

INFORMATION DESIGN IN MULTI-STAGE GAMES

MILTADIS MAKRIS AND LUDOVIC RENOU

ABSTRACT. This paper generalizes the concept of *Bayes correlated equilibrium* (Bergemann and Morris, 2016) to multi-stage games. We apply our characterization results to a number of illustrative examples and applications.

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

JEL CLASSIFICATION: C73, D82.

Date: September 29, 2022.

Ludovic Renou gratefully acknowledges the support of the Agence Nationale pour la Recherche under grant ANR CIGNE (ANR-15-CE38-0007-01) and through the ORA Project “Ambiguity in Dynamic Environments” (ANR-18-ORAR-0005). We thank Laura Doval, Stephen Morris, Sujoy Mukerji, Peter Norman, 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.

1. INTRODUCTION

This paper generalizes the concept of *Bayes correlated equilibrium* (Bergemann and Morris, 2016) to multi-stage games. In a multi-stage game, a set of players interact over several stages and, at each stage, players receive private signals 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.

Consider an analyst, who postulates a multi-stage game, which we call the *base game*, but also acknowledges that players may receive *additional* signals, which can depend on past and current states, past actions, and past and current signals (including the past additional ones). Define an *expansion* of the base game to be a multi-stage game that differs from the base game only in that players receive *additional* signals. Which predictions can the analyst make if he does not want to hypothesize a fixed expansion? Alternatively, consider an information designer who can design the additional information players receive. Which outcomes can the information designer achieve?

Bergemann and Morris (2016) address the above questions within the class of static games. These authors show that the Bayes correlated equilibria of the (static) base game characterize all the predictions the analyst can make or, equivalently, the outcomes the information designer can achieve. (See below for an informal definition of a Bayes correlated equilibrium.) In many economic applications, however, the interaction between the economic agents is best modeled as a dynamic game, where the agents receive information over time and have the opportunity to make multiple decisions.

As an example, consider the refinancing operations of central banks. Typically, central banks organize weekly tender auctions to provide short-term liquidities to financial institutions. While extensive regulations carefully

specify the auction formats central banks use, the information the financial institutions and the central banks receive over time as well as the communication between them are substantially harder to model. An analyst may thus want to postulate a base game, which captures all that is known to him – auction format, public announcements, public statistics – and to remain agnostic about the private information the financial institutions and central banks have. In other words, the analyst considers all possible *expansions* of the base game. Recent contributions in the econometrics literature on partial identification have adopted such an approach. See Bergemann, Brooks and Morris (2019), Gualdani and Sinha (2021), Magnolfi and Roncoroni (2017), Syrgkanis, Tamer, and Ziani (2018).

Our main contribution is methodological. We derive several generalizations of the concept of Bayes correlated equilibrium, where each generalization corresponds to a solution concept for multi-stage games. We focus primarily on the concept of Bayes-Nash equilibrium. While refinements are frequently used in applications, we do so for a simple reason: the logical arguments do not differ from one solution concept to another. Our main theorem (Theorem 1) states an equivalence between (i) the set of all distributions over states, signals in the base game, and actions induced by all *Bayes-Nash equilibria* of all expansions of the base game, and (ii) the set of all distributions over states, signals in the base game, 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, who knows everything that has occurred, 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 stress here

that the “omniscient” mediator is a metaphor, an abstract entity, which only serves as a tool to characterize all the equilibrium outcomes we can obtain by varying the information structures. Whenever we refer to the mediator making recommendations at a given history, it should be understood as the information structure –the statistical experiment– generating the additional signals conditional on the given history.

The logical arguments are simple. Fix an expansion of the base game and a (Bayes-Nash) equilibrium. We show that we can emulate the equilibrium of the expansion as an equilibrium of an auxiliary mediated game, where a dummy (additional) player makes reports to a mediator and the mediator sends messages to the original players. In that auxiliary mediated game, the dummy player knows the actions, signals and states and the messages the mediator sends are the additional signals of the expansion. We can then apply the classical revelation principle of Myerson (1986) and Forges (1986) to replicate the equilibrium of the auxiliary mediated game as a *canonical* equilibrium of the “direct” game, i.e., as an equilibrium of a mediated game where players report their private information to the mediator, the mediator recommends actions and players are truthful and obedient, provided they have been in the past. At that canonical equilibrium, the mediator is “omniscient” at truthful histories and players are obedient provided they have been in the past: we have a Bayes correlated equilibrium. The very same logic generalizes to a variety of other solution concepts. All we need is a revelation principle.

Finally, we provide two illustrations of the broad applicability of our results. In particular, we extend the characterization of de Oliveira and Lamba (2019).

The closest paper to ours is Bergemann and Morris (2016), henceforth BM. These authors characterize the set of distributions over actions, signals in the base game, 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.¹ Two main insights emerge from our generalization.

The first insight of our analysis is that we genuinely need the mediator to make recommendations at *all* histories. To understand the need for this, note that even in dynamic games where all the states and signals about the states are drawn ex-ante, 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. The reason is that players' signals at interim stages may also provide private information about the actions taken by players in earlier stages. For instance, if the base game is a repeated game with imperfect monitoring, a possible expansion is to perfectly inform players of past actions. As a result, if the mediator could not react to deviations that are unobserved by some players, it might not be able to induce the appropriate continuation play. In fact, as the introductory example (Section 2) demonstrates, applying the definition of BM on the strategic form of even the simplest multi-stage games does *not* characterize what we can obtain by considering all equilibria of all expansions of the base game – an extensive-form game. The need to work with games in extensive form in order to characterize their Bayes correlated equilibria echoes a similar observation in Myerson (1986) regarding the gradual release of information in communication games. This is not coincidental: Bayes correlated equilibria are communication equilibria (Forges, 1986; Myerson, 1986) of mediated games with an “omniscient mediator.”

The second insight is that our characterization of Bayes correlated equilibrium in multi-stage games generalizes to any solution concept for which a

¹Our work is related in spirit to Penta (2015) and Penta and Zuazo-Garin (2021). The first paper extends the belief-free approach to robust mechanism design to dynamic problems, in which agents obtain information over time about payoff-relevant states, but where actions are public. The second paper considers the question of robust predictions when common knowledge assumptions about players' information on each other moves are relaxed; payoffs remain common knowledge. In contrast to our paper, where we consider all expansions, Penta and Zuazo-Garin (2021) incorporate only “local” perturbations of the belief hierarchies.

revelation principle holds. In particular, this is true for the two versions of perfect Bayesian equilibrium we consider (Section 5). This is particularly important for many economic applications. Bargaining problems (e.g., Bergemann, Brooks and Morris, 2015), allocation problems with aftermarkets (e.g., Calzolari and Pavan, 2006, Giovannoni and Makris, 2014, and Dworzak, 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.

Doval and Ely (2020) is another generalization of the work of BM and nicely complements our own generalization. These authors take as given states, consequences and state-contingent payoffs over the consequences, and characterize all the distributions 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 less reasonable. 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,

2008), or Bayesian solution (Forges, 1993, 2006).² The concept of Bayes correlated equilibrium is a generalization of all these notions. Solan (2001) is a notable exception. In stochastic games with players perfectly informed of past actions and past and current states, Solan considers general communication devices, where the mediator sends messages to the players as a function of past messages sent and received, and the history of the game, i.e., the past actions, and the past and current states. Solan's mediator is omniscient. For that class of games, Solan shows that the set of Bayes correlated equilibrium payoffs is equal to the set of extensive-form correlated equilibrium payoffs. As we show in the introductory example, this equivalence does not hold if players are not perfectly informed of past actions. See Forges (1985) for a related result.

2. AN INTRODUCTORY EXAMPLE

This section illustrates our main results with the help of a simple example. The example highlights a novel and distinctive aspect of (unconstrained) information design in dynamic games: In addition to providing information about payoff-relevant states, the designer can choose the information players have about the past actions of others. E.g., in voting problems, the designer can choose how much information to reveal about past votes.

Example 1. There are two players, labelled 1 and 2, and two stages. Player 1 chooses either T or B at the first stage, while player 2 chooses either L or R at the second stage. In the base game, player 2 has no information about player 1's choice. Figure 1 depicts the base game, with the payoff of player 1 as the first coordinate.

Suppose that the information designer wants to maximize player 1's payoff. Which information structure(s) should it design?

²The concept of extensive-form correlated equilibrium was first introduced in Forges (1986). The concept introduced in von Stengel and Forges (2008) differs from the one in Forges (1986).




FIGURE 1. The base game

To address that question, we generalize the work of Bergemann and Morris (2016) to dynamic games. Bergemann and Morris (2016) characterize the distributions over outcomes we can induce by varying the information players have in static games of incomplete information—that is, as we vary the information a player has about the payoff-relevant states and the information of others. They prove that the *Bayes correlated equilibria* of the static base game characterize the distributions we can induce. At a Bayes correlated equilibrium, an omniscient mediator recommends actions, and the players have an incentive to be obedient.

Thus, a naive idea is to apply the concept of Bayes correlated equilibrium to the strategic form of the base game. In our example, the unique Bayes correlated equilibrium of the strategic form is (B, R) with a payoff profile of $(1, 1)$. Working with the strategic form is, however, too restrictive. E.g., if the designer perfectly informs player 2 of player 1's action, the induced game has an equilibrium with outcome (T, L) and associated payoff $(2, 2)$. (In static games with complete information, the set of Bayes correlated equilibria coincides with the set of correlated equilibria.)

Our approach is to have the omniscient mediator recommending actions to the players *not only* at the initial history, but *at each* history of the dynamic game. In addition, the players must have an incentive to be obedient, provided they have been obedient in the past. This approach generalizes the definition of Bayes correlated equilibria of Bergemann and Morris (2016) to multi-stage games and illustrates the need to work on the extensive-form

games. (Myerson, 1986, has already pointed out the insufficiency of the strategic form; see Section 2 of his paper.)

We now illustrate how our approach works in our example. Since the mediator is omniscient and makes recommendations at all histories, we need to consider two recommendation kernels. The first kernel specifies the probability of recommending an action to player 1 at the first stage. The second kernel specifies the probability of recommending an action to player 2 at the second stage as a function of the action recommended and chosen at the first stage. Players must have an incentive to be obedient. We claim that there exist such recommendation kernels with a payoff profile of $(5/2, 1)$.

To see this, assume that the mediator recommends with probability $1/2$ player 1 to play T and with the complementary probability to play B at the first stage, and recommends player 2 to play L at the second stage if, and only if, player 1 was obedient at the first stage. (Otherwise, the mediator recommends R .) We now prove that the players have an incentive to be obedient.

If player 2's recommendation is L , he believes that player 1 has played T with probability $1/2$ and, thus, expects a payoff of 1 if he plays L . He therefore has an incentive to be obedient. If player 2's recommendation is R , we are off the equilibrium path and any conjecture that puts probability of at least $1/2$ on player 1 having played B makes R optimal. As for player 1, he clearly has an incentive to be obedient when his recommendation is B since he gets his highest payoff. When his recommendation is T , a deviation to B is unprofitable because this leads player 2 to play R . Thus, we have a Bayes correlated equilibrium with a payoff profile of $(5/2, 1)$.

We now argue that no information structures can give player 1 a payoff higher than $5/2$, hence answering our initial question. Since player 2 can always play R , player 2's payoff cannot be lower than 1. Therefore, within

the set of feasible payoff profiles, conditional on player 2's getting a payoff of at least 1, player 1's highest payoff is $5/2$.

Finally, we now explain how we can use the Bayes correlated equilibrium to design an information structure, whose associated expansion generates an equilibrium payoff of $(5/2, 1)$. The idea is simple: Think of recommendations as signals. Accordingly, suppose that there are two equally likely signals, t and b , at the first stage, and two signals l and r at the second stage. Player 1 privately observes the first signal, while player 2 observes the signal l if only if either (t, T) or (b, B) is the profile of signal and action at the first stage. With such information structure, players have an incentive to play according to their signals and, thus, we obtain the payoff profile $(5/2, 1)$.

As a final observation, consider the alternative base game in Figure 2, where we reverse the order of play. (The first coordinate of a payoff vector refers to player 1's payoff.)




FIGURE 2. An alternative base game: reversing the order of play.

Note that the two base games have the same strategic form. Yet, their set of Bayes correlated equilibria differ. The alternative base game has the unique payoff profile $(1, 1)$ as Bayes correlated equilibrium's payoff. Indeed, on path, player 1 must play B since T is strictly dominated. And since player 2 can guarantee a payoff of 1 by playing R , the unique outcome is (B, R) .

