

Revelation Principles in Multistage Games*

Takuo Sugaya and Alexander Wolitzky

Stanford GSB and MIT

September 22, 2017

Abstract

We study the revelation principle in multistage games for perfect Bayesian equilibrium and sequential equilibrium. The question of whether or not the revelation principle holds is subtle, as expanding players' opportunities for communication expands the set of consistent beliefs at off-path information sets. In settings with adverse selection but no moral hazard, Nash, perfect Bayesian, and sequential equilibrium are essentially equivalent, and a virtual-implementation version of the revelation principle holds. In general, we show that the revelation principle holds for weak perfect Bayesian equilibrium and for a stronger notion of perfect Bayesian equilibrium introduced by Myerson (1986). Our main result is that the latter notion of perfect Bayesian equilibrium is outcome-equivalent to sequential equilibrium when the mediator can tremble; this provides a revelation principle for sequential equilibrium. When the mediator cannot tremble, the revelation principle for sequential equilibrium can fail.

Preliminary. Comments Welcome.

*For helpful comments, we thank Drew Fudenberg, Dino Gerardi, Shengwu Li, Roger Myerson, Alessandro Pavan, Harry Pei, Juuso Toikka, and Bob Wilson.

1 Introduction

The revelation principle states that any social choice function that can be implemented by any mechanism can also be implemented by a canonical mechanism where communication between players and the mechanism designer or mediator takes a circumscribed form: players communicate only their private information to the mediator, and the mediator communicates only recommended actions to the players. This cornerstone of modern microeconomics was developed by several authors in the 1970s and early 1980s, reaching its most general formulation in the principal-agent model of Myerson (1982), which treats one-shot games with both adverse selection and moral hazard.

More recently, there has been a surge of interest in the design of *dynamic* mechanisms and information systems.¹ The standard logic of the revelation principle applies immediately to dynamic models, if these models are studied under the solution concept of Nash equilibrium: this approach leads to the concept of *communication equilibrium* introduced by Forges (1986). But of course Nash equilibrium is not usually a satisfactory solution concept in dynamic models: following Kreps and Wilson (1982), economists prefer solution concepts that require rationality even after off-path events and impose “consistency” restrictions on players’ beliefs, such as sequential equilibrium or various versions of perfect Bayesian equilibrium. And it is not obvious whether the revelation principle holds for these stronger solution concepts, because—as we will see—expanding players’ opportunities for communication expands the set of consistent beliefs at off-path information sets. The time is therefore ripe to revisit the question of whether the revelation principle holds in dynamic games under the solution concept of sequential or perfect Bayesian equilibrium.

The key paper on the revelation principle in multistage games is Myerson (1986). In this beautiful paper, Myerson introduces the concept of *sequential communication equilibrium (SCE)*, which is a kind of perfect Bayesian equilibrium—what we will call a *conditional probability perfect Bayesian equilibrium (CPPBE)*—in a multistage game played with the canonical communication structure: in every period, players report their private informa-

¹For dynamic mechanism design, see for example Courty and Li (2000), Battaglini (2005), Esó and Szentes (2007), Bergemann and Välimäki (2010), Athey and Segal (2013), and Pavan, Segal, and Toikka (2014). For dynamic information design, see for example Kremer, Mansour, and Perry (2014), Ely, Frankel, and Kamenica (2015), Che and Hörner (2017), Ely (2017), and Renault, Solan, and Vieille (2017).

tion, and the mediator recommends actions. Myerson discusses how the logic of the revelation principle suggests that restricting attention to canonical communication structures is without loss of generality, but he does not state a formal revelation principle theorem. His main result instead provides an elegant and tractable characterization of SCE: a SCE is a communication equilibrium in which players avoid codominated actions, a generalization of dominated actions. Myerson's paper also proves that there is an equivalence between conditional probability systems—the key object used to restrict off-path beliefs in his solution concept—and limits of beliefs derived from full support probability distributions over moves. This result establishes an analogy between Myerson's belief restrictions and the consistency requirement of Kreps and Wilson. However, the analogy is not exact, because the probability distributions over moves used to generate beliefs in Myerson's approach need not be strategies: for example, some conditional probability systems can be generated only by supposing that a player takes different actions at two nodes in the same information set.²

Myerson's paper thus leaves open two important questions: First, when one formulates the CPPBE concept more generally—so that it can be applied to any communication system—is it indeed without loss of generality to restrict attention to canonical communication systems? Second, is there an equivalence between implementation under this perfect Bayesian equilibrium concept and implementation in sequential equilibrium, so that one can safely rely on Myerson's characterization while still imposing the more restrictive consistency requirement of Kreps and Wilson?

The main contribution of the current paper is to answer both of these questions in the affirmative, thus providing a more secure foundation for the use of the revelation principle in dynamic settings. That is, we prove that the revelation principle holds in dynamic games for the solution concept of CPPBE (as well as under weaker solution concepts, such as weak PBE), and we prove that a social choice function can be implemented in sequential equilibrium for some communication system if and only if it is a SCE. The first of these results may be viewed as a formalization of ideas implicit in Myerson (1986). The second is quite subtle and (to us) unexpected.

²This gap between Myerson's solution concept and sequential equilibrium has been noted before. See, for example, Fudenberg and Tirole (1991).

While the main message of this paper is a positive one, we must immediately add a significant caveat. The claimed equivalence between sequential equilibrium and SCE depends on a subtlety in the definition of sequential equilibrium in games with communication: whether or not the mediator is allowed to tremble, or more precisely whether players are allowed to attribute off-path observations to deviations by the mediator instead of or in addition to deviations by other players. Letting the mediator tremble yields a slightly more permissive version of sequential equilibrium, and it is this version that we show is outcome-equivalent to SCE. If instead the mediator cannot tremble—that is, if one insists on treating the mechanism as part of the immutable physical environment—then the revelation principle for sequential equilibrium fails, and there are SCE that cannot be implemented in sequential equilibrium.³ When interpreting these results, we think “whether the mediator can tremble” is better viewed as a choice of solution concept on the part of the modeler rather than a descriptive property of the mechanism. (After all, “mediator trembles” enter the solution concept only through the set of allowable player beliefs and not through the set of allowable mediator strategies.) We feel that the stronger version of sequential equilibrium is conceptually appealing, but that the equivalence with SCE and the resulting improvement in tractability are arguments in favor of the weaker version. We further discuss whether it is “reasonable” for the mediator to tremble once the solution concepts have been precisely defined.

In any case, the failure of the revelation principle when the mediator cannot tremble can be avoided if the game satisfies appropriate full support conditions. We clarify what conditions are required. For example, in settings with adverse selection but no moral hazard, we show that Nash, perfect Bayesian, and sequential equilibrium are essentially equivalent and a virtual-implementation version of the revelation principle always holds. As this class of games encompasses most of the recent literature on dynamic mechanism design, this virtual revelation principle may be viewed as a third main result of our paper.

By way of further motivation for our paper, we note that there seems to be some uncertainty in the literature as to what is known about the revelation principle in multistage

³In the context of one-shot games, Gerardi and Myerson (2007) emphasize the role of the mediator’s trembles and show that the revelation principle for sequential equilibrium can fail if the mediator cannot tremble. We build on this insight by showing that this failure can be even more dramatic in multistage games. We discuss Gerardi and Myerson’s results below.

games. A standard (and reasonable) approach in the dynamic mechanism design literature is to cite Myerson and then restrict attention to direct mechanisms without quite claiming that this is without loss generality. Pavan, Segal, and Toikka (2014, p. 611) is representative:

“Following Myerson (1986), we restrict attention to direct mechanisms where, in every period t , each agent i confidentially reports a type from his type space Θ_{it} , no information is disclosed to him beyond his allocation x_{it} , and the agents report truthfully on the equilibrium path. Such a mechanism induces a dynamic Bayesian game between the agents and, hence, we use perfect Bayesian equilibrium (PBE) as our solution concept.”

Our results provide a foundation for this approach, while also showing that Nash and perfect Bayesian equilibrium are essentially outcome-equivalent in pure adverse selection settings like this one.⁴

Other papers do claim that Myerson (1986) shows that restricting to direct mechanisms is without loss: see for example Kakade, Lobel, and Nazerzadhe (2013), Kremer, Mansour, and Perry (2014), and Board and Skrzypacz (2016). Indeed, we ourselves have incorrectly claimed that the revelation principle holds for sequential equilibrium in repeated games where the mediator cannot tremble (Sugaya and Wolitzky, 2017). The results in all of these papers remain valid, but—as the examples in the next section show—these models are similar to ones where the naïve application of the revelation principle can lead to serious problems. We hope this paper will allow the emerging literature on dynamic mechanism and information design to rely on the revelation principle with more confidence. To this end, we provide a compact summary of our results at the end of the paper.

⁴A caveat is that much of the dynamic mechanism design literature assumes continuous type spaces to facilitate the use of the envelope theorem, while we restrict attention to finite games to have a well-defined notion of sequential equilibrium. We do believe that our results for perfect Bayesian equilibrium concepts should extend to continuous type spaces.

2 Examples

We begin by noting some settings where the revelation principle for sequential equilibrium can fail when the mediator cannot tremble.⁵

Example 1: Pure Adverse Selection

Our first example is a game of pure adverse selection, where players have private information but only the mediator takes payoff-relevant actions. There are three players (in addition to the mediator) and two periods.

In period 1, player 1 observes a binary signal s_1 , which takes on value a or b with probability $1/2$ each. The mediator then takes an action $a_1 \in \{A, B\}$.

In period 2, player 2 observes signal $s_2 = (s_1, a_1)$; that is, he observes player 1's signal and the mediator's period 1 action. The mediator then takes an action $a_2 \in \{C, D\}$.

Players 1 and 2 have the same payoff function, given by

$$u(a_1, a_2) = \mathbf{1}_{\{a_1=A\}} + \mathbf{1}_{\{a_2=C\}},$$

where $\mathbf{1}_{\{\cdot\}}$ is the indicator function. Player 3 is a dummy player who observes nothing and is indifferent among all outcomes. Thus, in a canonical mechanism, player 1 reports her signal to the mediator in period 1; player 2 reports his signal in period 2; and player 3 does not communicate with the mediator.

Consider the outcome distribution given by $\Pr(A|a) = 1$, $\Pr(B|b) = 1$, $\Pr(C) = 1$. We claim that this is not implementable in a canonical mechanism. For suppose that it is, and suppose player 1 misreports $s_1 = b$ as a . The mediator must then take $a_1 = A$ with probability 1, as $\Pr(A|a) = 1$. Hence, for this misreport to be unprofitable for player 1, the mediator must take $a_2 = D$ with probability 1 conditional on the event that player 1 misreports b as a . Now suppose player 2 observes signal $s_2 = (b, A)$. As the mediator does not tremble and $\Pr(B|b) = 1$, in a canonical mechanism this signal can arise only if player 1 misreported b as a . So if player 2 follows his equilibrium strategy at this history, the

⁵This failure has nothing to do with the failure of revelation principle-like results in settings with hard evidence (Green and Laffont, 1986), common agency (Epstein and Peters, 1999; Martimort and Stole, 2002), limited commitment (Bester and Strausz, 2000, 2001), or computational limitations (Conitzer and Sandholm, 2004).

mediator will take $a_2 = D$ with probability 1. On the other hand, if player 2 misreports his signal as (a, A) , then the mediator's history will be the same as it would be if player 1's signal were a and players 1 and 2 reported truthfully, in which case the mediator takes $a_2 = C$ with probability 1. So player 2 will misreport.

