_{t+1}|h^{t+1}, \omega^{t+1}, \hat{a}^t),$$

with $m^{t+1} = \hat{a}^{t+1}$, for all $(a_t, h^t, m^t, \omega^t, \hat{a}^t, h_{t+1}, m_{t+1}, \omega_{t+1}, \hat{a}_{t+1})$ such that $(h^t, \omega^t, \hat{a}^t) \in H^t\Omega^t\overline{R}^t$. The constructed expansion is clearly admissible.

By construction, note also that any strategy $\bar{\tau}_t : H^t \times A^t \to \Delta(A_t)$ of $\mathcal{M}(\Gamma)$ is equivalent to a strategy $\sigma_t : H^t \times M^t \to \Delta(A_t)$ of Γ_π, i.e., $\sigma_t(a_t|h^t, m^t) := \bar{\tau}_t(a_t|h^t, \hat{a}^t)$ with $m^t = \hat{a}^t$. We write $\sigma \equiv \bar{\tau}$ for such an equivalence. We associate σ^* with τ^*. Again, this is a product of behavioral strategies.

By construction, for any σ, we have that

$$\begin{aligned}
\overline{\mathbb{P}}_{\sigma,\pi}(h_{t+1}, m_{t+1}, \omega_{t+1}|h^t, m^t, \omega^t) &= \\
\sum_{a_t} \pi_{t+1}(h_{t+1}, m_{t+1}, \omega_{t+1}|a_t, h^t, m^t, \omega^t)\sigma_t(a_t|h^t, m^t) &= \\
\sum_{a_t} p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t)\bar{\mu}_{t+1}(\hat{a}_{t+1}|h^{t+1}, \omega^{t+1}, \hat{a}^t)\bar{\tau}_t(a_t|h^t, \hat{a}^t) &= \\
\overline{\mathbb{P}}_{\bar{\mu}\circ\bar{\tau},p}(h_{t+1}, \omega_{t+1}, \hat{a}_{t+1}|h^t, \omega^t, \hat{a}^t),
\end{aligned}$$

with $\hat{a}^{t+1} = m^{t+1}$ and $\bar{\tau} \equiv \sigma$. This implies then that, for all t, all $(\boldsymbol{h}, \boldsymbol{m}, \boldsymbol{\omega}, \boldsymbol{\hat{a}})$ and $(h^t, \omega^t, \hat{a}^t) \in H^t\Omega^t\overline{R}^t$:

$$\overline{\mathbb{P}}_{\sigma,\pi}(\boldsymbol{h}, \boldsymbol{m}, \boldsymbol{\omega}|h^t, m^t, \omega^t) = \overline{\mathbb{P}}_{\bar{\mu}\circ\bar{\tau},p}(\boldsymbol{h}, \boldsymbol{\omega}, \boldsymbol{\hat{a}}|h^t, \omega^t, \hat{a}^t),$$

with $\boldsymbol{m} = \boldsymbol{\hat{a}}$ and $m^t = \hat{a}^t$, whenever $\bar{\tau} \equiv \sigma$.

Finally, notice that $(h, m, \omega) \in HM\Omega$ if and only if

$$p_1(h_1, \omega_1)\bar{\mu}_1(\hat{a}_1|h^1, \omega^1) \cdot \prod_{t=1}^{T-1} p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t)\bar{\mu}_{t+1}(\hat{a}_{t+1}|h^{t+1}, \omega^{t+1}, \hat{a}^t) > 0,$$

with $\hat{a} = m$, and A_{T+1} a singleton.

Claim: If $(h, \hat{a}, \omega) \in HM\Omega$, then $(h, \omega, \hat{a}) \in H\Omega\overline{R}$.

Proof of the claim. From the consistency of the CPS, we have that

$$\sum_{\hat{a}_{t+1}\in\times_i\overline{R}_{i,t+1}(h_i^{t+1},\hat{a}_i^t)} \bar{\mu}_{t+1}(\hat{a}_{t+1}|h^{t+1}, \omega^{t+1}, \hat{a}^t) = \sum_{\hat{a}_{t+1}\in\times_i\overline{R}_{i,t+1}(h_i^{t+1},\hat{a}_i^t)} \bar{\beta}(\hat{a}_{t+1}|h^{t+1}, \omega^{t+1}, \hat{a}^t).$$

Moreover,

$$\begin{aligned}
&\bar{\beta}(\hat{a}_{t+1}|h^{t+1}, \omega^{t+1}, \hat{a}^t) = \\
&\bar{\beta}(\{(\boldsymbol{h}, \boldsymbol{\omega}, \boldsymbol{\hat{a}}) \in H\Omega\overline{R} : (\boldsymbol{h^{t+1}}, \boldsymbol{\omega^{t+1}}, \boldsymbol{\hat{a}^{t+1}}) = (h^{t+1}, \omega^{t+1}, \hat{a}^{t+1})\}| \\
&\qquad \{(\boldsymbol{h}, \boldsymbol{\omega}, \boldsymbol{\hat{a}}) \in H\Omega\overline{R} : (\boldsymbol{h^{t+1}}, \boldsymbol{\omega^{t+1}}, \boldsymbol{\hat{a}^{t}}) = (h^{t+1}, \omega^{t+1}, \hat{a}^{t})\}).
\end{aligned}$$