3. MULTI-STAGE GAMES AND EXPANSIONS

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.

3.1. 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.

The description of the base game is very flexible. It can accommodate games with perfect or imperfect observation, repeated games with incomplete information, or stochastic games, among others. We stress, however, that the base game puts certain restrictions on what an information designer can do or what robustness exercises an analyst can conduct. For instance, once the order of moves is fixed, we cannot have a player receiving information (original or additional) about the moves of players who, according to the base game, move at a later stage.

We now introduce some notation. 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 $t \in \{2, \dots, T\}$, $H_{i,1} = S_{i,1}$ the set of initial information, and $H_{i,T+1} = A_{i,T}$ the set of new information at the end of the last stage.

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) .³ 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: $h_i = (h_{i,1}, \dots, h_{i,t}, \dots, h_{i,T+1})$ the private (terminal) history of player i , $h_i^t = (h_{i,1}, \dots, h_{i,t})$ the private history of player i at stage t , $h_t = (h_{1,t}, \dots, h_{n,t})$ the profile of actions played at stage $t - 1$ and signals received at stage t , $h^t = (h_1, \dots, h_t)$ the history of signals and actions at stage t , $\omega = (\omega_1, \dots, \omega_T)$ the profile of realized states, and $\omega^t = (\omega_1, \dots, \omega_t)$ the profile of states realized up to stage t , with corresponding sets $H_i = \{h_i : (h, \omega) \in H\Omega \text{ for some } \omega\}$, $H_i^t = \{h_i^t : (h, \omega) \in H\Omega \text{ for some } \omega\}$, $H_t = \{h_t : (h, \omega) \in H\Omega \text{ for some } \omega\}$, $H^t = \{h^t : (h, \omega) \in H\Omega \text{ for some } \omega\}$, $\Omega = \{\omega : (h, \omega) \in H\Omega \text{ for some } h\}$, $\Omega^t = \{\omega^t : (h, \omega) \in H\Omega \text{ for some } h\}$. 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 notation will apply to other sets. If there is no risk of confusion, we will not formally define these additional notation.

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 .⁴ Throughout, we refer to the signals in S as the *base signals*.

³The sets S_{T+1} and Ω_{T+1} are defined to be a singleton.

⁴This is without loss of generality as we can always redefine the states to include the signals.

3.2. 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 current and past (payoff-relevant) states (such as the valuations for objects in auction problems), but also about the past realization of actions (as in repeated games with imperfect monitoring). Throughout, we use the same notation as in the base game to denote relevant sub-vectors and their corresponding sets.

Formally, an expansion is a collection of sets of additional private signals $(M_{i,t})_{i,t}$ and probability kernels $(\xi_t)_t$ such that all sets of additional signals are non-empty and finite, $\xi_1 : H_1 \times \Omega_1 \rightarrow \Delta(M_1)$, and $\xi_t : H^t \times M^{t-1} \times \Omega^t \rightarrow \Delta(M_t)$ for all $t \geq 2$.⁵ Intuitively, at each stage t , player i receives the additional private signal $m_{i,t} \in M_{i,t}$, with

$$\xi_t(m_t | h^t, m^{t-1}, \omega^t)$$

the probability of m_t when (h^t, m^{t-1}, ω^t) is the history of actions, base signals, states, and past additional signals at the beginning of stage t . We write M for the collection $(M_{i,t})_{i,t}$ and ξ for $(\xi_t)_t$.

Together with the base game Γ , an expansion (M, ξ) induces a multi-stage game Γ_π , where at each stage t , a payoff-relevant state ω_t is realized, player i receives the private signal $(s_{i,t}, m_{i,t})$ and takes an action $a_{i,t}$. To complete the description of the induced multi-stage game, we let $\pi_1(h_1, m_1, \omega_1) := \xi_1(m_1 | h_1, \omega_1) p_1(h_1, \omega_1)$ be the probability of (h_1, m_1, ω_1) at the first stage and

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

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 . With a slight abuse of language, we use the word

⁵The set $M_{i,T+1}$ is a singleton.

“expansion” to refer to the collection of additional signals and kernels (M, ξ) as well as to the multi-stage game Γ_π induced by it.

It is worth stressing that our definition of expansions implies that the evolution of future states and base signals, as a function of earlier actions, states and base signals, is known to the analyst/designer in the sense of being pinned down by the base game (and, hence, remains constant across all expansions). Additional signals are just information and *do not cause* future states.

We denote $HM\Omega$ the set of all terminal histories, with $(h, m, \omega) \in HM\Omega$ if and only if $(h, \omega) \in H\Omega$. We do not require, however, histories $(h, m, \omega) \in HM\Omega$ to have strictly positive probability for some profile of actions, i.e., for all $(h, m, \omega) \in HM\Omega$, we do not require the existence of $a \in A$ such that

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

In other words, we do not prune out the additional messages, even when they have zero probability under all action profiles. (Note that $(h, \omega) \in H\Omega$ corresponds to $\sum_m \pi^a(h, m, \omega) > 0$ for some a .)

In closing, it is worth noting that an expansion ξ induces a collection of kernels π , with the property that $\text{marg}_{H\Omega} \pi^a = p^a$ for all $a \in A$, that is,

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

We call this property consistency. As explained in the introduction, in static problems, the converse is also true, that is, any consistent kernel π induces an expansion ξ . This equivalence breaks down in dynamic problems. This is the case when the additional signals are not just signals, but also *cause* the realization of future states and base signals, so that the base game stops pinning down their evolution (see Example 2 below).

4. A FIRST EQUIVALENCE THEOREM AND AN APPLICATION

This section contains our first characterization theorem—other characterizations will differ by the solution concepts adopted. In Section 4.1, we consider the concept of Bayes-Nash 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. As we will see, the main arguments extend almost verbatim to other solution concepts. In addition, if we are interested in proving an impossibility result, e.g., whether efficiency obtains, the weaker the solution concept, the stronger the result. Section 5 extends our analysis to two refinements of the concept of Bayes-Nash equilibrium, which all impose sequential rationality.

4.1. A First Equivalence Theorem. We first define the concepts of Bayes-Nash equilibrium and Bayes correlated equilibrium. Throughout, we fix an expansion Γ_π of Γ . A behavioral strategy σ_i is a collection of maps $(\sigma_{i,t})_{t \in T}$, with $\sigma_{i,t} : H_i^t M_i^t \rightarrow \Delta(A_{i,t})$.

Definition 1 (BNE): A profile σ of behavioral strategies is a *Bayes-Nash equilibrium* of Γ_π if

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

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 $\mathcal{BN}\mathcal{E}(\Gamma_\pi)$ be the set of distributions over $H\Omega$ induced by the Bayes-Nash equilibria of Γ_π .

We now state formally the main objective of our paper: we want to provide a characterization of the set $\bigcup_{\Gamma_\pi \text{ an expansion of } \Gamma} \mathcal{BN}\mathcal{E}(\Gamma_\pi)$, i.e., we want to characterize the distributions over the outcomes $H\Omega$ of the base game Γ that we can induce by means of some expansion Γ_π of the base game, *without*

any reference to particular expansions. To do so, we need to introduce the concept of Bayes correlated equilibrium of Γ .

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 an “omniscient” mediator – a mediator who knows everything that has occurred in $\mathcal{M}(\Gamma)$ – and chooses an action $a_{i,t}$. We let $\tau_{i,t} : H_i^t \times A_i^t \rightarrow \Delta(A_{i,t})$ be an action strategy at period t and write $\tau_{i,t}^*$ for the obedient strategy. We write τ_i for $(\tau_{i,t})_t$ and τ for $(\tau_i)_i$.

Definition 2 (BCE): A Bayes correlated equilibrium of Γ is a collection of recommendation kernels $\mu_t : H^t \Omega^t \times A^{t-1} \rightarrow \Delta(A_t)$ such that τ^* is an equilibrium of the mediated game $\mathcal{M}(\Gamma)$, that is,

$$\sum_{\mathbf{h}, \boldsymbol{\omega}, \hat{\mathbf{a}}} u_i(\mathbf{h}, \boldsymbol{\omega}) \mathbb{P}_{\mu \circ \tau^*, p}(\mathbf{h}, \boldsymbol{\omega}, \hat{\mathbf{a}}) \geq \sum_{\mathbf{h}, \boldsymbol{\omega}, \hat{\mathbf{a}}} u_i(\mathbf{h}, \boldsymbol{\omega}) \mathbb{P}_{\mu \circ (\tau_i, \tau_{-i}^*), p}(\mathbf{h}, \boldsymbol{\omega}, \hat{\mathbf{a}})$$

for all τ_i , for all i , with $\mathbb{P}_{\mu \circ \tilde{\tau}, p}$ denoting the distribution over profiles of actions, base signals, states and recommendations induced by $\mu \circ \tilde{\tau}$ and p .

We let $\mathcal{BCE}(\Gamma)$ be the set of distributions over $H\Omega$ induced by the Bayes correlated equilibria of Γ . The set $\mathcal{BCE}(\Gamma)$ is convex.

It is instructive to compare the concept of Bayes correlated equilibrium and communication equilibrium (Forges, 1986, Myerson, 1986). In a communication equilibrium, the mediator relies on the information provided 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—the mediator is omniscient. Let $\mathcal{CE}(\Gamma)$ be the distributions over $H\Omega$ induced by the communication equilibria of Γ . For all multi-stage games Γ , we have that $\mathcal{CE}(\Gamma) \subseteq \mathcal{BCE}(\Gamma)$ since the omniscient mediator can always replicate the Forges-Myerson mediator. Since we also have that $\mathcal{BNE}(\Gamma) \subseteq \mathcal{CE}(\Gamma)$, we have the inclusion $\mathcal{BNE}(\Gamma) \subseteq \mathcal{BCE}(\Gamma)$. However, it is a priori unclear whether $\mathcal{BNE}(\Gamma_\pi) \subseteq \mathcal{BCE}(\Gamma)$ for all expansions

Γ_π of Γ since players have additional signals in Γ_π , while the omniscient mediator of Γ has no additional signals. A consequence of our main result, Theorem 1, is that it is indeed the case.

Theorem 1. *We have the following equivalence:*

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

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 and (ii) the set of distributions over actions, base signals and states we can obtain by considering *all* Bayes-Nash equilibria of *all* expansions of Γ .⁶

Theorem 1 generalizes the work of BM to multi-stage games. As the introductory example demonstrates, our definition of a Bayes correlated equilibrium is, in general, weaker than applying the definition of BM to 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 base signals. Yet, in multi-stage base games where (i) all the states and base signals are drawn at the first stage, (ii) players observe their signals sequentially over time, and (iii) past actions are perfectly observable, the two formulations are equivalent. To see this, we first define the concept of “ex-ante” Bayes-correlated equilibrium. To ease notation, assume that the profile (s, ω) of base signals and states is drawn with probability $p(s, \omega)$ at the first stage, and no states and base signals are drawn at later stages. An “ex-ante” Bayes-correlated equilibrium is a kernel $\bar{\mu} : S \times \Omega \rightarrow \Delta(\Sigma)$ (where

⁶In a supplementary document, we also prove the equivalence with the set of distributions over actions, base signals and states we can obtain by considering *all* communication equilibria of *all* expansions of Γ .

Σ is the set of strategies of Γ), which satisfies

$$\begin{aligned} \sum_{\omega, s, \sigma_{-i}, a} p(s, \omega) \bar{\mu}(\sigma | s, \omega) \mathbb{P}_{(\sigma_i, \sigma_{-i})}(\mathbf{a} | s) u_i(\mathbf{a}, \omega) \geq \\ \sum_{\omega, s, \sigma_{-i}, a} p(s, \omega) \bar{\mu}(\sigma | s, \omega) \mathbb{P}_{(\sigma'_i, \sigma_{-i})}(\mathbf{a} | s) u_i(\mathbf{a}, \omega), \end{aligned}$$

for all σ'_i, σ_i, i . In words, players must have an incentive to be obedient.

Clearly, any distribution induced by an “ex-ante” Bayes-correlated equilibrium is in $\mathcal{BCE}(\Gamma)$. Indeed, we can interpret the kernel $\bar{\mu}$ as an expansion, where $M_1 = \Sigma$ and $\xi_1 = \bar{\mu}$, and an equilibrium of that expansion is to play according to one’s additional signal—the recommended strategy.