Contrast this with the situation under the non-canonical mechanism where, in addition to the usual communication between players 1 and 2 and the mediator, player 3 sends the mediator a binary message $r \in \{0, 1\}$ at the beginning of the game. Suppose player 3 sends $r = 0$ with equilibrium probability 1, but trembles to sending $r = 1$ with probability $1/k$ along a sequence of strategy profiles σ^k converging to the equilibrium, while players 1 and 2 report their signals truthfully and tremble with probability $1/k^2$. Suppose as well that the mediator takes $a_1 = A$ if player 1 reports $s_1 = a$ or $r = 1$, and takes $a_1 = B$ if player 1 reports $s_1 = b$ and $r = 0$. Finally, the mediator takes $a_2 = D$ if and only if player 1 reports $s_1 = a$, player 2 reports that $s_1 = b$, and $a_1 = A$.

This assessment implements the desired outcome distribution. In particular, after observing $s_2 = (b, A)$, player 2 believes with probability 1 that player 1 truthfully reported $s_1 = b$ but player 3 sent $r = 1$. Player 2 is therefore willing to truthfully report his signal. This behavior in turn deters player 1 from misreporting, as if she misreports $s_1 = b$ as a this leads (with probability 1) to a payoff gain of 1 in period 1 but a payoff loss of 1 in period 2.

Example 2: Repeated Games with Complete Information

The revelation principle can also fail in repeated games with complete information and perfect monitoring. Consider the following stage game:

	A_2	B_2	C_2		A_2	B_2	C_2
A_1	1, 1, 1	1, 1, 1	0, 1, 0	A_1	1, 1, 1	1, 1, 1	1, 1, 1
B_1	2, 1, 0	1, 1, 0	0, 1, 0	B_1	1, 1, 1	1, 1, 1	1, 1, 1
	$a_3 = A_3$				$a_3 = B_3$		

The stage game is played twice. In each period, players receive messages from the mediator before taking actions, and the actions played are perfectly observed by all players and

the mediator.⁶ A player’s payoff is the (undiscounted) sum of her two stage game payoffs. In this setting, an equilibrium is canonical if in each period the mediator’s message to each player consists of her recommended action only.

We claim that there is a non-canonical sequential equilibrium in which action profile (A_1, A_2, A_3) is played with positive probability in the first period, but that this cannot occur in a canonical sequential equilibrium (maintaining in both cases the assumption that the mediator cannot tremble). The intuition is as follows: If (A_1, A_2, A_3) is played with positive probability, then player 1 has an instantaneous benefit from deviating to B_1 . To deter this, player 1 must be punished with a play of (A_1, C_2, A_3) or (B_1, C_2, A_3) . But B_3 guarantees player 3 her best payoff of 1, so if player 3 believes that player 2 plays C_2 with positive probability, she prefers to play B_3 . Thus, to enforce (A_1, A_2, A_3) , we must be able to punish a deviation by player 1 without letting player 3 understand that this is occurring. The only way to do this is make player 3 believe that, when action profile (B_1, A_2, A_3) is played, this always corresponds to a simultaneous tremble by players 1 and 2 from recommendation profile (A_1, B_2, A_3) (in which case player 1’s tremble is not expected to be profitable, and thus does not need to be punished). In particular, player 3 must be made to believe that player 1’s tremble from A_1 to B_1 is correlated with player 2’s recommended action. Finally, in a sequential equilibrium, it is not possible to induce such a belief if players 2 and 3 observe only their own actions, but it is possible if they also observe an additional correlated signal. The details may be found in the appendix.

Further Examples: Gerardi and Myerson (2007)

Gerardi and Myerson (2007) also present examples where the revelation principle fails for sequential equilibrium when the mediator cannot tremble. Their examples involve one-shot games with both adverse selection and moral hazard, where some type profiles have 0 probability. The two examples above thus extend their insight by showing that, in multistage games, the revelation principle can fail even in settings with pure adverse selection or pure moral hazard. Our examples also seem a bit simpler.⁷

⁶The example also works if only the players observe actions. Details are in the appendix.

⁷In a related example, Dhillon and Mertens (1996; Example 1) show that the revelation principle fails for the solution concept of “perfect correlated equilibrium,” which describes outcomes that can be implemented with communication in trembling-hand perfect equilibrium.

In all of these examples, it is clear that the revelation principle would be restored if the mediator were allowed to tremble. Our main result is that this is a general phenomenon.

Note also that a canonical mechanism would suffice in Example 1 if we perturbed the target outcome distribution. We will see that, in general, a virtual-implementation version of the revelation principle holds in games of pure adverse selection. In contrast, Example 2 shows that no such result holds for games with moral hazard.

3 Multistage Games with Communication

3.1 Model

As in Forges (1986) and Myerson (1986), we consider multistage games with communication. A multistage game G is played by $N + 1$ players (indexed by $i = 0, 1, \dots, N$) over T periods (indexed by $t = 1, \dots, T$). Player 0 is a mediator who differs from the other players in three ways: (i) the players communicate only with the mediator and not directly with each other, (ii) the mediator is indifferent over outcomes of the game (and can thus “commit” to any strategy), (iii) depending on the solution concept, “trembles” by the mediator may be disallowed.⁸ In each period t , each player i (including the mediator) has a set of possible signals $S_{i,t}$ and a set of possible actions $A_{i,t}$, and each player other than mediator has a set of possible reports to send to the mediator $R_{i,t}$ and a set of possible messages to receive from the mediator $M_{i,t}$. These sets are all assumed finite. This formulation lets us capture settings where the mediator receives exogenous signals in addition to reports from the players, as well as settings where the mediator takes actions (such as choosing allocations for the players). Let $H^t = \prod_{\tau=1}^{t-1} \left(\prod_{i=0}^N S_{i,\tau}, \prod_{i=1}^N R_{i,\tau}, \prod_{i=1}^N M_{i,\tau}, \prod_{i=0}^N A_{i,\tau} \right)$ denote the set of possible histories of signals, reports, messages, and actions (“complete histories”) at the beginning of period t , with $H^1 = \emptyset$. Let $\hat{H}^t = \prod_{\tau=1}^{t-1} \prod_{i=0}^N (S_{i,\tau}, A_{i,\tau})$ denote the set of possible histories of signals and actions (“payoff-relevant histories”) at the beginning of period t . Given a complete history $h^t \in H^t$, let $\hat{h}^t = \prod_{\tau=1}^{t-1} \prod_{i=0}^N (s_{i,\tau}, a_{i,\tau})$ denote the projection of h^t onto \hat{H}^t . Let $X = \hat{H}^{T+1} = \prod_{\tau=1}^T \prod_{i=0}^N (S_{i,\tau}, A_{i,\tau})$ denote the set of final, payoff-relevant outcomes of the

⁸We also use male pronouns for the mediator and female pronouns for the players.

game. Let $u_i : X \rightarrow \mathbb{R}$ denote player i 's payoff function, where u_0 is a constant function.

The timing within each period t is as follows:

1. A signal $s_t \in S_t = \prod_{i=0}^N S_{i,t}$ is drawn with probability $p(s_t | \hat{h}^t)$, where $\hat{h}^t \in \hat{H}^t$ is the current history of signals and actions. Player i observes $s_{i,t}$, the i^{th} component of s_t .
2. Each player $i \neq 0$ chooses a report $r_{i,t} \in R_{i,t}$ to send to the mediator.
3. The mediator chooses a message $m_{i,t} \in M_{i,t}$ to send to each player $i \neq 0$.
4. Each player i takes an action $a_i \in A_{i,t}$.

We refer to the tuple $(N, T, S, A, u, p) = (N, T, \prod_{\tau=1}^T \prod_{i=0}^N S_{i,\tau}, \prod_{\tau=1}^T \prod_{i=0}^N A_{i,\tau}, \prod_{i=0}^N u_i, p)$ as the *base game* and refer to the pair $(R, M) = \prod_{\tau=1}^T \prod_{i=1}^N (R_{i,\tau}, M_{i,\tau})$ as the *message set*. Assume without loss of generality that $S_{i,t} = \bigcup_{\hat{h}^t \in \hat{H}^t} \text{supp } p_i(\cdot | \hat{h}^t)$ for all i, t , where p_i denotes the marginal distribution of p .

For $i \neq 0$, let $H_i^t = \prod_{\tau=1}^{t-1} (S_{i,\tau}, R_{i,\tau}, M_{i,\tau}, A_{i,\tau})$ denote the set of player i 's possible histories of signals, reports, messages, and actions at the beginning of period t . Let $H_0^t = \prod_{\tau=1}^{t-1} (S_{0,\tau}, \prod_{i=1}^N R_{i,\tau}, \prod_{i=1}^N M_{i,\tau}, A_{0,\tau})$ denote the set of the mediator's possible histories of signals, reports, messages, and actions at the beginning of period t . When a complete history $h^t \in H^t$ is understood, we let $h_i^t = \prod_{\tau=1}^{t-1} (s_{i,\tau}, r_{i,\tau}, m_{i,\tau}, a_{i,\tau})$ denote the projection of h^t onto H_i^t , and let $h_0^t = \prod_{\tau=1}^{t-1} (s_{0,\tau}, \prod_{i=1}^N r_{i,\tau}, \prod_{i=1}^N m_{i,\tau}, a_{0,\tau})$ denote the projection of h^t onto H_0^t . We use the notation \hat{h}_i^t and \hat{h}_0^t analogously.

A strategy for player $i \neq 0$ is a function $\sigma_i = (\sigma_i^R, \sigma_i^A) = (\sigma_{i,t}^R, \sigma_{i,t}^A)_{t=1}^T$, where $\sigma_{i,t}^R : H_i^t \times S_{i,t} \rightarrow \Delta(R_{i,t})$ and $\sigma_{i,t}^A : H_i^t \times S_{i,t} \times R_{i,t} \times M_{i,t} \rightarrow \Delta(A_{i,t})$. Let Σ_i be the set of player i 's strategies, and let $\Sigma = \prod_{i=1}^N \Sigma_i$. A belief for player $i \neq 0$ is a function $\beta_i = (\beta_i^R, \beta_i^A) = (\beta_{i,t}^R, \beta_{i,t}^A)_{t=1}^T$, where $\beta_{i,t}^R : H_i^t \times S_{i,t} \rightarrow \Delta(H^t \times S_t)$ and $\beta_{i,t}^A : H_i^t \times S_{i,t} \times R_{i,t} \times M_{i,t} \rightarrow \Delta(H^t \times S_t \times R_t \times M_t)$. A strategy for the mediator is a function $\mu = (\mu^M, \mu^A) = (\mu_t^M, \mu_t^A)_{t=1}^T$, where $\mu_t^M : H_0^t \times S_{0,t} \times R_t \rightarrow \Delta(M_t)$ and $\mu_t^A : H_0^t \times S_{0,t} \times R_t \times M_t \rightarrow \Delta(A_{0,t})$. We write $\sigma_{i,t}^R(r_{i,t} | h_i^t, s_{i,t})$ for $\sigma_{i,t}^R(h_i^t, s_{i,t})(r_{i,t})$, and similarly for $\sigma_{i,t}^A, \beta_{i,t}^R, \beta_{i,t}^A, \mu_t^M$, and μ_t^A . When the meaning is unambiguous, we omit the superscript R, A , or M and the subscript t from σ_i, β_i , and μ , so that, for example, σ_i can take as its argument either a pair $(h_i^t, s_{i,t})$ or a tuple $(h_i^t, s_{i,t}, r_{i,t}, m_{i,t})$. We extend players' payoff functions from terminal histories

to strategy profiles in the usual way, writing $\bar{u}_i(\sigma, \mu)$ for player i 's expected payoff at the beginning of the game under strategy profile (σ, μ) , and writing $\bar{u}_i(\sigma, \mu|h^t)$ for player i 's expected payoff conditional on reaching the complete history h^t .