Therefore,

$$\sum_{\hat{a}_{t+1} \in \times \overline{R}_{i,t+1}(h_i^{t+1}, \hat{a}_i^t)} \bar{\beta}(\hat{a}_{t+1} | h^{t+1}, \omega^{t+1}, \hat{a}^t) = 1,$$

since

$$\bigcup_{\hat{a}_{t+1} \in \times \overline{R}_{i,t+1}(h_i^{t+1}, \hat{a}_i^t)} \{(\boldsymbol{h}, \boldsymbol{\omega}, \hat{\boldsymbol{a}}) \in H\Omega\overline{R} : (\boldsymbol{h^{t+1}}, \boldsymbol{\omega^{t+1}}, \hat{\boldsymbol{a}}^{\boldsymbol{t+1}}) = (h^{t+1}, \omega^{t+1}, \hat{a}^{t+1})\} =$$

$$\{(\boldsymbol{h}, \boldsymbol{\omega}, \hat{\boldsymbol{a}}) \in H\Omega\overline{R} : (\boldsymbol{h^{t+1}}, \boldsymbol{\omega^{t+1}}, \hat{\boldsymbol{a}}^{\boldsymbol{t}}) = (h^{t+1}, \omega^{t+1}, \hat{a}^t)\}.$$

This proves the claim. ■

From the claim above, $(h, m, \omega) \in HM\Omega$ implies that $(h, \omega, \hat{a}) \in H\Omega\overline{R}$, with $m = \hat{a}$. We denote that subset $H\Omega\overline{R}^*$. Note that it is non-empty. We then identify $\hat{\beta}$ on $HM\Omega$ with $\overline{\beta}$ on the subset $HM\overline{R}^*$, that is,

$$\hat{\beta}(X|Y) := \overline{\beta}(\{(h, \omega, \hat{a}) \in HM\overline{R}^* : (h, \hat{a}, \omega) \in X\} | \{(h, \omega, \hat{a}) \in HM\overline{R}^* : (h, \hat{a}, \omega) \in Y\},$$

for all $X \subseteq HM\Omega$ and $\emptyset \neq Y \subseteq HM\Omega$. From the preliminary observations on conditional probability systems, $\hat{\beta}$ is a well-defined conditional probability system.

From these definitions and the consistency of $\overline{\beta}$ with $(\overline{\tau}^*, \overline{\mu}, p)$, we immediately have consistency of the CPS $\hat{\beta}$ with (σ^*, π).

The above properties of $\overline{\mathbb{P}}_{\sigma,\pi}$ and $\hat{\beta}$ imply that for all $(h_i^t, \hat{a}_i^t) \in H_i^t \overline{R}_i^t$:

$$\sum_{\boldsymbol{h}^t, \boldsymbol{m}^t, \boldsymbol{\omega}^t} U_i((\sigma_i, \sigma_{-i}^*) | \boldsymbol{h}^t, \boldsymbol{m}^t, \boldsymbol{\omega}^t) \hat{\beta}(\boldsymbol{h}^t, \boldsymbol{m}^t, \boldsymbol{\omega}^t | h_i^t, m_i^t) = \sum_{\boldsymbol{h}^t, \boldsymbol{\omega}^t, \hat{\boldsymbol{a}}^t} U_i(\overline{\mu} \circ (\overline{\tau}_i, \overline{\tau}_{-i}^*) | \boldsymbol{h}^t, \boldsymbol{\omega}^t, \hat{\boldsymbol{a}}^t) \overline{\beta}(\boldsymbol{h}^t, \boldsymbol{\omega}^t, \hat{\boldsymbol{a}}^t | h_i^t, \hat{a}_i^t),$$

with $\boldsymbol{m} = \hat{\boldsymbol{a}}$ and $m^t = \hat{a}^t$. It follows that $(\sigma^*, \hat{\beta})$ is a conditional probability perfect Bayesian equilibrium of Γ_π.

To complete the proof, it suffices to recall that $\mathcal{CPPBE}(\Gamma_\pi) \subseteq \mathcal{SCE}(\Gamma_\pi)$ for all expansions Γ_π of Γ.