Conversely, consider a Bayes-correlated equilibrium μ of such multi-stage games, where $\mu_t(\hat{a}_t | \hat{a}^{t-1}, a^{t-1}, s, \omega)$ is the probability of recommending \hat{a}_t , when the profile of past recommendations and actions is (\hat{a}^{t-1}, a^{t-1}) and the profile of states and signals is (s, ω) . The intuition is best understood when there is a single stage. At a Bayes-correlated equilibrium, the omniscient mediator, who knows (s, ω) , recommends a_i to player i with probability $\mu(a_i | s, \omega) = \sum_{a_{-i}} \mu((a_i, a_{-i}) | s, \omega)$, and player i has an incentive to be obedient when his signal is indeed s_i . An alternative, but equivalent, formulation would be to recommend a strategy $\sigma_i^{[a_i, s_i]}$ to player i , which stipulates to play a_i when the signal is s_i . Player i would indeed have an incentive to follow that strategy when his signal is s_i . When there are more stages, a similar construction is possible. We now sketch it and refer the reader to the supplementary materials for the details.

The idea is to associate with $(\mu_t(\cdot | \cdot, \cdot, s, \omega))_t$ an outcome-equivalent distribution over feedback rules $f = (f_t : A^{t-1} \times A^{t-1} \rightarrow A_t)_t$ (as in Kuhn’s theorem), and with each feedback rule f and signal s , a strategy $\sigma^{[f, s]}$ such that players have an incentive to be obedient. (When there is a single stage, feedback rules are actions.) Formally, for each feedback rule f , define recursively: $f^1(\emptyset) = f_1(\emptyset, \emptyset)$ and $f^t(a^{t-1}) = (f^{t-1}(a^{t-2}), f_t(a^{t-1}, f^{t-1}(a^{t-2})))$ for all a^{t-1} , for all t . We now associate a pure strategy $\sigma_i^{[f, s]}$ with the feedback rule f and the profile of signals s , as follows: $\sigma_{i,t}^{[f, s]}(s_i^t, a^{t-1}) = f_{i,t}(a^{t-1}, f^{t-1}(a^{t-2}))$ for all a^{t-1} ,

for all t , with $f_{i,t}(a^{t-1}, f^{t-1}(a^{t-2}))$ the i -th component of $f_t(a^{t-1}, f^{t-1}(a^{t-2}))$. In words, $\sigma_i^{[f,s]}$ dictates player i to play as f dictates when his signal's realizations are (s_i) . Let $\sigma^{[f,s]}$ be the profile $(\sigma_i^{[f,s]})_i$. It is important to stress that the strategy is well-defined because all past actions are perfectly observed. Without that assumption, this wouldn't be the case. Finally, if the mediator recommends $\sigma^{[f,s]}$ with probability

$$\bar{\mu}(\sigma^{[f,s]}|s, \omega) = \prod_t \prod_{a^{t-1}, \hat{a}^{t-1}} \mu_t(f_t(a^{t-1}, \hat{a}^{t-1})|a^{t-1}, \hat{a}^{t-1}, s, \omega),$$

then it is routine to verify that the players have an incentive to be obedient and the desired distribution over actions, signals and states, is implemented. The equivalence, however, breaks down if either actions are not perfectly observed (as in the introductory example) or states and signals are not drawn at the first stage, e.g., because the players control their evolutions through their actions.

As in BM, the proof of $\mathcal{BCE}(\Gamma) \subseteq \bigcup_{\Gamma_\pi \text{ an expansion of } \Gamma} \mathcal{BNE}(\Gamma_\pi)$ is constructive. However, unlike BM's constructive proof, our proof of $\bigcup_{\Gamma_\pi \text{ an expansion of } \Gamma} \mathcal{BNE}(\Gamma_\pi) \subseteq \mathcal{BCE}(\Gamma)$ is non-constructive: it utilizes the revelation principle of Forges (1986) and Myerson (1986). This approach has two main advantages: (i) it reveals the main logical arguments, which are somewhat hidden in constructive proofs, and (ii) its generalization to many other solution concepts is straightforward. The central arguments are the following. Consider an expansion Γ_π of Γ and an equilibrium distribution $\mu^d \in \mathcal{BNE}(\Gamma_\pi)$. By definition, there exists a Bayes-Nash equilibrium σ of Γ_π , which induces μ^d . The main idea is to replicate the expansion Γ_π and its equilibrium σ as a Bayes-Nash equilibrium of an *auxiliary mediated* game $\mathcal{M}^*(\Gamma)$, which we now describe. The game $\mathcal{M}^*(\Gamma)$ has one additional dummy player, called player 0, and a Forges-Myerson mediator, who receives reports by and sends messages to the players. At the first stage, Nature draws (h_1, ω_1) with probability $p_1(h_1, \omega_1)$, player i observes $h_{i,1}$ and player 0 observes (h_1, ω_1) . Player 0 then reports $(\hat{h}_1, \hat{\omega}_1)$ to the mediator; all other players do not make reports (their sets of reports are singletons). The mediator then draws the message

m_1 with probability $\xi_1(m_1|\hat{h}_1, \hat{\omega}_1)$ and sends $m_{i,1}$ to player i . Player 0 does not receive a message. Finally, player i chooses an action $a_{i,1}$; player 0 does not take an action. Consider now a history $(a_{t-1}, h^{t-1}, \omega^{t-1})$ of past actions, signals and states and a history $((\hat{h}^{t-1}, \hat{\omega}^{t-1}), m^{t-1})$ of reports and messages. Stage t unfolds as follows:

- Nature draws (h_t, ω_t) with probability $p_t(h_t, \omega_t|a_{t-1}, h^{t-1}, \omega^{t-1})$.
- Player $i \in I$ observes the signal $h_{i,t}$ and player 0 observes (h_t, ω_t) .
- Player 0 reports $(\hat{h}_t, \hat{\omega}_t)$ to the mediator. All other players do not make reports.
- The mediator draws the message m_t with probability $\xi_t(m_t|\hat{h}^t, m^{t-1}, \hat{\omega}^t)$ and sends the message $m_{i,t}$ to player i . Player 0 does not receive a message.
- Player i takes an action $a_{i,t}$. Player 0 does not take an action.

If player 0 is truthful and each player $i \in I$ follows σ_i , we have a Bayes-Nash equilibrium of the mediated game $\mathcal{M}^*(\Gamma)$, with equilibrium distribution μ^d . From the revelation principle of Forges (1986) and Myerson (1986), there exists a *canonical* equilibrium which implements μ^d , i.e., an equilibrium where players report their private information to the mediator, the mediator recommends actions and players are truthful and obedient provided that they have been in the past.⁷ At truthful histories, the mediator is omniscient and players have an incentive to be obedient provided they have been in the past: this is the Bayes correlated equilibrium.

Before applying Theorem 1, two additional remarks are worth making. First, the above arguments are not limited to the concept of Bayes-Nash equilibrium. The same arguments apply to all solution concepts, such as weak perfect Bayesian equilibrium or conditional probability perfect Bayesian equilibrium, which admit a revelation principle. We formally state these equivalences below. Second, the above arguments clearly demonstrate the

⁷Proposition 2 in Sugaya and Wolitzky (2021) is a restatement of the revelation principle of Forges (1986) and Myerson (1986). It also applies to mediated games, where the mediator receives private signals in addition to the players' reports. With such restatement, we can dispense with the dummy player and have the mediator directly learn (h_t, ω_t) .

role our definition of an expansion plays. It makes it possible for the mediator to replicate any expansion as the kernels ξ_t are assumed measurable with respect to the mediator's histories. With the alternative and weaker definition of an expansion as a consistent information structure, i.e., $\text{marg } \pi^a = p^a$ for all a , it is no longer guaranteed that the mediator, despite being omniscient, can simulate any expansion, as the next example illustrates.

Example 2. 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 $p_2(\omega_2 = 1|a_1 = 0) = 1/2$. 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 information structure: $M_1 = \{0, 1\}$, 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 information structure is consistent, but there are no kernels (ξ_1, ξ_2) that induce this information structure from the base game. The issue is that action a_1 and additional signal m_1 *jointly cause* the second-stage state ω_2 . In other words, the signal m_1 is not just additional information; it also determines the realization of the second-period state.

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 $\bar{\mu}_1(a_1)p_2(\omega_2|a_1)$ and there is no $\bar{\mu}_1$ that induces the distribution μ .

4.2. Application I: Rationalizing Dynamic Choices. Suppose that an analyst observes the choices of a decision-maker over a finite number of periods, but does not observe the information the decision-maker had. Suppose, furthermore, that the analyst assumes that the state does not change over time. Which profiles of choices can be rationalized? This question was recently addressed by de Oliveira and Lamba (2019), under the assumption that the information the decision-maker receives over time is independent of his actions. With our notation, this is equivalent to requiring that $\xi_t(\cdot|a^{t-1}, m^{t-1}, \omega) = \xi_t(\cdot|\tilde{a}^{t-1}, m^{t-1}, \omega)$ for all $(a^{t-1}, \tilde{a}^{t-1})$, for all m^{t-1} , for all t .⁸ We call such expansions *autonomous*.

We now show how Theorem 1 makes it possible to extend their result to *all* expansions. Throughout, we follow the terminology of de Oliveira and Lamba.

We say that the profile of choices $a^* := (a_1^*, \dots, a_T^*)$ is *rationalizable* if there exist a probability $p \in \Delta(\Omega)$, sets of signals M_t and kernels $\xi_t : A^{t-1} \times M^{t-1} \times \Omega \rightarrow \Delta(M_t)$ such that the decision-maker chooses optimally and (a_1^*, \dots, a_T^*) is optimal for some realizations (ω, m) of states and signals. As de Oliveira and Lamba, we assume the decision-maker payoff function u is known to the analyst.

From Theorem 1, the profile of choices if there exists a probability $p \in \Delta(\Omega)$ and a Bayes correlated equilibrium μ such that $\mathbb{P}_{\mu \circ \tau^*, p}(a^*) > 0$. Recall that μ is a Bayes correlation equilibrium if:

$$\sum_{a, \omega, \hat{a}} u(a, \omega) \mathbb{P}_{\mu \circ \tau^*, p}(a, \omega, \hat{a}) \geq \sum_{a, \omega, \hat{a}} u(a, \omega) \mathbb{P}_{\mu \circ \tau, p}(a, \omega, \hat{a}),$$

for all τ . The objective is to derive conditions on the primitives, which guarantee the existence of such a Bayes correlated equilibrium.

⁸As all sets of base signals are singletons, we do not denote them. Similarly, since the state does not change over time, we write ω for the fully persistent state.

We need some additional notation. We say that $D : A \rightarrow \Delta(A)$ is a *deviation plan* if there exists a behavioral strategy τ such that

$$D(a_1, \dots, a_T | \hat{a}_1, \dots, \hat{a}_T) := \prod_{t=1}^T \tau_t(a_t | \underbrace{(\hat{a}_1, \dots, \hat{a}_t)}_{\text{recommendations}}, \underbrace{(a_1, \dots, a_{t-1})}_{\text{choices}})$$

for all (\hat{a}, a) . A deviation plan specifies what the decision-maker would do if he were to face a fixed sequence of recommendations.

We say that $\gamma : A \rightarrow A$ is a *recommendation plan* if there exist maps $\gamma_t : A^{t-1} \rightarrow A_t$ such that

$$\gamma(a) = \left(\gamma_1(\emptyset), \gamma_2(a_1), \dots, \gamma_t(a_1, \dots, a_{t-1}), \dots, \gamma_T(a_1, \dots, a_{T-1}) \right)$$

for all a , with the convention that $A^0 = \{\emptyset\}$. A recommendation plan specifies a recommendation for each fixed sequence of choices. We write $\gamma^{\geq t}(a)$ for $(\gamma_t(a_1, \dots, a_{t-1}), \dots, \gamma_T(a_1, \dots, a_{T-1}))$.

Finally, we define the sets: $B_a^1 := (A_1 \setminus \{a_1\}) \times A_2 \times \dots \times A_T$, $B_a^t := \{(a_1, \dots, a_{t-1})\} \times (A_t \setminus \{a_t\}) \times A_{t+1} \times \dots \times A_T$ for all $t \in \{2, \dots, T-1\}$, and $B_a^T := \{(a_1, \dots, a_{T-1})\} \times A_T$. (By convention, if $T = 1$, then $B_a^1 = A_1$.) For all $t < T$, the set B_a^t is the set of all profiles of choices, which coincide with a up to period $t-1$ and differ from a at period t . Note that the sets are pairwise disjoint and $\bigcup_{t=1}^T B_a^t = A$ for all a . We are now ready to state our main definition.