To economize on notation, we let $\mathfrak{h}^t = (h^t, s_t)$, $\mathfrak{h}_i^t = (h_i^t, s_{i,t})$, etc..

3.2 Nash and Perfect Bayesian Equilibrium

A *Nash equilibrium (NE)* is a strategy profile (σ, μ) such that $\bar{u}_i(\sigma, \mu) \geq \bar{u}_i(\sigma'_i, \sigma_{-i}, \mu)$ for all $i \neq 0$, $\sigma'_i \in \Sigma_i$.

We consider two versions of perfect Bayesian equilibrium. A *weak perfect Bayesian equilibrium (WPBE)* is an assessment (σ, μ, β) such that

- [*Sequential rationality of messages*] For all $i \neq 0, t, \sigma'_i \in \Sigma_i$ and all $\mathfrak{h}_i^t \in (H_i^t, S_{i,t})$,

$$\sum_{\mathfrak{h}^t \in H^t \times S_t} \beta_i(\mathfrak{h}^t | \mathfrak{h}_i^t) \bar{u}_i(\sigma, \mu | \mathfrak{h}^t) \geq \sum_{\mathfrak{h}^t \in H^t \times S_t} \beta_i(\mathfrak{h}^t | \mathfrak{h}_i^t) \bar{u}_i(\sigma'_i, \sigma_{-i}, \mu | \mathfrak{h}^t). \quad (1)$$

- [*Sequential rationality of actions*] For all $i \neq 0, t, \sigma'_i \in \Sigma_i$ and all $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) \in (H_i^t, S_{i,t}, R_{i,t}, M_{i,t})$,

$$\begin{aligned} & \sum_{(\mathfrak{h}^t, r_t, m_t) \in H^t \times S_t \times R_t \times M_t} \beta_i(\mathfrak{h}^t, r_t, m_t | \mathfrak{h}_i^t, r_{i,t}, m_{i,t}) \bar{u}_i(\sigma, \mu | \mathfrak{h}^t, r_t, m_t) \\ \geq & \sum_{(\mathfrak{h}^t, r_t, m_t) \in H^t \times S_t \times R_t \times M_t} \beta_i(\mathfrak{h}^t, r_t, m_t | \mathfrak{h}_i^t, r_{i,t}, m_{i,t}) \bar{u}_i(\sigma'_i, \sigma_{-i}, \mu | \mathfrak{h}^t, r_t, m_t). \end{aligned} \quad (2)$$

- [*Bayes rule*] For all i , all $\mathfrak{h}_i^t \in (H_i^t, S_{i,t})$ such that $\Pr^{(\sigma, \mu)}(\mathfrak{h}_i^t) > 0$, and all $\mathfrak{h}^t \in (H^t, S_t)$,

$$\beta_i(\mathfrak{h}^t | \mathfrak{h}_i^t) = \frac{\Pr^{(\sigma, \mu)}(\mathfrak{h}^t)}{\Pr^{(\sigma, \mu)}(\mathfrak{h}_i^t)}.$$

Similarly, for all i , all $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) \in (H_i^t, S_{i,t}, R_{i,t}, M_{i,t})$ such that $\Pr^{(\sigma, \mu)}(\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) > 0$, and all $(\mathfrak{h}^t, r_t, m_t) \in (H^t, S_t, R_t, M_t)$,

$$\beta_i(\mathfrak{h}^t, r_t, m_t | \mathfrak{h}_i^t, r_{i,t}, m_{i,t}) = \frac{\Pr^{(\sigma, \mu)}(\mathfrak{h}^t, r_t, m_t)}{\Pr^{(\sigma, \mu)}(\mathfrak{h}_i^t, r_{i,t}, m_{i,t})}.$$

The notation above is that $\Pr^{(\sigma, \mu)}(\mathfrak{h}^t)$ is the probability of reaching history \mathfrak{h}^t under strategy profile (σ, μ) . Note that the requirement that (σ, μ) is a strategy profile implies that players' randomizations are stochastically independent and respect the information structure of the game, in that a player must use the same mixing probability at all nodes in the same information set.

A more refined notion of perfect Bayesian equilibrium requires that beliefs are derived from a common conditional probability system (CPS) on H^{T+1} , the set of complete histories of play (Myerson, 1986).⁹ We say that a *conditional probability perfect Bayesian equilibrium (CPPBE)* is a WPBE such that there exists a CPS f on H^{T+1} with

$$\begin{aligned} \sigma_i(r_{i,t}|\mathfrak{h}_i^t) &= f(r_{i,t}|\mathfrak{h}_i^t) \\ \sigma_i(a_{i,t}|\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) &= f(a_{i,t}|\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) \\ \mu(m_t|\mathfrak{h}_0^t, r_t) &= f(m_t|\mathfrak{h}_0^t, r_t) \\ \mu(a_{0,t}|\mathfrak{h}_0^t, r_t, m_t) &= f(a_{0,t}|\mathfrak{h}_0^t, r_t, m_t) \\ \beta_i(\mathfrak{h}^t|\mathfrak{h}_i^t) &= f(\mathfrak{h}^t|\mathfrak{h}_i^t) \\ \beta_i(\mathfrak{h}^t, r_t, m_t|\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) &= f(\mathfrak{h}^t, r_t, m_t|\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) \end{aligned}$$

for all $i \neq 0, t, \mathfrak{h}_i^t, \mathfrak{h}_0^t, \mathfrak{h}^t, r_{i,t}, r_t, m_{i,t}, m_t, a_{i,t}, a_{0,t}$.

Myerson did not explicitly formulate the CPPBE concept for general games, but this concept is not really new. For example, Fudenberg and Tirole (1991), Battigalli (1996), and Kohlberg and Reny (1997) study whether imposing additional independence conditions on CPPBE leads to an equivalence with sequential equilibrium in general games. In contrast, our main result is that CPPBE and sequential equilibrium are outcome-equivalent in games with communication (when the mediator can tremble). The basic reason why independence conditions are not required to obtain equivalence with sequential equilibrium in games with communication is that the correlation allowed by CPPBE can be replicated through correlation in the mediator's messages.

⁹Recall that a CPS on a finite set Ω is a function $f(\cdot|\cdot) : 2^\Omega \times 2^\Omega \rightarrow [0, 1]$ such that (i) for all $Z \subseteq \Omega$, $f(\cdot|Z)$ is a probability distribution on Z , and (ii) for all $X \subseteq Y \subseteq Z \subseteq \Omega$ with $Y \neq \emptyset$, $f(X|Y)f(Y|Z) = f(X|Z)$.

3.3 Sequential Equilibrium

We follow Kreps and Wilson's (1982) definition of sequential equilibrium. An important issue is whether the mediator is modeled as a player in the game who is allowed to tremble, or as a machine that executes its strategy perfectly. This distinction leads to two different notions of sequential equilibrium, which we refer to as *player sequential equilibrium (PSE)* and *machine sequential equilibrium (MSE)*. It will become clear that every MSE can be extended to a PSE, because, even if the mediator is allowed to tremble, the players can always believe that he does not tremble (or trembles with very small probability). In addition, Theorem 1 of Myerson (1986) shows that every PSE induces a CPS on H^{T+1} . This gives the chain of inclusions

$$NE \supseteq WPBE \supseteq CPPBE \supseteq PSE \supseteq MSE.$$

Formally, a PSE is an assessment (σ, μ, β) such that

- [*Sequential rationality of messages*] For all $i \neq 0, t, \sigma'_i \in \Sigma_i$ and all $\mathfrak{h}_i^t \in (H_i^t, S_{i,t})$, (1) holds.
- [*Sequential rationality of actions*] For all $i \neq 0, t, \sigma'_i \in \Sigma_i$ and all $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) \in (H_i^t, S_{i,t}, R_{i,t}, M_{i,t})$, (2) holds.
- [*Player Consistency*] There exists a sequence of completely mixed strategy profiles $(\sigma^n, \mu^n)_{n=1}^\infty$ such that $\lim_{n \rightarrow \infty} (\sigma^n, \mu^n) = (\sigma, \mu)$;

$$\beta_i(\mathfrak{h}^t | \mathfrak{h}_i^t) = \lim_{n \rightarrow \infty} \frac{\Pr^{(\sigma^n, \mu^n)}(\mathfrak{h}^t)}{\Pr^{(\sigma^n, \mu^n)}(\mathfrak{h}_i^t)}$$

for all i , all $\mathfrak{h}_i^t \in (H_i^t, S_{i,t})$, and all $\mathfrak{h}^t \in (H^t, S_t)$; and

$$\beta_i(\mathfrak{h}^t, r_t, m_t | \mathfrak{h}_i^t, r_{i,t}, m_{i,t}) = \lim_{n \rightarrow \infty} \frac{\Pr^{(\sigma^n, \mu^n)}(\mathfrak{h}^t, r_t, m_t)}{\Pr^{(\sigma^n, \mu^n)}(\mathfrak{h}_i^t, r_{i,t}, m_{i,t})}$$

for all i , all $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) \in (H_i^t, S_{i,t}, R_{i,t}, M_{i,t})$ and all $(\mathfrak{h}^t, r_t, m_t) \in (H^t, S_t, R_t, M_t)$.

Note that the requirement that (σ^n, μ^n) is a strategy profile implies that players' trembles are stochastically independent. Also, as usual, a single sequence of trembles must be used

to rationalize all off-path beliefs.

In defining an MSE, we impose sequential rationality and consistency only at histories consistent with the mediator's strategy. The reason is that the mediator can be identified with nature under the machine interpretation, so nodes inconsistent with mediator's strategy may be pruned from the game tree. An alternative, equivalent approach would be to define an MSE as a PSE in which players believe with probability 1 that the mediator has not deviated at any history consistent with the mediator's strategy.

Formally, let

$$\begin{aligned} H_i^t(\mu) &= \left\{ h_i^t : \Pr^{(\sigma, \mu)}(h_i^t) > 0 \text{ for some } \sigma \in \Sigma \right\}, \\ (H, S)_i^t(\mu) &= \left\{ \mathfrak{h}_i^t : \Pr^{(\sigma, \mu)}(\mathfrak{h}_i^t) > 0 \text{ for some } \sigma \in \Sigma \right\}, \text{ and} \\ (H, S, R, M)_i^t(\mu) &= \left\{ (\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) : \Pr^{(\sigma, \mu)}(\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) > 0 \text{ for some } \sigma \in \Sigma \right\} \end{aligned}$$

denote the set of period t histories for player i that can be reached under strategy profile (σ, μ) for some σ . An MSE is an assessment (σ, μ, β) such that

- *[Sequential rationality of messages]* For all $i \neq 0, t, \sigma'$ and all $\mathfrak{h}_i^t \in (H, S)_i^t(\mu)$, (1) holds.
- *[Sequential rationality of actions]* For all $i \neq 0, t, \sigma'$ and all $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) \in (H, S, R, M)_i^t(\mu)$, (2) holds.
- *[Machine Consistency]* There exists a sequence of completely mixed player strategy profiles $(\sigma^n)_{n=1}^\infty$ such that $\lim_{n \rightarrow \infty} \sigma^n = \sigma$;

$$\beta_i(\mathfrak{h}^t | \mathfrak{h}_i^t) = \lim_{n \rightarrow \infty} \frac{\Pr^{(\sigma^n, \mu)}(\mathfrak{h}^t)}{\Pr^{(\sigma^n, \mu)}(\mathfrak{h}_i^t)}$$

for all i , all $\mathfrak{h}_i^t \in (H, S)_i^t(\mu)$, and all $\mathfrak{h}^t \in (H^t, S_t)$; and

$$\beta_i(\mathfrak{h}^t, r_t, m_t | \mathfrak{h}_i^t, r_{i,t}, m_{i,t}) = \lim_{n \rightarrow \infty} \frac{\Pr^{(\sigma^n, \mu)}(\mathfrak{h}^t, r_t, m_t)}{\Pr^{(\sigma^n, \mu)}(\mathfrak{h}_i^t, r_{i,t}, m_{i,t})}$$

for all i , all $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) \in (H, S, R, M)_i^t(\mu)$ and all $(\mathfrak{h}^t, r_t, m_t) \in (H^t, S_t, R_t, M_t)$.