E A variation on Theorem 2

We say that $(\overline{\mu}, \overline{R}, \overline{\beta})$ is a strong sequential Bayes correlated equilibrium if it is a sequential Bayes correlated equilibrium and there exists a sequence $(\overline{\tau}^k)_k$ of completely mixed strategies converging to $\overline{\tau}^*$ with

$$\overline{\beta} = \lim_k \overline{\mathbb{P}}_{\bar{\mu}\circ\bar{\tau}^k, p}.$$

We let $\mathcal{SBCE}^*(\Gamma)$ be the distribution over $H\Omega$ induced by the strong sequential Bayes correlated equilibria of Γ. Note that $\mathcal{SBCE}^*(\Gamma) \subseteq \mathcal{SBCE}(\Gamma)$. We argue that

$$\mathcal{SBCE}^*(\Gamma) \subseteq \bigcup_{\Gamma_\pi \text{ an admissible expansion of } \Gamma} \mathcal{SE}(\Gamma_\pi) \subseteq \mathcal{SBCE}(\Gamma).$$

To prove the above claim, we revisit the only if part of the proof of Theorem 2. Define the sequence $(\sigma^k)_k$ with $\sigma^k \equiv \tau^k$ for all k. By construction of the expansion and the equivalence of σ^k with τ^k, we have that

$$\overline{\mathbb{P}}_{\sigma^k,\pi}(\boldsymbol{h}, \boldsymbol{m}, \boldsymbol{\omega} | h^t, m^t, \omega^t) = \overline{\mathbb{P}}_{\bar{\mu}\circ\bar{\tau}^k, p}(\boldsymbol{h}, \boldsymbol{\omega}, \hat{\boldsymbol{a}} | h^t, \omega^t, \hat{a}^t),$$

with $\boldsymbol{m} = \hat{\boldsymbol{a}}$ and $m^t = \hat{a}^t$ for all $(h^t, \omega^t, \hat{a}^t)$, for all t, and, therefore,

$$\beta = \lim_k \overline{\mathbb{P}}_{\sigma^k,\pi}.$$

Since $\sigma_t^k = \times \sigma_{i,t}^n$ for all (t, k), $(\sigma^k)_k$ is a sequence of completely mixed strategies. This completes the proof.

References

[1] Bergemann, Dirk and Stephen Morris, "Information Design: A Unified Perspective," 2018, forthcoming in Journal of Economic Literature.

[2] Bergemann, Dirk and Stephen Morris, "Bayes Correlated Equilibrium and The Comparison of Information Structures in Games," 2016, Theoretical Economics, 11, pp. 487-522.

[3] Bergemann, Dirk, Benjamin Brooks and Stephen Morris, "The Limits of Price Discrimination," 2015, American Economic Review, 105, pp. 921-957.

[4] Calzolari, Giacomo and Alessandro Pavan, "Monopoly with Resale," 2006, Rand Journal of Economics, 37, pp. 362-375.

[5] Condorelli, Daniele and Balász Szentes, "Information Design in the Hold-up Problem," 2018, mimeo.

[6] Doval, Laura and Jeffrey C., Ely, "Sequential Information Design," 2017, mimeo.

[7] Ely, Jeffrey C., "Beeps." 2017, American Economic Review, 107, pp. 31-53.

[8] Dworczak, Piotr, "Mechanism Design with Aftermarkets: Cutoff Mechanisms," 2017, mimeo.

[9] Kamenica, Emir and Matthew Gentzkow, "Bayesian Persuasion," 2011, American Economic Review, 101, pp. 2590-2615.

[10] Forges, Françoise, "An Approach to Communication Equilibria," 1986, Econometrica, 54, pp. 1375-1385.

[11] Forges, Françoise, "Five Legitimate Definitions of Correlated Equilibrium in Games with Incomplete Information," 1993, Theory and Decision, 35, pp. 277-310.

[12] Forges, Françoise, "Correlated Equilibrium in Games with Incomplete Information Revisited, 2006, Theory and Decision, 61, pp. 329-344.

[13] Giovannoni, Francesco, and Miltiadis Makris, "Reputational Bidding," 2014, International Economic Review, 55, pp. 693-710.

[14] Myerson, Roger, "Multistage Games with Communication," 1986, Econometrica, 54, pp. 323-358.

[15] Renault, Jérôme., Elon Solan, and Nicolas Vieille, "Optimal Dynamic Information Provision," 2017, Games and Economic Behavior, 104, pp. 329-349.

[16] Roesler, Anne-Katrin, and Balász Szentes, "Buyer-optimal Learning and Monopoly Pricing," 2017, American Economic Review, 107, pp. 2072-2080.

[17] Sugaya, Takuo, and Alexander Wolitzky, "Revelation Principles in Multistage Games," 2017, mimeo.

[18] von Stengel Bernhard and Françoise Forges, "Extensive form Correlated Equilibrium: Definition and Computational Complexity," 2008, Mathematics of Operations Research, 33, pp. 1002-1022.

School of Economics and Finance

This working paper has been produced by the School of Economics and Finance at Queen Mary University of London

School of Economics and Finance
Queen Mary University of London
Mile End Road
London E1 4NS
Tel: +44 (0)20 7882 7356
Fax: +44 (0)20 8983 3580
Web: www.econ.qmul.ac.uk/research/workingpapers/