Definition 1. *The profile a^* is surely dominated if there exists a deviation plan D such that for all ω , for all a , for all recommendation plans γ :*

$$u(a^*, \omega) < \sum_{t=1}^T \sum_{b \in B_{a^*}^t} u(b, \omega) D(b | a_1^*, \dots, a_t^*, \gamma^{\geq t+1}(b)), \quad (\text{D1})$$

$$u(a, \omega) \leq \sum_{t=1}^T \sum_{b \in B_a^t} u(b, \omega) D(b | a_1, \dots, a_t, \gamma^{\geq t+1}(b)). \quad (\text{D2})$$

Intuitively, a profile of choices is surely dominated if the decision-maker has a deviation plan, which guarantees an improvement regardless of the state ω , the period t at which the decision-maker is first disobedient, and

the subsequent recommendations $\gamma^{\geq t+1}(b)$, which may depend on the past choices. (By convention, the profile $(a_1, \dots, a_T, \gamma^{\geq T+1}(b)) = (a_1, \dots, a_T)$ for all b .) We have the following characterization.

Theorem 2. *The profile of choices a^* is rationalizable if, and only if, it is not surely dominated.*

To get some intuition for Theorem 2, let us first consider the special case of a *single* decision ($T = 1$). When the decision-maker makes a single decision, it is well-known that a^* is rationalizable if, and only if, it is not strictly dominated, that is, there does not exist a mixed action $\alpha \in \Delta(A)$ such that $u(a^*, \omega) < \sum_b u(b, \omega)\alpha(b)$ for all ω . We now argue that sure dominance is equivalent to strict dominance. To see this, if a^* is strictly dominated by α , then choosing $D(\cdot|a^*) = \alpha$ and $D(\cdot|a) = \mathbb{1}(a)$ for all $a \neq a^*$ guarantees that a^* is surely dominated. Conversely, choosing $\alpha = D(\cdot|a^*)$ guarantees that a^* is strictly dominated.

To understand the role of (D2), we need to consider genuine dynamic problems. So, let us assume that the decision-maker has to choose twice ($T = 2$). To play $a^* = (a_1^*, a_2^*)$ with positive probability, the decision-maker must find it optimal to play a_1^* at the first period given that he will play a_2^* with some probability at the second period and other actions with the complementary probability. Condition (D2) guarantees that the decision-maker would not find it optimal to play a_1^* at the first period, regardless of what he would play at the second period. Finally, the role of γ is to capture the fact that the recommendations and, therefore, the information the decision-maker receives may depend on his past actions.

As already mentioned, de Oliveira and Lamba (2019) were the first to consider the problem of rationalizing dynamic choices, restricting attention to autonomous expansions. We extend their work in that we consider *all* expansions. It is instructive to compare their characterization with ours. Their main result states that the profile a^* is rationalizable if, and only if, it

is not *truly* dominated. An action profile is truly dominated if there exists a deviation plan that strictly increases the payoff along the action sequence without worsening payoffs at other parts of the decision tree, regardless of the state. Formally, the profile a^* is truly dominated if there exists a deviation rule D such that:

$$\begin{aligned} u(a^*, \omega) &< \sum_b u(b, \omega) D(b|a^*), \\ u(a, \omega) &\leq \sum_b u(b, \omega) D(b|a), \end{aligned}$$

for all ω , for all a .

Clearly, if a profile a^* is *surely* dominated, then it is *truly* dominated. Indeed, if we choose $\gamma^{\geq t+1}(b)$ to be equal to (a_{t+1}, \dots, a_T) for all (a, b, t) , then we recover the conditions for true dominance.

To understand the differences, recall that de Oliveira and Lamba (2019) restrict attention to autonomous expansions. With such a restriction, it is without loss of generality to assume that *all* signals are drawn ex-ante and then they are gradually released to the decision-maker, independently of what he does. It is easy to prove that this is equivalent to restricting attention to recommendation kernels, which depend on past recommendations, base signals and states, but *not* on past actions, i.e., $\mu_t(\cdot|a^{t-1}, s^t, \omega^t, \hat{a}^{t-1}) = \mu_t(\cdot|b^{t-1}, s^t, \omega^t, \hat{a}^{t-1})$ for all (a^{t-1}, b^{t-1}) . (See the supplementary material for a proof.) We therefore have:

$$\begin{aligned} \mathbb{P}_{\mu \circ \tau, p}(a, \omega, \hat{a}) &= \mathbb{P}_{\mu \circ \tau, p}(\omega, \hat{a}) \mathbb{P}_{\mu \circ \tau, p}(a|\omega, \hat{a}) \\ &= \mathbb{P}_{\mu \circ \tau^*, p}(\omega, \hat{a}) [\tau_1(a_1|\hat{a}^1) \times \dots \times \tau_t(a_t|\hat{a}^t, a^{t-1}) \times \dots \times \tau_T(a_T|\hat{a}^T, a^{T-1})] \\ &= \mathbb{P}_{\mu \circ \tau^*, p}(\omega, \hat{a}) D(a|\hat{a}), \end{aligned}$$

where D is a deviation plan. The equality $\mathbb{P}_{\mu \circ \tau, p}(\omega, \hat{a}) = \mathbb{P}_{\mu \circ \tau^*, p}(\omega, \hat{a})$ follows from the fact that the recommendations made and the state realized are independent of the decision-maker's choices and, thus, do not depend on τ .

We can thus rewrite the obedience constraint as:

$$\sum_{\hat{a}, \omega} \mathbb{P}_{\mu \circ \tau^*, p}(\omega, \hat{a}) u(\hat{a}, \omega) \geq \sum_{\hat{a}, a, \omega} \mathbb{P}_{\mu \circ \tau^*, p}(\omega, \hat{a}) D(a|\hat{a}) u(a, \omega),$$

for all deviation plans D . To recover the characterization of de Oliveira and Lamba, it then suffices to follow the same steps as in the proof below, starting with the above rewriting of the obedience constraint.

We conclude this discussion with an example, which demonstrates that a profile can be truly dominated and not surely dominated. There are two states, ω and ω' , three actions, ℓ (eft), c (enter), r (ight), and two periods. The inter-temporal payoff is the sum of the per-period payoff in Table 1.

TABLE 1. (ℓ, c) is rationalizable and truly dominated

	ℓ	c	r
ω	0	1	0
ω'	0	0	1

We now argue that (ℓ, c) is *truly* dominated. Intuitively, since ℓ is strictly dominated, the decision-maker benefits from playing a mixture of c and r instead of ℓ in the first period. More formally, consider the behavioral strategy τ given by $\tau_1(c|\ell) = \tau_1(r|\ell) = 1/2$, $\tau_1(r|r) = \tau_1(c|c) = 1$ and $\tau_2 = \tau_2^*$. The induced deviation rule is $D(c\ell|\ell\ell) = D(r\ell|\ell\ell) = D(cc|\ell c) = D(rc|\ell c) = D(cr|\ell r) = D(rr|\ell r) = 1/2$ and $D(a_1 a_2|\hat{a}_1 \hat{a}_2) = 1$ for all other profiles (a_1, a_2) and (\hat{a}_1, \hat{a}_2) such that $(a_1, a_2) \neq (\hat{a}_1, \hat{a}_2)$. It is then easy to verify that (ℓ, c) is indeed truly dominated.

Yet, it is not surely dominated and, therefore, is rationalizable. Intuitively, if the decision-maker learns the state after playing ℓ in the first period but does not get any additional information otherwise, he has an incentive to play ℓ . A Bayes correlated equilibrium is as follows: the mediator recommends ℓ at the first period, regardless of the state, and recommends c (resp., r) at the second period if and only if the decision-maker has been obedient

and the state is ω (resp., ω'). If the decision-maker disobeys the recommendation, the mediator recommends then either c or r , independently of the state.

Proof of Theorem 2. We first rewrite the obedience constraint. Let $f = (f_1, \dots, f_T)$ be a feedback rule, with $f_t : A^{t-1} \times A^{t-1} \times \Omega \rightarrow A$. A feedback rule specifies a deterministic recommendation at each history of past actions, recommendations and states. A feedback rule is a pure strategy of the mediator. Let \mathcal{F} be the finite set of all feedback rules and \mathcal{F}_ω^* be the non-empty subset of feedback rules, which recommend a^* on path when the state is ω . That is, $f^* \in \mathcal{F}_\omega^*$ if $f_1^*(\omega) = a_1^*$ and

$$f_t^*((a_1^*, \dots, a_{t-1}^*), ((a_1^*, \dots, a_{t-1}^*)), \omega) = a_t^*,$$

for all $t \geq 2$. We write $\mathcal{F}^*\Omega$ for $\cup_{\omega \in \Omega} (\mathcal{F}_\omega^* \times \{\omega\})$.

Similarly, we associate a pure strategy τ with an action rule $g = (g_1, \dots, g_T)$, where $g_t : A^{t-1} \times A^{t-1} \times A \rightarrow A$. The action rule g specifies a pure action at each history of past actions and past and current recommendations. We associate τ^* with the rule g^* , where $g_t^*(a^{t-1}, \hat{a}^{t-1}, \hat{a}) = \hat{a}$. Let \mathcal{G} be the set of action rules.

Thanks to Kuhn's theorem, we can rewrite the condition for rationalization as: there exists $\mu \in \Delta(\mathcal{F} \times \Omega)$ such that $\mu(\mathcal{F}^*\Omega) > 0$ and

$$\sum_{f, \omega, g} \sum_{a, \hat{a}} u(a, \omega) \left(\mathbb{P}_{f, \omega, g^*}(a, \hat{a}) - \mathbb{P}_{f, \omega, g}(a, \hat{a}) \right) \mu(f, \omega) \nu(g) \geq 0,$$

for all $\nu \in \Delta(\mathcal{G})$, where $\mathbb{P}_{f, \omega, g}$ is the degenerate distribution over actions and recommendations induced by the feedback rule f and the action rule g when the state is ω .⁹

We first prove necessity. We prove that if a^* is not rationalizable, then a^* is surely dominated. So, assume that a^* is not rationalizable. For all μ such

⁹E.g., $\hat{a} = (f_1(\omega), f_2(f_1(\omega)), g_1(f_1(\omega)), \omega, \dots)$ and $a = (g_1(f_1(\omega)), g_2(f_2(f_1(\omega)), g_1(f_1(\omega))), \omega, \dots)$

that $\mu(\mathcal{F}^*\Omega) > 0$, there exists ν such the obedience constraint is violated, i.e.,

$$\sup_{\mu:\mu(\mathcal{F}^*\Omega)>0} \min_{\nu} \sum_{f,\omega,g} \sum_{a,\hat{a}} u(a,\omega) \left(\mathbb{P}_{f,\omega,g^*}(a,\hat{a}) - \mathbb{P}_{f,\omega,g}(a,\hat{a}) \right) \mu(f,\omega) \nu(g) < 0.$$

Since the set of μ such that $\mu(\mathcal{F}^*\Omega) > 0$ is non-empty and convex (but not compact) and the objective is bi-linear in (μ, ν) , we can apply Proposition I.1.3 from Mertens, Sorin and Zamir, (2015, p. 6) to obtain:

$$\min_{\nu} \sup_{\mu:\mu(\mathcal{F}^*)>0} \sum_{f,\omega,g} \sum_{a,\hat{a}} u(a,\omega) \left(\mathbb{P}_{f,\omega,g^*}(a,\hat{a}) - \mathbb{P}_{f,\omega,g}(a,\hat{a}) \right) \mu(f,\omega) \nu(g) < 0.$$

Hence, there exists $\bar{\nu}$ such that for all μ with $\mu(\mathcal{F}^*\Omega) > 0$,

$$\sum_{f,\omega,g} \sum_{a,\hat{a}} u(a,\omega) \left(\mathbb{P}_{f,\omega,g^*}(a,\hat{a}) - \mathbb{P}_{f,\omega,g}(a,\hat{a}) \right) \mu(f,\omega) \bar{\nu}(g) < 0.$$

Necessity then follows by first constructing the behavioral strategy $\bar{\tau}$ induced by $\bar{\nu}$ and, therefore, its associated deviation plan \bar{D} and second considering all (f, ω) .