Let us comment on the “reasonableness” of allowing mediator trembles. In several recent papers, a mediator is introduced as a mathematical stand-in for the set of all possible information structures in a game, on the logic that the set of outcomes that can be implemented for some information structure coincides with the set of outcomes that can be implemented with the assistance of an “omniscient” mediator who knows the realizations of all payoff-relevant variables. See Bergemann and Morris (2013, 2016) in the context of static incomplete information games and Sugaya and Wolitzky (2017) in the context of repeated complete information games. This approach is valid only under the assumption that the mediator does not tremble: letting the mediator tremble would be analogous to letting nature tremble in the unmediated game, which is not allowed under any standard solution concept.

Outside of the special context where the mediator is seen as a stand-in for a collection of information structures, we view the issue of whether the mediator should be allowed to tremble as one of taste and tractability. In particular, while we find the stronger consistency requirement of MSE attractive, our main result shows that in general the set of PSE-implementable outcomes has a more tractable structure. In terms of applications, one possible approach is to use our results to characterize the set of PSE-implementable outcomes, and then attempt to verify directly that the “optimal” PSE outcome is also implementable in MSE.

Gerardi and Myerson (2007) make a similar point in the context of one-shot games. Their main result is a characterization of MSE in one-shot games. The characterization is rather complicated, and the authors suggest that “it may be simpler to admit the possibility that the mediator makes mistakes and use the concept of SCE,” (p. 124). We agree: our theorem shows that this approach is also valid in multistage games.¹⁰

3.4 Restricted Solution Concepts

The equilibrium definitions given above are the natural ones from a game-theoretic perspective. However, towards introducing the revelation principle, it will be helpful to consider additional restrictions on the mediator’s possible messages. These restrictions will let us re-

¹⁰The current paper does not attempt to characterize MSE in multistage games. This seems to be a difficult problem.

quire that players obey the mediator’s recommendations even when the mediator is allowed to tremble, and will also clarify that letting the mediator tremble can only expand the set of implementable outcomes.

Following Myerson (1986), a *mediation range* $Q = (Q_i)_{i \neq 0}$ specifies a set of possible messages $Q_i \left((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t} \right) \subseteq M_i^t$ that can be received by each player i when the history of communications between player i and the mediator is given by $\left((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t} \right)$. Given a game G and a mediation range Q , let $G|_Q$ denote the game that results when the mediator is restricted to sending messages in Q : that is, when M_i^t is replaced by $Q_i \left((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t} \right)$ at every history $\left(\tilde{h}_0^t, \tilde{s}_{0,t}, \tilde{r}_t \right)$ with $\left((\tilde{r}_{i,\tau}, \tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \tilde{r}_{i,t} \right) = \left((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t} \right)$. A *restricted NE* (resp., WPBE, CPPBE, PSE, MSE) in G is a mediation range Q together with a strategy profile (σ, μ) (resp., assessment (σ, μ, β)) that forms a NE (resp., WPBE, CPPBE, PSE, MSE) in game $G|_Q$.

As one can always take $Q_i \left((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t} \right) = M_i^t$, it is clear that every equilibrium outcome is also a restricted equilibrium outcome. The following lemma says that the converse also holds, so that restricting the mediator’s possible messages does not expand the set of implementable outcomes. The proof is deferred to the appendix.

Lemma 1 *An outcome distribution $\rho \in \Delta(X)$ arises in a NE (resp., WPBE, CPPBE, PSE, MSE) of G if and only if there exists a mediation range Q such that ρ arises in a NE (resp., WPBE, CPPBE, PSE, MSE) of $G|_Q$.*

An implication of Lemma 1 is that every MSE outcome distribution is also a PSE outcome distribution, as every MSE (σ, μ) in G is a PSE in $G|_Q$, where $Q_i \left((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t} \right)$ is the set of messages that can arise under some strategy profile (σ', μ) in which the mediator follows his equilibrium strategy.

3.5 The Revelation Principle

Given a base game (N, T, S, A, u, p) , the *canonical message set* (R, M) is given by $R_{i,t} = A_{i,t-1} \times S_{i,t}$ and $M_{i,t} = A_{i,t}$, for all $i \neq 0$ and t . That is, a message set is canonical if players’ reports are actions and signals and the mediator’s messages are “recommended” actions. A game is *canonical* if its message set is canonical. For any game G , let $C(G)$ denote the

canonical game with the same base game as G , and let $C(\Sigma)$ denote the set of player strategy profiles in the canonical game. Given a canonical game, a strategy profile (σ, μ) together with a mediation range Q is *canonical* if the following conditions hold:

1. [Players are truthful if they have been truthful in the past] $\sigma_i^R(h_i^t, s_{i,t}) = (a_{i,t-1}, s_{i,t})$ for all $(h_i^t, s_{i,t}) \in H_i^t \times S_{i,t}$ such that $r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$ and $m_{i,\tau} \in Q_i(\prod_{\tau'=1}^{\tau-1} ((a_{i,\tau'-1}, s_{i,\tau'}), a_{i,\tau'}), (a_{i,\tau-1}, s_{i,\tau}))$ for all $\tau < t$.
2. [Players obey all possible recommendations if they have been truthful in the past] $\sigma_i^A(h_i^t, s_{i,t}, r_{i,t}, m_{i,t}) = m_{i,t}$ for all $(h_i^t, s_{i,t}, r_{i,t}, m_{i,t}) \in H_i^t \times S_{i,t} \times R_{i,t} \times M_{i,t}$ such that $r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$ and $m_{i,\tau} \in Q_i(\prod_{\tau'=1}^{\tau-1} ((a_{i,\tau'-1}, s_{i,\tau'}), a_{i,\tau'}), (a_{i,\tau-1}, s_{i,\tau}))$ for all $\tau \leq t$.

An assessment (σ, μ, β) of a canonical game is *canonical* if the strategy profile (σ, μ) is canonical and in addition each player believes that her opponents have been truthful if she herself has been truthful (but possibly not obedient): if $(h^t, s_t) \in \text{supp } \beta_{i,t}^R(h_i^t, s_{i,t})$, $r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$, and $m_{i,\tau} \in Q_i(\prod_{\tau'=1}^{\tau-1} ((a_{i,\tau'-1}, s_{i,\tau'}), a_{i,\tau'}), (a_{i,\tau-1}, s_{i,\tau}))$ for all $\tau < t$, then $r_{j,\tau} = (a_{j,\tau-1}, s_{j,\tau})$ for all $j \neq i$ and all $\tau < t$ (and similarly for histories $(h^t, s_t, r_t, m_t) \in \text{supp } \beta_{i,t}^A(h_i^t, s_{i,t}, r_{i,t}, m_{i,t})$). We will also sometimes refer to a strategy profile or assessment as canonical without specifying a mediation range; in this case, the mediation range should be understood to be the trivial one where no recommendations are ruled out.

We wish to assess the validity of the following proposition:

Revelation Principle For any game G , any distribution over outcomes X that arises in any equilibrium of G also arises in a canonical equilibrium of $C(G)$.

A remark: Townsend (1988) extends the revelation principle by requiring a player to be truthful and obedient even if she has previously lied to the mediator, and correspondingly lets players report their entire history of actions and signals every period (thus giving players opportunities to “confess” to any lie). While one could conduct a similar exercise for Townsend’s version of the revelation principle, in this paper we follow Myerson (1986) in imposing truthfulness and obedience only for previously-truthful players.

4 The Revelation Principle for Full-Support Equilibria and PBE

We establish several versions of the revelation principle of gradually increasing complexity. This section first presents full-support conditions under which all solution concepts we consider are outcome-equivalent and the revelation principle holds, and then proves the revelation principle for WPBE and CPPBE.

4.1 Full-Support Equilibria

We start with a simple but fundamental result: any NE outcome distribution under which no player can perfectly detect another's deviation is a canonical MSE outcome distribution.

Let $\rho^{(\sigma, \mu)} \in \Delta(X)$ denote the outcome distribution induced by strategy profile (σ, μ) . Recall that ρ_i is the projection of ρ onto \hat{H}_i^{T+1} .

Proposition 1 *For any game G , if (σ, μ) is a NE and $\text{supp } \rho_i^{(\sigma, \mu)} = \bigcup_{\sigma'_{-i} \in \Sigma_{-i}} \text{supp } \rho_i^{(\sigma_i, \sigma'_{-i}, \mu)}$ for all $i \neq 0$, then $\rho^{(\sigma, \mu)}$ is a canonical MSE outcome distribution.*

The proof relies on two lemmas. The first is the revelation principle for NE, due to Forges (1986). Since we will build on this construction, we include a complete proof in the appendix.

Lemma 2 (Forges (1986; Proposition 1)) *For any game G , if (σ, μ) is a NE then $\rho^{(\sigma, \mu)}$ is the outcome distribution of a canonical NE $(\tilde{\sigma}, \tilde{\mu})$ such that, for all $i \neq 0$, $\bigcup_{\sigma'_{-i} \in \Sigma_{-i}} \text{supp } \rho_i^{(\sigma_i, \sigma'_{-i}, \mu)} \supseteq \bigcup_{\tilde{\sigma}'_{-i} \in C(\Sigma_{-i})} \text{supp } \rho_i^{(\tilde{\sigma}_i, \tilde{\sigma}'_{-i}, \tilde{\mu})}$.*

The second lemma states that any NE outcome can be implemented in strategies where each player is sequentially rational following her own deviations.¹¹ (To be precise, by a *deviation* from σ_i we mean a report $r_{i,t}$ or action $a_{i,t}$ such that $r_{i,t} \notin \text{supp } \sigma_i(\mathfrak{h}_i^t)$ or $a_{i,t} \notin \text{supp } \sigma_i(\mathfrak{h}_i^t, r_{i,t}, m_{i,t})$, for some history \mathfrak{h}_i^t or $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t})$. A history *follows* a deviation if it is a successor of a history at which a deviation occurred.)

¹¹This is a generalization of the result in repeated games that if signals have full support then Nash equilibrium and sequential equilibrium are outcome-equivalent. See, e.g., Mailath and Samuelson (2006; Proposition 12.2.1).

Lemma 3 For any game G , if (σ, μ) is a NE then $\rho^{(\sigma, \mu)}$ is the outcome distribution of an assessment $(\hat{\sigma}, \mu, \hat{\beta})$ such that

1. $(\hat{\sigma}, \mu)$ is a NE.
2. For all $i \neq 0$, (1) and (2) are satisfied at all histories \mathfrak{h}_i^t and $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t})$ that follow a deviation from $\hat{\sigma}_i$.
3. $\hat{\beta}$ is (machine) consistent.
4. For all $i \neq 0$ and $\sigma'_{-i} \in \Sigma_{-i}$, $\rho_i^{(\hat{\sigma}_i, \sigma'_{-i}, \mu)} = \rho_i^{(\sigma_i, \sigma'_{-i}, \mu)}$.

Proof. Fix a game G and a NE (σ, μ) . Consider the auxiliary game $\hat{G}(k)$ derived from G by fixing the mediator's strategy at μ and specifying that, for every player $i \neq 0$ and every history where player i has not yet deviated from σ_i , player i must follow σ_i with probability at least $k/(k+1)$. By standard results, $\hat{G}(k)$ has a MSE $(\hat{\sigma}^k, \mu, \hat{\beta}^k)$. Let $(\hat{\sigma}, \mu, \hat{\beta}) = \lim_k (\hat{\sigma}^k, \mu, \hat{\beta}^k)$, taking a convergent subsequence if necessary. Note that, for each $i \neq 0$, $\hat{\sigma}_i$ and σ_i differ only at histories that follow a deviation by player i . By continuity, $\tilde{\sigma}_i$ is sequentially rational at all histories that follow a deviation by player i (given beliefs $\hat{\beta}_i$), and $\hat{\beta}$ is consistent. Moreover, $(\sigma'_i, \sigma_{-i}, \mu)$ and $(\sigma'_i, \hat{\sigma}_{-i}, \mu)$ induce the same outcome distribution for every $\sigma'_i \in \Sigma_i$. In particular, (σ, μ) and $(\hat{\sigma}, \mu)$ induce the same outcome distribution, and $(\hat{\sigma}, \mu)$ is a NE. ■

We can now complete the proof of Proposition 1.

Proof of Proposition 1. Fix a game G and a NE (σ, μ) . By Lemmas 2 and 3, there exists a canonical NE $(\tilde{\sigma}, \hat{\mu})$ and a canonical assessment $(\hat{\sigma}, \hat{\mu}, \hat{\beta})$ such that $\rho^{(\hat{\sigma}, \hat{\mu})} = \rho^{(\sigma, \mu)}$, points 1 through 3 of Lemma 3 hold with $\hat{\mu}$ in place of μ , and, for all $i \neq 0$,

$$\bigcup_{\sigma'_{-i} \in \Sigma_{-i}} \text{supp } \rho_i^{(\sigma_i, \sigma'_{-i}, \mu)} \supseteq \bigcup_{\tilde{\sigma}'_{-i} \in C(\Sigma_{-i})} \text{supp } \rho_i^{(\tilde{\sigma}_i, \tilde{\sigma}'_{-i}, \hat{\mu})} = \bigcup_{\hat{\sigma}'_{-i} \in C(\Sigma_{-i})} \text{supp } \rho_i^{(\hat{\sigma}_i, \hat{\sigma}'_{-i}, \hat{\mu})},$$

where the equality holds by point 4 of Lemma 3. By the full-support assumption, we have

$$\text{supp } \rho_i^{(\hat{\sigma}, \hat{\mu})} = \text{supp } \rho_i^{(\sigma, \mu)} = \bigcup_{\sigma'_{-i} \in \Sigma_{-i}} \text{supp } \rho_i^{(\sigma_i, \sigma'_{-i}, \mu)} \supseteq \bigcup_{\hat{\sigma}'_{-i} \in C(\Sigma_{-i})} \text{supp } \rho_i^{(\hat{\sigma}_i, \hat{\sigma}'_{-i}, \hat{\mu})},$$

and hence $\text{supp } \rho_i^{(\hat{\sigma}, \hat{\mu})} = \bigcup_{\hat{\sigma}'_i \in C(\Sigma_{-i})} \text{supp } \rho_i^{(\hat{\sigma}_i, \hat{\sigma}'_i, \hat{\mu})}$. Therefore, every history for player i consistent with $\hat{\mu}$ either arises with positive probability along the equilibrium path of $(\hat{\sigma}, \hat{\mu})$ or follows a deviation from $\hat{\sigma}_i$. Hence, $\hat{\sigma}_i$ is sequentially rational at all histories, and therefore $(\hat{\sigma}, \hat{\mu}, \hat{\beta})$ is a MSE. ■

We mention two useful corollaries of Proposition 1.

First, the revelation principle holds in single-agent settings. This result is applicable to many models of dynamic moral hazard (e.g., Garrett and Pavan, 2012) and dynamic information design (e.g., Ely, 2017).

Corollary 1 *If $N = 1$, then any NE outcome distribution is a canonical MSE outcome distribution.*

Proof. If $N = 1$, the condition $\text{supp } \rho_i^{(\sigma, \mu)} = \bigcup_{\sigma'_i \in \Sigma_{-i}} \text{supp } \rho_i^{(\sigma_i, \sigma'_i, \mu)}$ is vacuous. ■

Second, say a game is one of *pure adverse selection* if $|A_i| = 1$ for all $i \neq 0$. Thus, in a pure adverse selection game, players report types to the mediator, the mediator chooses allocations, and players take no further actions. Much of the dynamic mechanism design literature assumes pure adverse selection (e.g., Pavan, Segal, and Toikka (2014) and references therein). We show that a virtual-implementation version of the revelation principle holds for pure adverse selection games. In particular, the distinction between Nash, perfect Bayesian, and sequential equilibrium is essentially immaterial in these games.

Let $\|\cdot\|$ denote the sup norm on $\Delta(X)$: for distributions $\rho, \rho' \in \Delta(X)$, $\|\rho - \rho'\| = \max_{\hat{h}^{T+1} \in X} \left| \rho(\hat{h}^{T+1}) - \rho'(\hat{h}^{T+1}) \right|$.

Corollary 2 *Let G be a game of pure adverse selection. For any NE outcome distribution $\rho \in \Delta(X)$ and any $\varepsilon > 0$, there exists a canonical MSE outcome distribution $\rho' \in \Delta(X)$ with $\|\rho - \rho'\| < \varepsilon$.*

Proof. Let (σ, μ) be a NE in G that induces distribution ρ . Define a strategy $\tilde{\mu}$ for the mediator in G as follows: Messages are given by $\tilde{\mu}^M(h_0^t, s_{0,t}, r_t) = \mu^M(h_0^t, s_{0,t}, r_t)$. As for actions, at the beginning of the game, with probability ε the mediator selects a path of actions $\prod_{t=0}^T a_{0,t} \in \prod_{t=0}^T A_{0,t}$ uniformly at random and deterministically follows this path of actions for the remainder of the game. With probability $1 - \varepsilon$, actions in every period are

given by $\mu^A(h_0^t, s_{0,t}, r_t, m_t)$. When the mediator follows strategy $\tilde{\mu}$, a player's reports are irrelevant in the event that the mediator follows a deterministic path of actions, so players can condition on the event that the mediator follows μ when making reports. The fact that (σ, μ) is a NE of G thus implies that $(\sigma, \tilde{\mu})$ is also a NE of G . In addition, $\|\rho - \rho^{(\sigma, \tilde{\mu})}\| < \varepsilon$. Finally, since only the mediator takes actions and $S_{i,t} = \bigcup_{\hat{h}^t \in \hat{H}^t} \text{supp } p_i(\cdot | \hat{h}^t)$ for all i, t , $\text{supp } \rho_i^{(\sigma', \tilde{\mu})} = \hat{H}_i^{T+1}$ for all $i \neq 0$ and $\sigma' \in \Sigma$. The result now follows from Proposition 1. ■

Note that Example 1 and Corollary 2 imply that the set of outcome distributions that are implementable in a canonical MSE is not compact.

4.2 PBE

We first record the revelation principle for WPBE.

Proposition 2 *The revelation principle holds for WPBE.*

The proof is a fairly straightforward extension of the proof of Lemma 2 and is deferred to the appendix. The required construction combines the strategy profile constructed in the proof of Lemma 2 with the belief system that results from projecting beliefs in the original equilibrium onto the set of payoff-relevant histories.

Our next result formalizes the revelation principle for CPPBE implicit in Myerson (1986). This is more subtle than the corresponding result for NE or WPBE. The difficulty is ensuring that beliefs in the canonical game are derived from a CPS. In particular, Myerson's Theorem 1 shows that every CPS is the limit of completely mixed distributions over moves, but these distributions need not be strategy profiles, in that move probabilities can differ at nodes in the same information set and players' trembles can be correlated. To prove the revelation principle, we must translate these correlated trembles in the original game to the canonical game. Intuitively, this is achieved by (i) insisting that players report truthfully with probability 1, (ii) delegating responsibility for all trembles in communication to the mediator, and (iii) translating players' action trembles as a function of messages in the original game to action trembles as a function of recommendations in the canonical game. (Note that in general action trembles cannot be attributed to the mediator: for example, if a player is seen to have played a dominated action, she must have deviated from her

equilibrium strategy. In contrast, trembles in communication can always be attributed to the mediator, as no player observes another's report.)

Proposition 3 *The revelation principle holds for CPPBE.*

Proof. Fix a game G and a CPPBE (σ, μ, β) . Let f denote the corresponding CPS on H^{T+1} . Let us call a tuple $\alpha = (\alpha_t^R, \alpha_t^M, \alpha_t^A)_{t=1}^T$, where $\alpha_t^R : H^t \times S_t \rightarrow \Delta(R_t)$, $\alpha_t^M : H^t \times S_t \times R_t \rightarrow \Delta(M_t)$, and $\alpha_t^A : H^t \times S_t \times R_t \times M_t \rightarrow \Delta(A_t)$, a *move distribution*. A move distribution is a more general object than a strategy profile, as it allows for the possibility that moves may be correlated or may not respect the information structure of the game. Note that every CPS on H^{T+1} induces a move distribution. Let α be the move distribution induced by f .

We construct a mediation range \tilde{Q} and a canonical assessment $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta})$ in $C(G)|_{\tilde{Q}}$. Players are truthful and obedient whenever they have been truthful in the past and receive recommendations in the mediation range. The mediator's strategy is constructed as follows:

In period 1, if the vector of reports $\hat{s}_1 := (s_{0,1}, \hat{s}_{i,1})_{i \neq 0}$ is in $\text{supp } p(\cdot|\emptyset)$, then the mediator draws a vector of fictitious reports $r_1 \in R_1$ according to $\alpha_1^R(\hat{s}_1)$. Given the resulting vector r_1 , the mediator draws a vector of fictitious messages $m_1 \in M_1$ according to $\alpha_1^M(\hat{s}_1, r_1)$.

If instead $\hat{s}_1 \notin \text{supp } p(\cdot|\emptyset)$, then for each $i \neq 0$ the mediator draws $r_{i,1}$ according to $\sigma_i^R(\hat{s}_{i,1})$. Given the resulting vector $r_1 = (r_{i,1})_{i \neq 0}$, the mediator draws $m_1 \in M_1$ according to $\mu(s_{0,1}, r_1)$.

Next, given $\hat{s}_{i,1}, r_{i,1}, m_{i,1}$, the mediator draws $b_{i,1}$ according to $\sigma_i^A(\hat{s}_{i,1}, r_{i,1}, m_{i,1})$. Finally, the mediator sends message $b_{i,1}$ to player i .