More precisely, fix an arbitrary ω and any feedback rule f^* such that $f_1^*(\omega) = a_1^*$ and

$$f_t^*((a_1^*, \dots, a_{t-1}^*), ((a_1^*, \dots, a_{t-1}^*)), \omega) = a_t^*,$$

for all $t \geq 2$. Let μ be degenerate on (f^*, ω) . Note that $(f^*, \omega) \in \mathcal{F}^*\Omega$. We have that:

$$\sum_{f,\omega,g} \sum_{a,\hat{a}} u(a,\omega) \mathbb{P}_{f,\omega,g^*}(a,\hat{a}) \mu(f,\omega) \bar{\nu}(g) = u(a^*, \omega),$$

while

$$\begin{aligned} \sum_{f,\omega,g} \sum_{a,\hat{a}} u(a,\omega) \mathbb{P}_{f,\omega,g}(a,\hat{a}) \mu(f,\omega) \bar{\nu}(g) &= \sum_g \sum_{a,\hat{a}} u(a,\omega) \mathbb{P}_{f^*,\omega,g}(a,\hat{a}) \bar{\nu}(g) \\ &= \sum_{a,\hat{a}} u(a,\omega) \mathbb{P}_{f^*,\bar{\tau}}(a,\hat{a}) \\ &= \sum_a u(a,\omega) \bar{D}(a|\gamma^*(a)), \end{aligned}$$

where γ^* is the recommendation plan induced by $\gamma_1^*(\emptyset) = f_1^*(\omega)$ and

$$\gamma_t^*(a_1, \dots, a_{t-1}) = f_t^*((a_1, \dots, a_{t-1}), (\gamma_1(\emptyset), \dots, \gamma_{t-1}(a_1, \dots, a_{t-2})), \omega).$$

for all $t \geq 2$. Since $\gamma_t^*(a_1^*, \dots, a_{t-1}^*) = a_t^*$, this is readily seen to be equivalent to the right-hand side of (D1).

We now prove the necessity of (D2). The arguments are nearly identical to the above ones. Fix an arbitrary ω , an arbitrary profile a and a feedback rule f such that $f_1(\omega) = a_1$ and

$$f_t((a_1, \dots, a_{t-1}), ((a_1, \dots, a_{t-1})), \omega) = a_t,$$

for all $t \geq 2$. If $f \in \mathcal{F}^*\Omega$, we can repeat the above arguments. If, however, $f \notin \mathcal{F}^*\Omega$, choose μ such that $\mu(f, \omega) = 1 - \varepsilon$ and $\mu(f^*, \omega) = \varepsilon$, where $1 > \varepsilon > 0$ and f^* is the feedback rule defined above. From the above steps, for all $\varepsilon > 0$, we have that

$$u(a, \omega) < (1 - \varepsilon) \sum_b u(b, \omega) \bar{D}(b|\gamma(b)) + \varepsilon \sum_b u(b, \omega) \bar{D}(b|\gamma^*(b)),$$

where γ is the recommendation plan induced by $\gamma_1(\emptyset) = f_1(\omega)$ and

$$\gamma_t(b_1, \dots, b_{t-1}) = f_t((b_1, \dots, b_{t-1}), (\gamma_1(\emptyset), \dots, \gamma_{t-1}(b_1, \dots, b_{t-2})), \omega).$$

for all $t \geq 2$. Taking the limit as $\varepsilon \rightarrow 0$ and noting that $\gamma_t(a_1, \dots, a_{t-1}) = a_t$, we obtain the condition (D2).

The proof of sufficiency is immediate and left to the reader. ■

To conclude, this application illustrates how we can apply our results to derive testable implications in dynamic decision problems. We stress that our results apply equally to dynamic games, including games with evolving states, and thus offer a wide scope for applications.

5. ADDITIONAL EQUIVALENCE THEOREMS AND ANOTHER APPLICATION

The objective of this section is to enrich our analysis by requiring rational behavior on and off the equilibrium path. The main message is that Theorem 1 generalizes to stronger solution concepts. All we need is a revelation principle for these solution concepts. While the definitions of these revelation principles are rather complex, they share some salient features. First, these revelation principles require players to be obedient and truthful at other histories than the on-path histories, but not at all histories. A player is required to be obedient and truthful only if he has not lied in the past. That player may have disobeyed past recommendations, however. In addition, not even all these histories are considered. A further consistency requirement is imposed, the so-called consistency with mediation ranges. Second, these revelation principles postulate that players assign probability zero to the event that others have lied to the mediator. As Myerson (1986, p. 342) put it:

This begs the question of whether we could get a larger set of sequentially rational communication equilibria if we allowed players to assign positive probability to the event that others have lied to the mediator. Fortunately, by the revelation principle, this set would not be any larger. (Myerson, 1986, p. 342)

In what follows, we do not provide the reader with a restatement of these revelation principles—this would take too much space. We refer to Myerson (1986) and the recent work of Sugaya and Wolitzky (2021). We start with the concept of weak perfect Bayesian equilibrium, one of the most widely used solution concepts in applications and, probably, the easiest refinement to state.

5.1. Weak Perfect Bayesian Equilibrium. Throughout, we fix an expansion Γ_π of Γ . We denote $\bar{\mathbb{P}}_{\sigma,\pi}(\cdot|h^t, m^t, \omega^t)$ 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) . The distribution $\bar{\mathbb{P}}_{\sigma, \pi}(\cdot | h^t, m^t, \omega^t)$ 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 distribution 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 $\bar{\mathbb{P}}_{\sigma, \pi}(\mathbf{h}, \mathbf{m}, \boldsymbol{\omega}) := \bar{\mathbb{P}}_{\sigma, \pi}(\mathbf{h}, \mathbf{m}, \boldsymbol{\omega} | h^0, m^0, \omega^0)$. At any given history (h^t, m^t, ω^t) , player i 's expected payoff is

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

To complete the description, we need to specify the belief player i has at any private history (h_i^t, m_i^t) . To do so, we specify a belief system β . Player i believes that the history is (h^t, m^t, ω^t) with probability $\beta(h^t, m^t, \omega^t | h_i^t, m_i^t)$ at the private history (h_i^t, m_i^t) . At the private history (h_i^t, m_i^t) , player i 's expected payoff is therefore:

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

Definition 3 (wPBE): A profile σ of behavioral strategies is a *weak perfect Bayesian equilibrium* of Γ_π if there exists a belief system β on $HM\Omega$ such that:

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

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

for all σ'_i .

- (ii) *Belief consistency:* The belief system β is consistent with σ , that is, for all $(h, m, \omega) \in HM\Omega$, for all (i, t) ,

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

whenever $\mathbb{P}_{\sigma, \pi}(h_i^t, m_i^t) > 0$.

We let $w\mathcal{PBE}(\Gamma_\pi)$ be the set of distributions over $H\Omega$ induced by the weak perfect Bayes equilibria of Γ_π .

As before, the objective is to characterize the set $\bigcup_{\Gamma_\pi \text{ an expansion of } \Gamma} w\mathcal{PBE}(\Gamma_\pi)$, i.e., we want to characterize the distributions over the outcomes $H\Omega$ of the base game Γ that we can induce by means of some expansion Γ_π of the base game, *without* any reference to particular expansions. To do so, we need to introduce the concept of weak perfect Bayes correlated equilibrium of Γ .

Weak Perfect 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 signals s_i^t , past actions a_i^{t-1} and past recommendations \hat{a}_i^{t-1} , $R_{i,t}(h_i^t, \hat{a}_i^{t-1}) \subseteq A_{i,t}$ is the set of possible recommendations to player i . We refer to the function $R_{i,t}$ as the *mediation range* of player i at stage t . We denote $\mathcal{H}(R)$ the set of all terminal histories consistent with the mediation ranges in the mediated extension $\mathcal{M}(\Gamma)$, i.e., $(h, \omega, \hat{a}) \in \mathcal{H}(R)$ if and only if $(h, \omega) \in H\Omega$ and $\hat{a}_{i,t} \in R_{i,t}(h_i^t, \hat{a}_i^{t-1})$ for all i , for all t .

We denote $\bar{\mathbb{P}}_{\mu \circ \tau, p}(\cdot | h^t, \omega^t, \hat{a}^t)$ the distribution over $\mathcal{H}(R)$ induced by the profile of strategies τ , the recommendation kernels μ and the kernels p , given the history $(h^t, \omega^t, \hat{a}^t)$. At any history $(h^t, \omega^t, \hat{a}^t)$, player i 's expected payoff is

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

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

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

where β is a belief system. 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 . We are now ready to define the concept of weak perfect Bayes correlated equilibrium.

Definition 4 (wPBCE): A weak perfect Bayes correlated equilibrium of Γ is a collection of mediation ranges $R_{i,t} : H_i^t \times A_i^{t-1} \rightarrow 2^{A_{i,t}} \setminus \{\emptyset\}$ for all (i, t) , a collection of recommendation kernels $\mu_t(h^t, \omega^t, \hat{a}^{t-1}) : \times_{i \in I} R_{i,t}(h_i^t, \hat{a}_i^{t-1}) \rightarrow [0, 1]$, where $\sum_{\hat{a}_t \in \times_{i \in I} R_{i,t}(h_i^t, \hat{a}_i^{t-1})} \mu_t(h^t, \omega^t, \hat{a}^{t-1})[\hat{a}_t] = 1$, for all $(h^t, \omega^t, \hat{a}^{t-1})$ in $\mathcal{H}(R)$ and a belief system β such that:

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

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

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

- (ii) *Belief consistency:* β is consistent with (τ^*, μ, p) , that is, for all $(h, \omega, \hat{a}) \in \mathcal{H}(R)$, for all (i, t) ,

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

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

We let $wPBCE(\Gamma)$ be the set of distributions over $H\Omega$ induced by the weak perfect Bayes correlated equilibria of Γ .

It is worth pausing over the role of the mediation ranges. A weak perfect Bayes correlated equilibrium constrains the mediator to only recommend actions consistent with the mediation ranges, that is, the only recommendations the mediator can make to player i are in $R_{i,t}(h_i^t, \hat{a}_i^{t-1})$ at history $(h^t, \omega^t, \hat{a}^{t-1})$.¹⁰ In addition, players must have an incentive to be obedient at all histories consistent with the mediation ranges. The role of mediation ranges is precisely to ensure that players can be obedient at all histories of the mediated game. Without constraining the recommendations the mediator can make, it wouldn't be possible to ensure that players are obedient at all histories. E.g., no player would ever have an incentive to play a strictly dominated action. An equivalent formulation is to consider weak perfect

¹⁰This explains why the domain of $\mu_t(h^t, \omega^t, \hat{a}^{t-1})$ is $\times_{i \in I} R_{i,t}(h_i^t, \hat{a}_i^{t-1})$ in our definition.

Bayesian equilibria of the mediated game $\mathcal{M}(\Gamma)$, where the mediator is omniscient and unconstrained in its recommendations, and players are obedient on path. The drawback of this alternative formulation is that players do not have to be obedient off path and, therefore, requires to explore all possible behaviors off path. The advantage is that no mediation ranges are required.

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

Theorem 3. *We have the following equivalence:*

$$w\mathcal{PBC}\mathcal{E}(\Gamma) = \bigcup_{\Gamma_\pi \text{ an expansion of } \Gamma} w\mathcal{PBE}(\Gamma_\pi).$$

Theorem 3 states an equivalence between (i) the set of distributions over actions, base signals and states induced by all weak perfect Bayes correlated equilibria of Γ , and (ii) the set of distributions over actions, base signals and states we can obtain by considering *all* weak perfect Bayesian equilibria of *all* expansions of Γ .

The logic behind Theorem 3 is identical to the the one behind Theorem 1. We can replicate any weak perfect Bayesian equilibrium of Γ_π as a weak perfect Bayesian equilibrium of the auxiliary mediated game $\mathcal{M}^*(\Gamma)$ and then invoke the revelation principle for weak perfect Bayesian equilibria, which was recently proved by Sugaya and Wolitzky (2018, Proposition 2). More precisely, their revelation principle states that it is without loss of generality to assume that players report their private information to the mediator, that the mediator recommends actions to the players, and that players have an incentive to be truthful and obedient provided they have been truthful in the past. The mediator cannot recommend actions outside the mediation ranges and a belief system gives the player's beliefs. Thus, unlike the classical revelation principle of Forges (1986) and Myerson (1986), players are

required to continue to be truthful and obedient even if they have disobeyed in the past, so long as they have been truthful in the past.