Inductively, given $h_{i,0}^t = (\hat{a}_{i,\tau-1}, \hat{s}_{i,\tau}, r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}$ and $\hat{r}_{i,t} = (\hat{a}_{i,t-1}, \hat{s}_{i,t})$ for all $i \neq 0$, if $\hat{h}^t = (\hat{a}_{\tau-1}, \hat{s}_{\tau})_{\tau=1}^t$ is a possible history then the mediator draws $r_t \in R_t$ according to $\alpha_t^R(h_{I,0}^t, a_{0,t-1}, s_{0,t}, \hat{r}_t)$, where

$$h_{I,0}^t := \prod_{\tau=1}^{t-1} \left(\left(a_{0,\tau-1}, \prod_{i=1}^N \hat{a}_{i,\tau-1} \right), \left(s_{0,\tau}, \prod_{i=1}^N \hat{s}_{i,\tau} \right), \prod_{i=1}^N r_{i,\tau}, \prod_{i=1}^N m_{i,\tau} \right).$$

Given r_t , the mediator draws $m_t \in M_t$ according to $\alpha_t^M(h_{I,0}^t, a_{0,t-1}, s_{0,t}, \hat{r}_t, r_t)$.

If instead \hat{h}^t is not a possible history, then for each $i \neq 0$ the mediator draws $r_{i,t} \in R_{i,t}$ according to $\sigma_i^R(h_{i,0}^t, \hat{r}_{i,t})$. Given the resulting vector $r_t = (r_{i,t})_{i \neq 0}$, the mediator draws

$m_t \in M_t$ according to $\mu \left((s_{0,\tau}, r_\tau, m_\tau, a_{0,\tau})_{\tau=1}^{t-1}, s_{0,t}, r_t \right)$.

Next, given $\mathfrak{h}_{i,0}^t = (h_{i,0}^t, \hat{a}_{i,t-1}, \hat{s}_{i,t}), r_{i,t}, m_{i,t}$, the mediator draws $b_{i,t}$ according to $\sigma_i^A(\mathfrak{h}_{i,0}^t, r_{i,t}, m_{i,t})$. (If $\hat{\mathfrak{h}}^t$ is not a possible history, the mediator draws $b_{i,t}$ randomly.) The mediator sends message $b_{i,t}$ to player i .

Given this construction of $\tilde{\mu}$, let \tilde{Q} be the mediation range that restricts the mediator to sending recommendations in $\text{supp } \tilde{\mu}$ (see the proof of Proposition 2).

It remains to construct a belief system $\tilde{\beta}$ in $C(G)|_{\tilde{Q}}$ and verify that $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta})$ is a restricted CPPBE. Say that a move distribution α is *completely mixed* if $\alpha_t^R(h^t, s_t), \alpha_t^M(h^t, s_t, r_t)$, and $\alpha_t^A(h^t, s_t, r_t, m_t)$ are everywhere in the interior of $\Delta(R_t), \Delta(M_t)$, and $\Delta(A_t)$, respectively. It is clear that every completely mixed move distribution α induces a CPS $F(\alpha)$ on H^{T+1} , and Theorem 1 of Myerson (1986) shows that every CPS on H^{T+1} is the limit (pointwise over conditional probabilities) of CPS's induced by completely mixed move distributions. Let (α^k) be a sequence of completely mixed move distributions such that $\lim_k F(\alpha^k) = f$.

To construct a belief system $\tilde{\beta}$, we begin by constructing a sequence of move distributions $\tilde{\alpha}^k$ as follows:

First, construct a strategy for the mediator $\tilde{\mu}^k$ by following the construction of μ^k while everywhere replacing α^R and α^M with $\alpha^{R,k}$ and $\alpha^{M,k}$. Note that, for every fictitious history $(h_0^t, s_{0,t}, r_t)$ (where $h_0^t = (s_{0,\tau}, r_\tau, m_\tau, a_{0,\tau})_{\tau=1}^{t-1}$ and $r_\tau = \prod_{i=1}^N (a_{i,\tau-1}, s_{i,\tau})$), if $b_{i,t} \in \tilde{Q}_i((a_{i,\tau-1}, s_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, a_{i,t-1}, s_{i,t})$ then $\tilde{\mu}_i^k(b_{i,t}|h_0^t, s_{0,t}, r_t) > 0$. That is, every possible message—each message in the mediation range given the history—is received with positive probability at every fictitious history.

Second, given a fictitious history $(h_0^t, s_{0,t}, r_t, m_t)$, let

$$d^k = \frac{1}{2} \min \left\{ \min_{b_t \in A_t} \alpha^{A,k}(b_t|h_0^t, s_{0,t}, r_t, m_t), \frac{1}{k} \right\}.$$

Note that $d^k > 0$ for all k , $\lim_k d^k = 0$, and $d^k < \alpha^{A,k}(b_t|h_0^t, s_{0,t}, r_t, m_t)$ for all $b_t \in A_t$. For every vector of recommendations $b_t = (b_{i,t})_{i \neq 0}$ with $\alpha^A(b_t|h_0^t, s_{0,t}, r_t, m_t) > 0$, let

$$\phi^k(a_t|h_0^t, s_{0,t}, r_t, m_t, b_t) = \begin{cases} \min \left\{ \frac{\alpha^{A,k}(b_t|h_0^t, s_{0,t}, r_t, m_t)}{\alpha^A(b_t|h_0^t, s_{0,t}, r_t, m_t)}, 1 \right\} - \frac{d^k}{\alpha^A(b_t|h_0^t, s_{0,t}, r_t, m_t)} & \text{if } a_t = b_t \\ \frac{\max\{\alpha^{A,k}(a_t|h_0^t, s_{0,t}, r_t, m_t) - \alpha^A(a_t|h_0^t, s_{0,t}, r_t, m_t), 0\}}{\sum_{a'_t \in A_t} \max\{\alpha^{A,k}(a'_t|h_0^t, s_{0,t}, r_t, m_t) - \alpha^A(a'_t|h_0^t, s_{0,t}, r_t, m_t), 0\}} \max \left\{ 1 - \frac{\alpha^{A,k}(b_t|h_0^t, s_{0,t}, r_t, m_t)}{\alpha^A(b_t|h_0^t, s_{0,t}, r_t, m_t)}, 0 \right\} \\ \quad + \frac{1}{|A_t|-1} \frac{d^k}{\alpha^A(b_t|h_0^t, s_{0,t}, r_t, m_t)} & \text{if } a_t \neq b_t \end{cases}.$$

Note that $\phi^k(\cdot|h_0^t, s_{0,t}, r_t, m_t, b_t)$ is a full-support probability distribution on A_t , $\lim_k \phi^k(b_t|h_0^t, s_{0,t}, r_t, m_t, b_t) = 1$, and

$$\sum_{b_t \in A_t} \phi^k(a_t|h_0^t, s_{0,t}, r_t, m_t, b_t) \alpha^A(b_t|h_0^t, s_{0,t}, r_t, m_t) = \alpha^{A,k}(a_t|h_0^t, s_{0,t}, r_t, m_t) \text{ for all } a_t \in A_t.$$

The interpretation is that, if the players “intend” to play each action profile $b_t \in A_t$ with probability $\alpha^A(b_t|h_0^t, s_{0,t}, r_t, m_t)$ but tremble from b_t to a_t with probability $\phi^k(a_t|h_0^t, s_{0,t}, r_t, m_t, b_t)$, then each action profile $a_t \in A_t$ ends up being played with probability $\alpha^{A,k}(a_t|h_0^t, s_{0,t}, r_t, m_t)$.

Third, let

$$\tilde{\alpha}^{A,k}(a_t|\hat{h}^t, \hat{s}_t, b_t) = \sum_{h_0^t, s_{0,t}, r_t, m_t} \Pr^{\alpha^k}(h_0^t, s_{0,t}, r_t, m_t|\hat{h}^t, \hat{s}_t, b_t) \phi^k(a_t|h_0^t, s_{0,t}, r_t, m_t, b_t).$$

Fourth, let $\tilde{\alpha}^k$ be the move distribution that results when the mediator follows strategy $\tilde{\mu}^k$, players report truthfully, and players take actions according to $\tilde{\alpha}^{A,k}(\hat{h}^t, \hat{s}_t, b_t)$. Note that $\tilde{\alpha}^k$ is completely mixed in the game where players are restricted to report truthfully. Fifth, define a CPS \tilde{f} in the restricted game by $\tilde{f} := \lim_k F(\tilde{\alpha}^k)$. Finally, let $\tilde{\beta}$ be the belief system induced by \tilde{f} in the unrestricted game, where a player’s beliefs after she has lied to the mediator are arbitrary.

We claim that $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta})$ is a restricted CPPBE. We first argue that honesty (resp., obedience) is optimal in $C(G)$ for any fictitious history $\mathfrak{h}_{i,0}^t$ (resp., $(\mathfrak{h}_{i,0}^t, r_{i,t}, m_{i,t})$), for a player who has been truthful in the past. To see this, it is straightforward to check the following claims: (i) If a previously truthful player i has a profitable misreport in $C(G)$ when the

mediator's history is $\mathfrak{h}_{i,0}^t$, then there is a profitable deviation in G in which she misreports at history $\mathfrak{h}_i^t = \mathfrak{h}_{i,0}^t$. (ii) If a previously truthful player i has a profitable deviation in $C(G)$ when the mediator's history is $(\mathfrak{h}_{i,0}^t, r_{i,t}, m_{i,t})$ and she receives recommendation $b_{i,t} \in \text{supp } \tilde{\mu}_i^k(\mathfrak{h}_{i,0}^t, r_{i,t}, m_{i,t})$, then there is a profitable deviation in G in which she deviates at history $(\mathfrak{h}_{i,0}^t, r_{i,t}, m_{i,t})$. Hence, the fact that (σ, μ, β) is sequentially rational in G implies that a player would not have a profitable deviation under $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta})$ in $C(G)$ even if she knew the mediator's fictitious history. As a player has less information than this when deciding to misreport or disobey the mediator's recommendation in $C(G)$, it follows that $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta})$ is sequentially rational in $C(G)$. Finally, by construction, $\tilde{\beta}$ is derived from a CPS and truthful players believe their opponents have also been truthful. ■

Myerson refers to an outcome distribution that arises as a CPPBE in the canonical game $C(G)$ as a *sequential communication equilibrium (SCE)* of G . Thus, Proposition 3 says that any outcome distribution that arises as a CPPBE in any game G is a SCE. Myerson shows that the set of SCE has a remarkably tractable structure: it equals the set of communication equilibria (Forges, 1986) that never assign positive probability to a codominated action.

5 Outcome-Equivalence of PSE and SCE

Our main result is the following:

Theorem 1 *For any base game, the set of outcome distributions that arise in PSE for some message set coincides with the set of SCE.*

The problem of characterizing the set of PSE outcomes is thus equivalent to characterizing SCE, which by Myerson (1986) is in turn equivalent to characterizing static communication equilibria along with the set of codominated actions. It is therefore possible to combine the tractability of Myerson's characterization with the appealing belief restrictions of PSE. In contrast, if one requires the more demanding restrictions of MSE, the equivalence with SCE is lost: Example 3 of Gerardi and Myerson (2007) shows that for some games there exist SCE that cannot be implemented in MSE for any message set.

In terms of the revelation principle, Theorem 1 shows that, if one wants to know what outcome distributions are implementable in PSE, it is without loss of generality to restrict

attention to canonical games and use the weaker solution concept of CPPBE (which admits a tractable characterization). Theorem 1 thus establishes the substance of the revelation principle for PSE. However, Theorem 1 does not resolve the more technical question of whether every PSE outcome distribution can be implemented in a canonical PSE. In particular, the proof proceeds by constructing a single message set M^* such that any SCE can be implemented in a PSE where players report their actions and signals truthfully and the mediator sends messages in M^* . Furthermore, M^* embeds the canonical message set $M = A$, and along the equilibrium path the mediator recommends actions and players are obedient. However, at certain off-path histories, the mediator instead sends a special message \star , which is in effect a signal for the players to tremble with high probability. The mediator also sends additional information at some off-path histories. Whether these extra messages can be dispensed with is an open question, albeit one that is not relevant for characterizing the set of PSE outcomes.

Theorem 1 does not claim an equivalence between the sets of CPPBE and PSE assessments. No such equivalence holds. To see this, consider a game where players 1 and 2 act simultaneously and each have a strictly dominated action a^* , and player 3 then observes a binary signal $s \in \{A, B\}$, where $s = A$ if and only if at least one of players 1 and 2 played a^* . In CPPBE, after observing $s = A$ player 3 can believe with probability 1 that both players 1 and 2 played a^* , as CPPBE allows players' deviations from equilibrium to be correlated. In contrast, in PSE, after observing $s = A$ player 3 must believe with probability 1 that exactly one of players 1 and 2 played a^* ; this follows because players tremble independently conditional on their information sets in PSE, and players 1 and 2 both play a^* with equilibrium probability 0 at every information set. This example shows that CPPBE and PSE are distinct solution concepts, and that the equivalence in Theorem 1 holds only in terms of implementable outcome distributions.

Let us briefly preview the proof of Theorem 1. One direction is immediate: Every PSE induces a CPS by Myerson's Theorem 1, so any PSE outcome distribution is a CPPBE outcome distribution. By Proposition 3, such a distribution is a SCE.

The converse is much more subtle. Fix a canonical game G and a SCE ρ . Let (σ, μ, β) be a corresponding canonical CPPBE, and let f be the corresponding CPS. To prove the theo-

rem, it would suffice to construct a sequence of completely mixed strategy profiles (σ^k, μ^k) converging to (σ, μ) such that the corresponding belief system β^k converges to β . By Myerson’s Theorem 1, there exists a sequence of completely mixed move distributions (α^k) such that $\lim_k F(\alpha^k) = f$. But, as we have emphasized, the α^k ’s are not strategy profiles, and there is no simple way to translate the move distribution sequence (α^k) into an appropriate strategy profile sequence (σ^k, μ^k) . Indeed, such a translation is not always possible: as we have seen, the equivalence between CPPBE and PSE holds only at the level of outcomes, not assessments.

The proof approach is thus to implement ρ via a rather different construction. The basic idea is to have the mediator initially set out to implement ρ , while specifying that, if a player observes a 0-probability event, the player believes that the mediator has trembled and is now issuing recommendations according to an arbitrary PSE distribution π . Assuming that the marginal distributions of ρ and π on each player’s actions have the same support, endowing players with this belief allows the mediator to enforce the target strategy profile (σ, μ) .

In particular, the following simple proof is valid under the assumption that the set of enforceable actions at every history is the same in CPPBE and PSE. Fix an arbitrary PSE distribution π with the same support as ρ , and assume that at the beginning of the game the mediator’s “state” (that is, his target outcome distribution) is either ρ or π . The state is perfectly persistent and equals ρ with (limit) probability 1, but can equal π as the result of a relatively likely tremble—say a tremble of probability $1/k$ along a sequence $(\tilde{\sigma}^k, \tilde{\mu}^k)$ converging to (σ, μ) . If the mediator’s state is π , he informs each player of this fact with probability 1, but he can again tremble by failing to inform a player:¹² along the sequence $(\tilde{\sigma}^k, \tilde{\mu}^k)$, he informs each player with probability $1 - 1/k$, independently across players. Subsequently, the mediator issues action recommendations according to his state, and players follow their recommendations but tremble with much higher probability if they have been informed that the state is π : players tremble with probability $1/k$ if they have been informed that the state is π and tremble with probability $1/k^K$ otherwise, for K large. Now, as the game progresses in state ρ , each player believes that the state is ρ with probability 1 until she observes an off-path signal. At that point, she suddenly starts believing that with probability

¹²This information that the state is π plays roughly the same role as the extra message \star .

1 (i) the state is π , (ii) the mediator failed to inform her of this fact, and (iii) another player has trembled. She is then content to continue to follow the mediator’s recommendations for the remainder of the game, believing that these recommendations are being issued according to a PSE supporting distribution π (and this belief is never subsequently disconfirmed, by the assumption that ρ and π have the same support). Finally, a player’s willingness to behave in this way following an off-path signal implies that the mediator can enforce the desired strategy profile (σ, μ) .

Unfortunately, we cannot rely on this simple argument, because we do not know a priori that the set of enforceable actions is the same in CPPBE and PSE. Suppose that the set of enforceable actions were strictly larger for CPPBE, and that supporting the distribution ρ requires recommending an off-path action a_i that is never played in any PSE. Then a player who observes an off-path signal and is then recommended action a_i will not follow this recommendation if she believes that the mediator is following the PSE supporting π . This difficulty forces us to rely on a more complicated construction, where—for example—a player who observes an off-path signal switches from believing that the state is ρ to believing it is π , but then switches *back* to believing that the state is ρ if she receives an action recommendation outside the support of the PSE supporting π . The need to control players’ beliefs in this construction requires in turn that the likelihood ratio between the “main state” ρ and the “auxiliary state” π is carefully controlled throughout the proof. Finally, this problem is simplified somewhat by constructing not just a single auxiliary state into which the mediator can switch at the beginning of the game, but a different auxiliary state for each period t , where in every period the mediator can either remain in the main state or switch to the current period’s auxiliary state.

6 Proof of Theorem 1

6.1 Preliminaries

Fix a canonical game G and a SCE ρ . Let (σ, μ, β) be a corresponding canonical CPPBE, let $\alpha^A = (\alpha^A)_{t=1}^T$ be the corresponding action distribution (where $\alpha_t^A : H^t \times S_t \times R_t \times M_t \rightarrow$

$\Delta(A_t)$), and let f be the corresponding CPS on H^{T+1} . By Proposition 3, there exists a sequence of completely mixed action distributions $\tilde{\alpha}^{A,k}$ and a sequence of strategies for the mediator $\tilde{\mu}^k$ such that, letting $\tilde{\alpha}^k$ be the move distribution that results when the mediator follows strategy $\tilde{\mu}^k$, players report truthfully, and players take actions according to $\tilde{\alpha}^{A,k}$, we have $F(\tilde{\alpha}^k) \rightarrow f$ on all events where players have reported truthfully.

Our goal is to construct a sequence of completely mixed strategy profiles (σ^k, μ^k) with corresponding belief system β^k such that $(\sigma^k, \mu^k, \beta^k)$ converges to a PSE $(\sigma^*, \mu^*, \beta^*)$ that generates outcome distribution ρ . The construction of (σ^k, μ^k) begins by specifying a collection of sequences of tremble probabilities $\left\{f_1(k), f_2(k), \bar{f}_3(k), \underline{f}_3(k), f_4(k)\right\}_{k=1}^{\infty}$. This specification proceeds in two steps.

First, if ρ has full support, then by Proposition 1 ρ arises in an MSE, and hence in a PSE. So assume that ρ does not have full support, and let

$$\begin{aligned}\underline{f}_3(k) &= \min_{E', E'' \subseteq E \subseteq H^{T+1}: \lim_k \frac{F(\tilde{\alpha}^k)(E'|E)}{F(\tilde{\alpha}^k)(E''|E)} = 0} \frac{F(\tilde{\alpha}^k)(E'|E)}{F(\tilde{\alpha}^k)(E''|E)} \\ \bar{f}_3(k) &= \max_{E', E'' \subseteq E \subseteq H^{T+1}: \lim_k \frac{F(\tilde{\alpha}^k)(E'|E)}{F(\tilde{\alpha}^k)(E''|E)} = 0} \frac{F(\tilde{\alpha}^k)(E'|E)}{F(\tilde{\alpha}^k)(E''|E)}\end{aligned}$$

be the smallest and largest likelihood ratio between any two events E', E'' whose likelihood ratio converges to 0 conditional on some event E . Note that $\underline{f}_3(k)$ and $\bar{f}_3(k)$ are strictly positive (as ρ does not have full support), $\underline{f}_3(k) \leq \bar{f}_3(k)$ for all k , and $\lim_k \underline{f}_3(k) = \lim_k \bar{f}_3(k) = 0$.

Second, given $\left\{\underline{f}_3(k), \bar{f}_3(k)\right\}$, let $\left\{f_1(k), f_2(k), f_4(k)\right\}_{k=1}^{\infty}$ be arbitrary sequences that all converge to 0 as $k \rightarrow \infty$ and satisfy the relationship

$$f_1(k) \gg f_2(k) \gg \bar{f}_3(k) \geq \underline{f}_3(k) \gg f_4(k), \quad (3)$$

where $f_{n-1}(k) \gg f_n(k)$ means $\lim_k f_n(k) (f_{n-1}(k))^{-2(N+1)T} = 0$.

6.2 Auxiliary Equilibria

As discussed above, the construction of (σ^k, μ^k) makes use of T distinct auxiliary sequential equilibria, denoted $(\psi[t], \beta^\psi[t])_{t=1}^T$. Intuitively, in each period t the mediator will either follow the original equilibrium (supporting outcome distribution ρ) or switch to following the auxiliary equilibrium $(\psi[t], \beta^\psi[t])$ (if the mediator has not already switched to following a different auxiliary equilibrium $(\psi[\tau], \beta^\psi[\tau])$ in some period $\tau \leq t - 1$).

To construct $(\psi[t], \beta^\psi[t])$, first define a full support probability distribution $F^k[t] \in \Delta(H^t \times S_t \times A_t)$ by $F^k[t](h^t, s_t, \tilde{a}_t) = F(\tilde{a}^k)(h^t, s_t, \tilde{a}_t)$ for each $(h^t, s_t, \tilde{a}_t) \in (H^t \times S_t \times A_t)$, where H^t is the set of period t histories in the canonical game. Now define an auxiliary game $G^k[t]$ as follows: (i) Nature draws an “initial state” $(h^t, s_t, \tilde{a}_t) \in H^t \times S_t \times A_t$ according to $F^k[t]$. (ii) Player i observes $(h_i^t, s_{i,t}, \tilde{a}_{i,t})$. (iii) Play proceeds as in the original game G , starting at history (h^t, s_t, \tilde{a}_t) . (iv) The mediator mechanically sends a fixed message $m_{i,\tau} = \star$ for all i and all $\tau \geq t$, where \star denotes a new, “null” message outside the original message set $M = A$. (That is, in game $G^k[t]$, all meaningful messages from the mediator are ruled out.) (v) In every period $\tau \geq t$, each player i is constrained to play every action $a_{i,\tau} \in A_{i,\tau}$ with probability at least $f_1(k)$. (Recall that $f_1(k)$ is an arbitrary function satisfying $f_1(k) \rightarrow 0$ and (3).) The game $G^k[t]$ admits a sequential equilibrium $(\psi^k[t], \beta^{\psi,k}[t])$ by standard arguments; choose one arbitrarily. Letting $(\psi[t], \beta^\psi[t]) = \lim_k (\psi^k[t], \beta^{\psi,k}[t])$ (taking a convergent subsequence if necessary), it follows from standard arguments that $(\psi[t], \beta^\psi[t])$ is a sequential equilibrium of the game $G^\infty[t]$ where the initial state is distributed according to $\lim_k F^k[t]$ but nature can tremble to any initial state in $H^t \times S_t \times A_t$.