We conclude with few additional remarks. First, the set $wPBC\mathcal{E}(\Gamma)$ is convex.¹¹ Second, despite its theoretical shortcomings, we have considered the concept of weak perfect Bayesian equilibrium as our solution concept.¹² We did so for two two main reasons. First, it is simple and indeed widely used in applications. Second, it generalizes to continuous games, a common assumption in applications. In what follows, we present another solution concept, which alleviates some of the theoretical shortcomings of weak perfect Bayesian equilibrium. However, it comes at a cost: it is “harder” to state and to use in applications.

5.2. Conditional Probability Perfect Bayesian Equilibrium. 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 β 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) $\beta(Z|Z) = 1$ and $\beta(\mathcal{X}|Z) = 1$,
- (ii) if $X \cap Y = \emptyset$, then $\beta(X \cup Y|Z) = \beta(X|Z) + \beta(Y|Z)$,
- (iii) if $X \subseteq Y \subseteq Z$ and $Y \neq \emptyset$, then $\beta(X|Z) = \beta(X|Y)\beta(Y|Z)$.

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

¹¹See the working paper version for a proof.

¹²It is well-known that weak perfect Bayesian equilibria may not be subgame perfect, may rely on “irrational” beliefs, and may not satisfy the one-shot deviation principle.

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.¹³

We now define the concept of conditional probability perfect Bayesian equilibrium, a concept introduced by Sugaya and Wolitzky (2021). We first give an informal definition. A conditional probability perfect Bayesian equilibrium is a profile of strategies and a conditional probability system such that (i) sequential rationality holds given the belief system induced by the conditional probability system and (ii) the conditional probability system is consistent with the profile of strategies and the data of the game. It is a stronger concept than the concept of weak perfect Bayesian equilibrium and a weaker concept than the concept of sequential equilibrium. We now turn to a formal definition.

In what follows, we use notation, which parallel the one used in previous definitions, and thus do not rehash formal definitions.

Definition 5 (CPPBE): A conditional probability perfect Bayesian equilibrium of Γ_π is a profile σ of behavioral strategies and a CPS β on $HM\Omega$, which satisfy:

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

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

for all σ'_i .

¹³Myerson shows that for any conditional probability system β , 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) $\beta = \lim_n \mathbb{P}^n$, that is, $\beta(X|Y) = \lim_n \frac{\mathbb{P}^n(X \cap Y)}{\mathbb{P}^n(Y)}$ for all X , for all $Y \neq \emptyset$.

(ii) *CPS consistency*: The CPS β is consistent with (σ, p, ξ) , that is, for all $(h, m, \omega) \in HM\Omega$, for all (i, t) ,

$$\begin{aligned}\beta(a_t|h^t, m^t, \omega^t) &= \prod_{i \in I} \sigma_{i,t}(a_{i,t}|h_i^t, m_i^t), \\ \beta(h_{t+1}, \omega_{t+1}|a_t, h^t, m^t, \omega^t) &= p_{t+1}(h_{t+1}, \omega_{t+1}|a_t, h^t, \omega^t), \\ \beta(m_{t+1}|h^{t+1}, m^t, \omega^{t+1}) &= \xi_{t+1}(m_{t+1}|h^{t+1}, m^t, \omega^{t+1}).\end{aligned}$$

We let $CPPE(\Gamma_\pi)$ be the set of distributions over $H\Omega$ induced by the conditional probability perfect Bayesian equilibria of Γ_π .

Few comments are worth making. First, to ease notation, we have written $\beta(a_t|h^t, m^t, \omega^t)$ for

$$\beta\left(\left\{(\mathbf{h}, \mathbf{m}, \boldsymbol{\omega}) \in HM\Omega : (\mathbf{a}_t, \mathbf{h}^t, \mathbf{m}^t, \boldsymbol{\omega}^t) = (a_t, h^t, m^t, \omega^t)\right\} \mid \left\{(\mathbf{h}, \mathbf{m}, \boldsymbol{\omega}) \in HM\Omega : (\mathbf{h}^t, \mathbf{m}^t, \boldsymbol{\omega}^t) = (h^t, m^t, \omega^t)\right\}\right).$$

We use similar abuse of notation throughout. Second, the consistency of the CPS implies that

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

whenever $\mathbb{P}_{\sigma, \pi}(h_i^t, m_i^t) > 0$. Third, a conditional probability perfect Bayesian equilibrium is subgame perfect. Fourth, since the belief a player has is induced by the CPS, two players with the same information have the same belief. However, the CPS does not impose a “don’t signal what you don’t know” condition. To do so, we would need to require the CPS to maintain the relative likelihood of any two histories before and after players taking actions.

As before, the objective is to characterize the set $\bigcup_{\Gamma_\pi \text{ an expansion of } \Gamma} CPPE(\Gamma_\pi)$, i.e., we want to characterize the distributions over the outcomes $H\Omega$ of the base game Γ that we can induce by means of some expansion Γ_π of the base game, *without* any reference to particular expansions. To do so, we need to introduce the concept of sequential Bayes correlated equilibrium of Γ .

Sequential Bayes correlated equilibrium. As in the previous section, 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. We use the same notation and do not rehash them.

A feedback rule $f := (f_1, \dots, f_T)$ is a deterministic recommendation kernel, which recommends the action $f_t(h^t, \omega^t)$ at history $(h^t, \omega^t) \in H^t \Omega^t$. Note that given f , the history (h^t, ω^t) encodes the profile of recommendations \hat{a}^t as $(f_1(h^1, \omega^1), f_2(h^2, \omega^2), \dots, f_t(h^t, \omega^t))$. A feedback rule f is consistent with the mediation ranges R if $f_{i,t}(h^t, \omega^t) \in R_{i,t}(h_i^t, \hat{a}_i^{t-1})$ for all i , for all (h^t, ω^t) , for all t , where \hat{a}^{t-1} is the profile of recommendations encoded by f at (h^{t-1}, ω^{t-1}) . We let \mathcal{F} be the set of feedback rules and $\mathcal{F}(R)$ the subset of feedback rules consistent with the mediation ranges R .

We denote $\bar{\mathbb{P}}_{f \circ \tau, p}(\cdot | h^t, \omega^t)$ the distribution over $\mathcal{H}(R)$ induced by the profile of strategies $\bar{\tau}$, the feedback rule f and the kernels p , given the history (h^t, ω^t) . At any history (h^t, ω^t) , player i 's expected payoff is

$$U_i(f \circ \tau | h^t, \omega^t) := \sum_{\mathbf{h}, \boldsymbol{\omega}} u_i(\mathbf{h}, \boldsymbol{\omega}) \bar{\mathbb{P}}_{f \circ \tau, p}(\mathbf{h}, \boldsymbol{\omega} | h^t, \omega^t),$$

when the feedback rule is f . Finally, at any private history (h_i^t, \hat{a}_i^t) , player i 's expected payoff is:

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

where β is a CPS on $\mathcal{F}(R) \times H\Omega$. Here, we write $\beta(f, h^t, \omega^t | h_i^t, \hat{a}_i^t)$ for:

$$\beta\left(\left\{(\mathbf{f}, \mathbf{h}, \boldsymbol{\omega}) : (\mathbf{f}, \mathbf{h}^t, \boldsymbol{\omega}^t = f, h^t, \omega^t)\right\} \middle| \left\{(\mathbf{f}, \mathbf{h}, \boldsymbol{\omega}) : (\mathbf{f}_{i,1}(h^1, \omega^1), \dots, \mathbf{f}_{i,t}(h^t, \omega^t)) = \hat{a}_i^t, \mathbf{h}_i^t = h_i^t\right\}\right)$$

Definition 6 (SBCE): A communication mechanism $\mu \in \Delta(\mathcal{F})$ is a sequential Bayes correlated equilibrium of Γ if there exist mediation ranges R and a conditional probability system β on $\mathcal{F}(R) \times H\Omega$ such that:

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

$$U_i(\tau^*, \beta | h_i^t, a_i^t) \geq U_i((\tau_i, \tau_{-i}^*), \beta | h_i^t, a_i^t)$$

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

- (ii) *CPS consistency*: For all f, h, ω, t ,

$$\beta(f, h, \omega) = \mu(f) \bar{\mathbb{P}}_{f \circ \tau^*, p}(h, \omega)$$

$$\beta(f, h, \omega | (f_1, \dots, f_t), (h^t, \omega^t)) = \beta(f | (f_1, \dots, f_t), (h^t, \omega^t)) \bar{\mathbb{P}}_{f \circ \tau^*, p}(h, \omega | h^t, \omega^t).$$

We let $SBC\mathcal{E}(\Gamma)$ be the set of distributions over $H\Omega$ induced by the sequential Bayes correlated equilibria of Γ . The set $SBC\mathcal{E}(\Gamma)$ is convex.

Few remarks are worth making. First, in a sequential Bayes correlated equilibrium, players have an incentive to be obedient at all histories consistent with the mediation ranges. Second, unlike previous definitions, the definition asserts that the omniscient mediator selects a feedback rule f with probability μ , i.e., as if the mediator chooses a mixed strategy (and not a behavioral strategy). In addition, the conditional probability system is required to be consistent with μ . Third, we may wonder whether an equivalent formulation exists where the mediator chooses recommendation kernels $(\mu_t)_t$ (behavioral strategies) and consistency is imposed with respect to $(\mu_t)_t$, as we did in the definition of a weak perfect Bayes correlated equilibrium. As Sugaya and Wolitzky (2021) show, the answer is unfortunately no. Intuitively, the current formulation allows more flexibility in choosing beliefs, which is needed for a revelation principle to hold. Lastly, sequential Bayes correlated equilibria are sequential *communication* equilibria (Myerson, 1986) of mediated games, where the mediator is omniscient.¹⁴ Readers should not confuse these concepts with that of sequential equilibrium, which does not involve a device such as a mediator.

¹⁴Sequential Bayes correlated equilibria are the subsets of Bayes correlated equilibria, where the mediator never recommends co-dominated actions, a generalization of the concept of dominance. We refer the reader to Myerson (1986) for more detail.

Theorem 4. *We have the following equivalence:*

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

Theorem 4 states an equivalence between (i) the set of distributions over actions, base signals and states induced by all sequential Bayes correlated equilibria of Γ , and (ii) the set of distributions over actions, base signals and states we can obtain by considering *all* conditional probability perfect Bayesian equilibria of *all* expansions of Γ . The logic behind Theorem 4 and its proof are the same as in previous sections. (More precisely, the revelation principle we invoke is stated in Proposition 8 in Sugaya and Wolitzky (2021).)

5.3. Application II: 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).¹⁵ Unlike these papers, we consider *all* 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 characterize the set of *sequential Bayes correlated equilibria*. A communication system μ is a sequential Bayes correlated equilibrium if there exist mediation ranges (R_1, R_2) and a conditional probability system β , which jointly satisfy the following constraints. First, if the omniscient mediator recommends $f_1(\omega) \in R_1$ to the seller, the seller must have an incentive to be obedient, i.e.,

$$\sum_{f, \omega} f_1(\omega) f_2(f_1(\omega), \omega) \beta(f, \omega | f_1(\omega)) \geq \sum_{f, \omega} a_1 f_2(a_1, \omega) \beta(f, \omega | f_1(\omega))$$

for all a_1 . Second, if the offer made to the buyer is a_1 and the mediator recommends $f_2(a_1, \omega) \in R_2(a_1)$ to the buyer, the buyer must have an incentive to be obedient, i.e.,

$$\sum_{f, \omega} (\omega - a_1) f_2(a_1, \omega) \beta(f, \omega | a_1, f_2(a_1, \omega)) \geq \sum_{f, \omega} (\omega - a_1) (1 - f_2(a_1, \omega)) \beta(f, \omega | a_1, f_2(a_1, \omega)).$$

Third, the conditional probability system must be consistent, that is, for all $f \in \mathcal{F}(R)$, for all a_1, a_2, ω ,

$$\begin{aligned} \beta(f, a_1, a_2, \omega) &= \mu(f) p(\omega) \mathbb{1}\{f_1(\omega) = a_1, f_2(f_1(\omega), \omega) = a_2\}, \\ \beta(f, a_1, a_2, \omega | f_1, a_1, \omega) &= \beta(f | f_1, a_1, \omega) \mathbb{1}\{(f_2(a_1, \omega), \omega) = a_2\}. \end{aligned}$$

There are immediate bounds on the equilibrium payoffs: the sum of the buyer and seller's payoffs is bounded from above by $\mathbb{E}(\omega) = \sum_{\omega} p(\omega) \omega$, the

¹⁵In related work, Kartik and Zhang (2019) characterize all wPBE of the bilateral bargaining model with interdependent values as one varies the information structure. They do so under three different scenarios regarding base signals: one like ours, one like that in Bergemann, Brooks and Morris, and one where the buyer is better informed than the seller. Under the first scenario, they also find that the set of implementable payoffs is the one that satisfies the immediate bounds on equilibrium payoffs we discuss prior to Proposition 1.

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*

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

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




FIGURE 3. Payoffs at all sequential Bayes correlated equilibria

Notice that the set of BCE here differs from the one in Bergemann, Brooks and Morris (2013) in that the lowest seller payoff is ω_L^- , and not the monopoly profit.

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 $\sum_{f,\omega} \mu(f)p(\omega)\mathbb{1}\{f_1(\omega) = a_1\} = 0$. To see this, for any conditional probability system, $\beta(a_1, \omega) = \beta(\omega, a_1|a_1)\beta(a_1)$. Moreover, from the consistency of β , we have that $\beta(a_1, \omega) = \sum_f \beta(f, a_1, \omega) = \sum_f \mu(f)p(\omega)\mathbb{1}\{f_1(\omega) = a_1\} = 0$. Since $\beta(a_1) = 0$, $\beta(\omega, a_1|a_1)$ is arbitrary and, thus, 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 $f_2(a_1, \omega)$.

We are now ready to state how to obtain the payoff profile $(\mathbb{E}(\omega) - \omega_L^-, \omega_L^-)$. We first start with an informal description. 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. If any offer $a_1 > \omega_L^-$ is made, the mediator recommends the buyer to reject the offer, regardless of the state. 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. Formally, the communication system puts probability one to f , given by $f_1(\omega) = \omega_L^-$, $f_2(a_1, \omega) = 0$ if $a_1 > \omega_L^-$ and $f_2(a_1, \omega) = 1$ if $a_1 \leq \omega_L^-$ for all ω . 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 $a_1 > \omega_L$.

We now turn our attention to the two other payoff profiles $(0, \mathbb{E}(\omega))$ and $(0, \omega_L^-)$. The profile $(0, \mathbb{E}(\omega))$ corresponds to full surplus extraction, which can be obtained with $f_1(\omega) = \omega$ for all ω and $f_2(a_1, \omega) = 1$ whenever $a_1 \leq \omega$ (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, when $\mathbb{E}(\omega) \in A_1$ (which we assume), the profile $(0, \omega_L^-)$ is implementable as follows. Consider two feedback rules f and f' such that for all ω , $f_1(\omega) = f'_1(\omega) = \mathbb{E}(\omega)$, $f_2(a_1, \omega) = f'_2(a_1, \omega) = 0$ if $a_1 > \mathbb{E}(\omega)$, $f_2(a_1, \omega) = f'_2(a_1, \omega) = 1$ if $a_1 < \mathbb{E}(\omega)$, $f_2(\mathbb{E}(\omega), \omega) = 1$ while $f'_2(\mathbb{E}(\omega), \omega) = 0$. Assume that $\mu(f) = \omega_L^-/\mathbb{E}(\omega)$, $\mu(f') = 1 - \mu(f)$, and that $R_1 = \{\mathbb{E}(\omega)\}$, $R_2(a_1) = \{1\}$ if $a_1 < \omega_L$, $R_2(a_1) = \{0, 1\}$ if $a_1 = \mathbb{E}(\omega)$, and $R_2(a_1) = \{0\}$, otherwise. In effect, 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. 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.

6. CONCLUSION

This paper generalizes the concept of Bayes correlated equilibrium to multi-stage games and offers two applications, which are suggestive of the usefulness of our characterization results. The main contribution is methodological.

The reader may wonder why we have not considered the concept of sequential equilibrium. The main reason is that a revelation principle does not hold for this concept. To be more precise, Sugaya and Wolitzky (2021) show that the set of sequential communication equilibria of a multi-stage game characterizes the set of equilibrium distributions we can obtain by considering all mediated extensions of the multi-stage game, where the solution concept is sequential equilibrium. However, their definition of a sequential equilibrium treats the mediator as a player and, thus, allows for the mediator to tremble. When we consider an expansion and its emulation by a mediator with the mediated game $\mathcal{M}^*(\Gamma)$, players do not expect the mediator to tremble. If a player observes an unexpected additional signal, that player must believe with probability one that one of his opponents has deviated. He cannot believe that none of his opponents deviated, but the mediator did. This would be inconsistent with the expansion being the game actually played. Extending the analysis to other solution concepts such as sequential equilibrium or rationalizability or to general extensive-form games is challenging and left for future research.

APPENDICES

A. Proof of Theorem 1. (\Leftarrow .) We first prove that $\bigcup_{\Gamma_\pi \text{ an expansion of } \Gamma} \mathcal{BN}\mathcal{E}(\Gamma_\pi) \subseteq \mathcal{BC}\mathcal{E}(\Gamma)$. Throughout, we fix an expansion Γ_π of Γ . Recall that 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 .

Let σ^* be a Bayes-Nash equilibrium of Γ_π . We now construct an auxiliary mediated game $\mathcal{M}^*(\Gamma)$, which emulates the distribution $\mathbb{P}_{\sigma^*, \pi}$ as an equilibrium distribution.

The game $\mathcal{M}^*(\Gamma)$ has one additional player, labelled player 0, and a (Forges-Myerson) mediator. Player 0 is a dummy player: his payoff is identically zero.

The game unfolds as follows: At stage $t = 1$,

- Nature draws (h_1, ω_1) with probability $p_1(h_1, \omega_1)$.
- Player $i \in I$ observes the signal $h_{i,1}$ and player 0 observes (h_1, ω_1) .
- Player 0 reports $(\hat{h}_1, \hat{\omega}_1)$ to the mediator. All other players do not make reports.
- The mediator draws the message m_1 with probability $\xi_1(m_1 | \hat{h}_1, \hat{\omega}_1)$ and sends the message $m_{i,1}$ to player i . Player 0 does not receive a message.
- Player i takes an action $a_{i,1}$. Player 0 does not take an action.

Consider now a history $(a_{t-1}, h^{t-1}, \omega^{t-1})$ of past actions, signals and states and a history $((\hat{h}^{t-1}, \hat{\omega}^{t-1}), m^{t-1})$ of reports and messages. At stage t :

- Nature draws (h_t, ω_t) with probability $p_t(h_t, \omega_t | a_{t-1}, h^{t-1}, \omega^{t-1})$.
- Player $i \in I$ observes the signal $h_{i,t}$ and player 0 observes (h_t, ω_t) .
- Player 0 reports $(\hat{h}_t, \hat{\omega}_t)$ to the mediator. All other players do not make reports.
- The mediator draws the message m_t with probability $\xi_t(m_t | \hat{h}^t, m^{t-1}, \hat{\omega}^t)$ and sends the message $m_{i,t}$ to player i . Player 0 does not receive a message.
- Player i takes an action $a_{i,t}$. Player 0 does not take an action.

In the above description, when we say that player i does not make a report, we implicitly assume that the set of reports player i can make to the mediator is a singleton. Similarly, when we say that player 0 does not take an action. In the rest of the proof, we omit these trivial reports and actions.

We restrict attention to the histories of $\mathcal{M}^*(\Gamma)$, where $(h, \omega) \in H\Omega$. At stage t , player i 's private history is (h_i^t, m_i^t) , which is also player i 's private history in Γ_π . In addition, any private history $(h_{i,t}, m_{i,t})$ in $\mathcal{M}^*(\Gamma)$ is also a private history in Γ_π . Thus, σ_i^* is a well-defined strategy for player i in $\mathcal{M}^*(\Gamma)$. Moreover, if player 0 truthfully reports his private information (h_t, ω_t) at all histories $((h_t, \omega_t), (h^{t-1}, \omega^{t-1}), (\hat{h}^{t-1}, \hat{\omega}^{t-1}))$, the conditional probability of the message m_t is the same as in Γ_π . It follows immediately that σ^* together with the truthful strategy for player 0 is a Bayes-Nash equilibrium of the auxiliary mediated game $\mathcal{M}^*(\Gamma)$.

From the revelation principle of Forges (1986) and Myerson (1986), there exists a canonical equilibrium μ , where the mediator recommends actions and players are truthful and obedient, provided they have been in the past. At truthful histories, the mediator recommends \hat{a}_t with probability

$$\mu_t(\hat{a}_t | \underbrace{(h^t, \omega^t)}_{\text{player 0}}, \underbrace{(h_1^t, \dots, h_n^t)}_{\text{players in } I}, \underbrace{\hat{a}^{t-1}}_{\text{past recommendations}}).$$

It is then routine to verify that we have a Bayes-correlated equilibrium with the recommendation kernel μ_t given by

$$\mu_t(\hat{a}_t | h^t, \omega^t, \hat{a}^{t-1}) := \mu_t(\hat{a}_t | (h^t, \omega^t), (h_1^t, \dots, h_n^t), \hat{a}^{t-1}),$$

for all $(h^t, \omega^t, \hat{a}^{t-1})$ for all t .

(\Rightarrow). We now prove that $\mathcal{BC}\mathcal{E}(\Gamma) \subseteq \bigcup_{\Gamma_\pi} \text{an expansion of } \Gamma \mathcal{BN}\mathcal{E}(\Gamma_\pi)$.

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

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

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

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

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

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 well-defined: $\xi_1(m_1|h_1, \omega_1) = \mu_1(\hat{a}_1|h_1, \omega_1)$ with $m_1 = \hat{a}_1$, and, for $t > 1$, $\xi_{t+1}(m_{t+1}|h^{t+1}, m^t, \omega^{t+1}) = \mu_{t+1}(\hat{a}_{t+1}|h^{t+1}, \omega^{t+1}, \hat{a}^t)$ with $(m^t, m_{t+1}) = (\hat{a}^t, \hat{a}_{t+1})$.

By construction, any strategy $\tau_t : H^t \times A^t \rightarrow \Delta(A_t)$ of $\mathcal{M}(\Gamma)$ is equivalent to a strategy $\sigma_t : H^t \times M^t \rightarrow \Delta(A_t)$ of Γ_π , i.e., $\sigma_t(a_t|h^t, m^t) := \times_i \sigma_{i,t}(a_{i,t}|h_i^t, m_i^t) = \times_i \tau_{i,t}(a_{i,t}|h_i^t, \hat{a}_i^t)$ with $m^t = \hat{a}^t$, with the property that $\mathbb{P}_{\sigma,\pi}(h^t, m^t, \omega^t) = \mathbb{P}_{\mu \circ \tau, p}(h^t, \hat{a}^t, \omega^t)$ 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 $\mathbb{P}_{\sigma,\pi}(h_1, m_1, \omega_1) = \mathbb{P}_{\mu \circ \tau, p}(h_1, \omega_1, \hat{a}_1)$ with $m_1 = \hat{a}_1$, for all (h_1, m_1, ω_1) . By induction, assume that $\mathbb{P}_{\sigma,\pi}(h^t, m^t, \omega^t) = \mathbb{P}_{\mu \circ \tau, p}(h^t, \omega^t, \hat{a}^t)$ 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

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

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

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

This completes the proof.

B. Proof of Theorem 3. The proof is nearly identical to the proof of Theorem 1 and is, therefore, omitted. We only sketch the minor differences.

(\Leftarrow .) Fix an expansion Γ_π and a weak perfect Bayesian equilibrium (σ^*, β) of Γ_π . We need to construct a weak perfect Bayesian equilibrium of the auxiliary game $\mathcal{M}^*(\Gamma)$, which replicates the distribution $\mathbb{P}_{\sigma, \pi}$. To do so, we define a belief system β^* of the auxiliary game $\mathcal{M}^*(\Gamma)$ as follows:

$$\beta^*(h^t, m^t, \omega^t, (h^t, \omega^t), (h^t, \omega^t) | h_i^t, m_i^t) := \beta(h^t, m^t, \omega^t | h_i^t, m_i^t)$$

for all (h^t, m^t, ω^t) , for all (i, t) . Note that the above formulation implies that player i believes with probability 1 that player 0 truthfully report (h^t, ω^t) . (In $\mathcal{M}^*(\Gamma)$, player i also has beliefs about the signals (h^t, ω^t) player 0 receives and the reports $(\hat{h}^t, \hat{\omega}^t)$ by player 0 to the mediator.) It is immediate to verify that $(\sigma_0^*, \sigma^*, \beta^*)$ is a weak perfect Bayesian equilibrium of $\mathcal{M}^*(\Gamma)$, where σ_0^* is the truthful reporting strategy of player 0. The proof then follows from the revelation principle for weak perfect Bayesian equilibrium. Three remarks are worth making. First, the mediation ranges and the belief system come from the revelation principle—the revelation principle precisely states the existence of mediation ranges and belief system such that players have an incentive to be obedient and truthful, provided they have been truthful in the past. Second, since $\mathcal{M}^*(\Gamma)$ is a mediated extension of Γ , it is a common belief that states and base signals evolve according to $(p_t)_t$. At the truthful histories of the direct mediated extension of Γ , players have an incentive to be obedient at all recommendations consistent with the mediation ranges, the mediator is omniscient and beliefs are as in the base game Γ . (Beliefs of player i about the action of player 0 are trivial—player 0 has no actions.) Third, the revelation principle asserts that players assign probability zero to the event others have lied to the mediator, hence we can use the belief $\bar{\beta}$ such that:

$$\bar{\beta}(h^t, \omega^t, \hat{a}^t | h_i^t, \hat{a}_i^t) := \beta^{**}(h^t, \omega^t, \hat{a}^t, (h_j^t)_j, (h^t, \omega^t) | h_i^t, h_i^t, \hat{a}_i^t).$$

for all $h^t, \omega^t, \hat{a}^t, t$, to sustain obedience, where β^{**} is the belief system inherited from the revelation principle.

We therefore have a weak perfect Bayes correlated equilibrium.

(\Rightarrow .) We construct the expansions as in the the proof of Theorem 1, i.e., defining the additional signals as the recommendations. Since the additional signals player i can receive are the recommendations, player i can only receive additional signals consistent with the mediation ranges. Thus, we can use the belief system of the weak perfect Bayes correlated equilibrium to construct the weak perfect Bayesian equilibrium of Γ_π .

C. Proof of Theorem 4. The proof is yet again nearly identical to the proof of Theorem 1. We only sketch the main differences.

(\Leftarrow .) Fix an expansion Γ_π and a conditional probability perfect Bayesian equilibrium (σ^*, β) of Γ_π . As in the previous proofs, we construct a conditional probability perfect Bayesian equilibrium of the mediated game $\mathcal{M}^*(\Gamma)$, which replicates the distribution $\mathbb{P}_{\sigma^*, \pi}$. As in the proof of Theorem 3, we construct a conditional probability system β^* of the mediated game $\mathcal{M}^*(\Gamma)$ from the conditional probability system β of the game Γ_π such that $((\sigma_0^*, \sigma^*), \beta^*)$ is a conditional probability perfect Bayesian equilibrium of $\mathcal{M}^*(\Gamma)$, with player 0, the dummy player, truthfully reporting his private information (h^t, ω^t) at each stage t . Since β is a conditional probability system, there exists a sequence β^n of fully supported probabilities such that $\beta(X|Y) = \lim_n \frac{\beta^n(X \cap Y)}{\beta^n(Y)}$ for all X and all non-empty Y . Consider now the sequence of fully supported kernels $\gamma^n : HM\Omega \rightarrow \Delta(H\Omega \times H\Omega)$, where γ^n converges to $\gamma((h, \omega), (h, \omega)|(h, m, \omega)) = 1$ for all (h, m, ω) . The interpretation is that player 0 learns and truthfully report (h, ω) , when the profile of actions, signals, and states is (h, m, ω) . Let β^* be the CPS resulting from taking the limit of $\beta^n \times \gamma^n$. By construction, $\beta^*(h^t, m^t, \omega^t|h_i^t, m_i^t) = \beta(h^t, m^t, \omega^t|h_i^t, m_i^t)$, so that σ_i^* remains sequentially rational for player i . Moreover, the newly constructed conditional probability system is consistent with the kernels p and ξ . The rest of the proof follows from the revelation principle.

(\Rightarrow). Let (μ, R, β) be a sequential Bayes correlated equilibrium. The difference with the previous proofs is that the definition of a sequential Bayes correlated equilibrium does not specify recommendation kernels (μ_1, \dots, μ_T) , which can then be used as expansions. However, as in the proof of Kuhn's theorem, we can construct such recommendation kernels from μ .

The construction is iterative. In the sequel, we slightly abuse notation and write (μ_1, \dots, μ_T) for the kernels. For all $(h^1, \omega^1, \hat{a}^0)$ such that $p_1(h_1, \omega_1) > 0$,

$$\mu_1(\hat{a}_1 | \hat{a}^0, h^1, \omega^1) := \sum_f \mu(f) \mathbb{1}\{f_1(h^1, \omega^1) = \hat{a}_1\}.$$

Note that $\mu_1(\hat{a}_1 | h^1, \omega^1, \hat{a}^0) > 0$ and $\sum_{\hat{a}_1} \mu_1(\hat{a}_1 | h^1, \omega^1, \hat{a}^0) = 1$. The kernel μ_1 is thus well-defined. (Recall that $(h^1, \omega^1) = (h_1, \omega_1)$ and that \hat{a}^0 is a singleton.)

We proceed iteratively. For all $(h^t, \omega^t, \hat{a}^{t-1})$ such that

$$p_1(h_1, \omega_1) \mu_1(\hat{a}_1 | \hat{a}^0, h^1, \omega^1) \times \dots \times p_{t-1}(h_{t-1}, \omega_{t-1} | a_{t-1}, h^{t-2}, \omega^{t-2}) \mu_{t-1}(\hat{a}_{t-1} | \hat{a}^{t-2}, h^{t-1}, \omega^{t-1}) p_t(h_t, \omega_t | a_t, h^{t-1}, \omega^{t-1}) > 0$$

for some (a_1, \dots, a_t) ,

$$\mu_t(\hat{a}_t | \hat{a}^{t-1}, h^t, \omega^t) := \frac{\sum_f \mu(f) \mathbb{1}\{f_1(h^1, \omega^1) = \hat{a}_1, \dots, f_t(h^t, \omega^t) = \hat{a}_t\}}{\sum_f \mu(f) \mathbb{1}\{f_1(h^1, \omega^1) = \hat{a}_1, \dots, f_{t-1}(h^{t-1}, \omega^{t-1}) = \hat{a}_{t-1}\}}.$$

It is immediate to verify that the kernel is well-defined.

Two remarks are in order. First, since we consider histories $(h, \omega) \in H\Omega$, we already have that

$$p_1(h_1, \omega_1) \mu_1(\hat{a}_1 | \hat{a}^0, h^1, \omega^1) \times \dots \times p_{t-1}(h_{t-1}, \omega_{t-1} | a_{t-1}, h^{t-2}, \omega^{t-2}) p_t(h_t, \omega_t | a_t, h^{t-1}, \omega^{t-1}) > 0$$

for some (a_1, \dots, a_t) . Second, since we only consider feedback rules in the support of μ , all the recommendations with positive probabilities are consistent with the mediation ranges. Hence, players are obedient at these recommendations.

We define the conditional probability system on $H\Omega A$ as

$$\beta(h, \omega, \hat{a}) = \sum_f \beta(f, h, \omega) \mathbf{1}\{f(h, \omega) = \hat{a}\}.$$

To complete the proof, we repeat the same steps as in the proof of Theorem 1, that is, the additional messages are the recommendations, the kernels (ξ_1, \dots, ξ_T) are the recommendation kernels (μ_1, \dots, μ_T) , and the conditional probability system is the one on $H\Omega A$ defined above. Since we consider the restriction to recommendations with positive probabilities, the recommendations are consistent with the mediation ranges and players have an incentive to play according to their signals.

REFERENCES

- [1] Bergemann, Dirk and Stephen Morris, “Bayes Correlated Equilibrium and The Comparison of Information Structures in Games,” 2016, *Theoretical Economics*, 11, pp. 487-522.
- [2] Bergemann, Dirk, Benjamin Brooks and Stephen Morris, “The Limits of Price Discrimination,” 2015, *American Economic Review*, 105, pp. 921-957.
- [3] Bergemann, Dirk, Benjamin Brooks, and Stephen Morris, “Counterfactuals with Latent Information,” 2019, Cowles Foundation Discussion Paper 2162.
- [4] Calzolari, Giacomo and Alessandro Pavan, “Monopoly with Resale,” 2006, *Rand Journal of Economics*, 37, pp. 362-375.
- [5] de Oliveira, and Rohit Lamba, “Rationalizing Dynamic Choices,” 2019, mimeo.
- [6] Doval, Laura and Jeffrey C., Ely, “Sequential Information Design,” 2020, *Econometrica*, 86, pp. 2575-2608.
- [7] Dworczak, Piotr, “Mechanism Design with Aftermarkets: Cutoff Mechanisms,” 2017, mimeo.
- [8] Ely, Jeffrey C., “Beeps.” 2017, *American Economic Review*, 107, pp. 31-53.
- [9] Forges, Françoise, “Correlated Equilibria in a Class of Repeated Games with Incomplete Information,” 1985, *International Journal of Game Theory*, 14(3), pp. 29-150.
- [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] Gualdani, Christina and Sruti Sinha, “Identification in Discrete Choice Models with Imperfect Information,” 2021, mimeo.
- [14] Giovannoni, Francesco, and Miltiadis Makris, “Reputational Bidding,” 2014, *International Economic Review*, 55, pp. 693-710.
- [15] Forges, Françoise, “Correlated Equilibria in a Class of Repeated Games with Incomplete Information,” 1985, *International Journal of Game Theory*, 14(3), pp. 29-150.
- [16] Forges, Françoise, “An Approach to Communication Equilibria,” 1986, *Econometrica*, 54, pp. 1375-1385.
- [17] Forges, Françoise, “Five Legitimate Definitions of Correlated Equilibrium in Games with Incomplete Information,” 1993, *Theory and Decision*, 35, pp. 277-310.
- [18] Forges, Françoise, “Correlated Equilibrium in Games with Incomplete Information Revisited, 2006, *Theory and Decision*, 61, pp. 329-344.
- [19] Magnolfi, L., and C. Roncoroni, “Estimation of Discrete Games with Weak Assumptions on Information,” 2017, mimeo.
- [20] Mertens, Jean-François, and Sylvain Sorin and Shmuel Zamir, “Repeated Games,” 2015, Cambridge University Press.
- [21] Myerson, Roger, “Multistage Games with Communication,” 1986, *Econometrica*, 54, pp. 323-358.
- [22] Penta, Antonio, “Robust Dynamic Implementation,” 2015, *Journal of Economic Theory*, 160, pp. 280-316.
- [23] Penta, Antonio, and Peio Zuazo-Garin, “Rationalizability, Observability and Common Knowledge,” *Review of Economic Studies*, forthcoming.
- [24] Renault, Jérôme, Elon Solan, and Nicolas Vieille, “Optimal Dynamic Information Provision,” 2017, *Games and Economic Behavior*, 104, pp. 329-349.
- [25] Roesler, Anne-Katrin, and Balász Szentes, “Buyer-optimal Learning and Monopoly Pricing,” 2017, *American Economic Review*, 107, pp. 2072-2080.
- [26] Syrgkanis, Vasilis, Elie Tamer, and Juba Ziani, “Inference on Auctions with Weak Assumptions on Information,” 2018, arXiv:1710.03830.
- [27] Solan, Elon, “Characterization of Correlated Equilibria in Stochastic Games”, 2001, *International Journal of Game Theory*, 30, pp. 1259-277.
- [28] Sugaya, Takuo, and Alexander Wolitzky, “Revelation Principles in Multistage Games,” 2021, *The Review of Economic Studies*, 88, pp. 1503-1540.
- [29] Sugaya, Takuo, and Alexander Wolitzky, “Revelation Principles in Multistage Games,” 2018, mimeo.

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

MILTADIS MAKRIS, DEPARTMENT OF ECONOMICS, UNIVERSITY OF ESSEX, UK

E-mail address: mmakris.econ(at)gmail.com

LUDOVIC RENO, QUEEN MARY UNIVERSITY OF LONDON , CEPR AND UNIVERSITY OF ADELAIDE, MILES END, E1 4NS, LONDON, UK

E-mail address: lrenou.econ(at)gmail.com