6.3 Construction of (σ^k, μ^k)

We now construct (σ^k, μ^k) . To construct the mediator’s strategy, in each period t we will specify for the mediator a “provisional state” $\hat{\theta}_t \in \{\rho, \pi[1], \dots, \pi[t-1]\}$, a “final state” $\theta_t \in \{\rho, \pi[1], \dots, \pi[t-1]\}$, and a “provisional recommendation” $\tilde{a}_t \in A_t$. These are all artificial variables used only to construct the mediator’s strategy and differ from the mediator’s actual messages and actions. Intuitively, being in state ρ means that the mediator is following the original equilibrium, while being in state $\pi[\tau]$ means that the mediator switched to following

an auxiliary equilibrium in period τ . The difference between the provisional and final states is explained below.

The message set used in the construction is given by $R_{i,t} = A_{i,t-1} \times S_{i,t}$ and $M_{i,t} = (A_{i,t} \cup \{\star\}) \times \{\rho, \pi\} \times \{0, 1, \dots, t-1\}$ for all i, t . Thus, the set of reports to the mediator is canonical, while the set of messages from the mediator is augmented to allow the mediator to send message \star rather than an action recommendation and to convey information about whether and in what period the state switched from ρ to π .

Say that the mediator's "fictitious history" is the concatenation of his actual history in the game and the history of the artificial states and recommendations, so that the set of all period t fictitious histories is

$$\underbrace{(A_{\tau-1}, S_\tau)}_{\text{reports}}, \underbrace{M_\tau}_{\text{messages}}, \underbrace{\{\rho, \bar{\pi}[1], \dots, \bar{\pi}[\tau]\}}_{\text{provisional state}}, \underbrace{\{\rho, \bar{\pi}[1], \dots, \bar{\pi}[\tau]\}}_{\text{final state}}, \underbrace{A_\tau}_{\text{provisional recommendation}})_{\tau=1}^{t-1}.^{13}$$

In a slight abuse of notation, for the remainder of the proof we let H_0^t denote this set of fictitious histories, rather than (as usual) the smaller set $(A_{\tau-1}, S_\tau, M_\tau)_{\tau=1}^{t-1}$. We also let $\mathfrak{h}_0^t = (h_0^t, a_{t-1}, s_t)$, where $h_0^t \in H_0^t$ is a fictitious history, $(a_{i,t-1}, s_{i,t})$ is player i 's period t report, $a_{t-1} = (a_{0,t-1}, (a_{i,t})_{i \neq 0})$, and $s_t = (s_{0,t}, (s_{i,t})_{i \neq 0})$. Finally, given a fictitious history $h_0^t = (a_{\tau-1}, s_\tau, m_\tau, \hat{\theta}_\tau, \theta_\tau, \tilde{a}_\tau)_{\tau=1}^{t-1}$, we let $\hat{h}_0^t = (a_{\tau-1}, s_\tau, \tilde{a}_\tau)_{\tau=1}^{t-1}$ be the history in the canonical game where signals and reports are as in h_0^t but the mediator's messages are given by \tilde{a}_τ rather than m_τ . Similarly, let $\hat{\mathfrak{h}}_0^t = (\hat{h}_0^t, a_{t-1}, s_t)$.

6.3.1 Mediator's State and Provisional Recommendation: $\hat{\theta}_t$, θ_t , and \tilde{a}_t

We first define $\hat{\theta}_t$ and θ_t , the mediator's provisional and final state in period t . Intuitively, the mediator's action recommendations depend only on the final state, but both the final state and the provisional state affect the mediator's state transition.

The mediator's provisional state in period 1 equals ρ . The mediator's final state in period 1 equals ρ with probability $1 - f_2(k)$ and equals $\pi[1]$ with probability $f_2(k)$.

Recursively, given the mediator's provisional and final states in period t and the medi-

¹³In the expression, the mediator's period $\tau - 1$ action and period τ signal are included in the term labeled "reports."

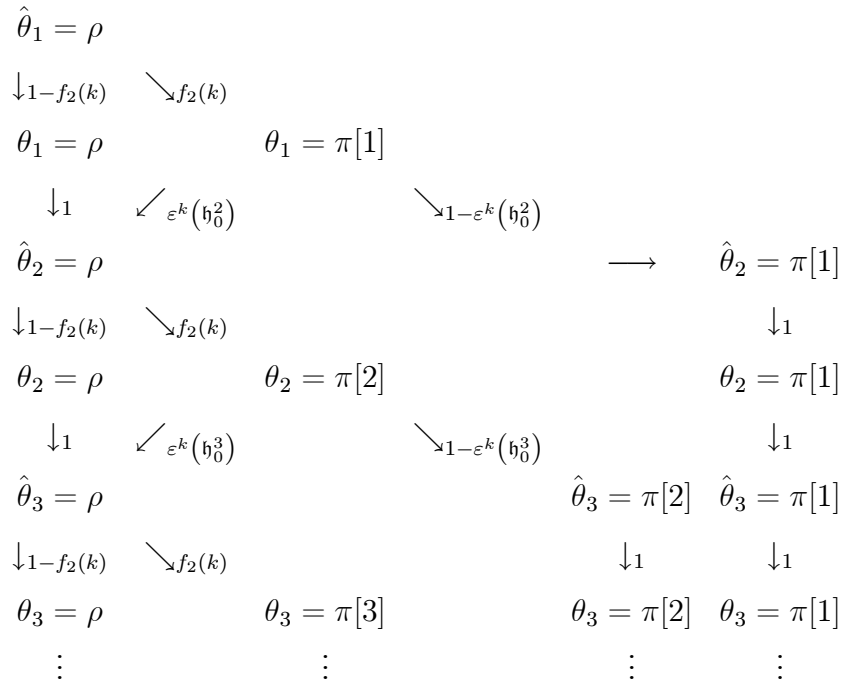
ator's fictitious history \mathfrak{h}_0^{t+1} , the mediator's provisional state in period $t + 1$ is defined as follows:

1. If $\theta_t = \rho$, then $\hat{\theta}_{t+1} = \rho$.
2. If $\theta_t = \pi[t]$, then
 - (a) If $\hat{\theta}_t = \rho$, then $\hat{\theta}_{t+1} = \rho$ with probability $\varepsilon^k (\mathfrak{h}_0^{t+1})$ and $\hat{\theta}_{t+1} = \pi[t]$ with probability $1 - \varepsilon^k (\mathfrak{h}_0^{t+1})$, where the transition probability $\varepsilon^k (\mathfrak{h}_0^{t+1})$ is determined below.
 - (b) If $\hat{\theta}_t = \pi[\tau]$ for some $\tau \leq t - 1$, then $\hat{\theta}_{t+1} = \pi[\tau]$.

Next, given the mediator's provisional state $\hat{\theta}_{t+1}$, the mediator's final state in period $t + 1$ is defined as follows:

1. If $\hat{\theta}_{t+1} = \rho$, then $\theta_{t+1} = \rho$ with probability $1 - f_2(k)$ and $\theta_{t+1} = \pi[t + 1]$ with probability $f_2(k)$.
2. If $\hat{\theta}_{t+1} = \pi[\tau]$ for some $\tau \leq t$, then $\theta_{t+1} = \pi[\tau]$.

The transition of the mediator's state is illustrated in the following flow-chart:



In particular, note that if the mediator's provisional state ever equals $\pi[t]$, then the mediator's provisional and final state both equal $\pi[t]$ in every subsequent period.

Finally, given his period t fictitious history \mathfrak{h}_0^t , the mediator draws the provisional recommendation $\tilde{a}_t \in A_t$ according to $\tilde{\mu}^k(\tilde{a}_t | \widehat{\mathfrak{h}}_0^t)$.

6.3.2 Mediator's Messages: $m_{i,t}^A$, $m_{i,t}^\Theta$, and $m_{i,t}^{\leq t-1}$

We next define the mediator's messages. We denote a message by $m_{i,t} = (m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1})$, where $m_{i,t}^A \in A_{i,t} \cup \{\star\}$, $m_{i,t}^\Theta \in \{\rho, \pi\}$, and $m_{i,t}^{\leq t-1} \in \{0, \dots, t-1\}$. The mediator's message always takes one of the following three forms:

1. $(m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \rho, 0)$ for some $\tilde{a}_{i,t} \in A_{i,t}$.

(As will be seen below, such a message is sent if $\theta_t = \rho$, and only if $\hat{\theta}_t = \rho$.)

2. $(m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0)$ for some $\tilde{a}_{i,t} \in A_{i,t}$.

(Such a message is sent only if $\theta_t = \pi[t]$ and $\hat{\theta}_t = \rho$.)

3. $(m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1}) = (\star, \pi, \tau)$ with $0 < \tau < t$.

(Such a message is sent if and only if $\hat{\theta}_t = \pi[\tau]$.)

Specifically, the message components $m_{i,t}^A$, $m_{i,t}^\Theta$, and $m_{i,t}^{\leq t-1}$ are determined as follows:

- $m_{i,t}^A$: Given the provisional recommendation \tilde{a}_t , we define $m_{i,t}^A = \tilde{a}_{i,t}$ if $\hat{\theta}_t = \rho$ and $m_{i,t}^A = \star$ if $\hat{\theta}_t = \pi[\tau]$ for some $\tau \leq t-1$. That is, the provisional recommendation is actually recommended, unless the mediator's provisional state is equal to $\pi[\tau]$ for some $\tau \leq t-1$ (in which case the mediator's provisional and final states are absorbed at $\pi[\tau]$, as noted above).
- $m_{i,t}^\Theta$: If $\theta_t = \rho$ then $m_{i,t}^\Theta = \rho$ for all i ; if $\theta_t = \pi[\tau]$ for some $\tau \leq t-1$ then $m_{i,t}^\Theta = \pi$ for all i ; and if $\theta_t = \pi[t]$ then, for each player i , $m_{i,t}^\Theta = \rho$ with probability $f_2(k)$ and $m_{i,t}^\Theta = \pi$ with probability $1 - f_2(k)$, independently across players. In particular, $m_{i,t}^\Theta$ is equal to the mediator's final state (except the time index) with high probability, but a player may receive message $m_{i,t}^\Theta = \rho$ when $\theta_t = \pi[t]$. Intuitively, this failure to inform

a player that the mediator's final state is $\pi[t]$ corresponds to the failure to inform a player that the mediator's state is π in the simplified proof described earlier.

- $m_{i,t}^{\leq t-1}$: If $\hat{\theta}_t = \pi[\tau]$ for some $\tau \leq t-1$ then $m_{i,t}^\ominus = \tau$ for all i ; otherwise, $m_{i,t}^\ominus = 0$ for all i . Thus, once the mediator's provisional state equals $\pi[\tau]$, the mediator informs all players of this fact. In all other cases—including the case where the final state equals $\pi[t]$ but the provisional state equals ρ —the mediator sends message $m_{i,t}^\ominus = 0$, indicating that his state has not yet been absorbed.

6.3.3 Players' Strategies

We start with some terminology. We say that player i takes an $f_4(k)$ action in period t if she receives message $(m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (a_{i,t}, \rho, 0)$ but plays $a'_{i,t} \neq a_{i,t}$. (As we will see, this event occurs with probability at most $f_4(k)$.) Player i is *faithful* in period t if she reports $(a_{i,t-1}, s_{i,t})$ truthfully and does not take an $f_4(k)$ action. A history $\mathfrak{h}_i^t = ((a_{i,\tau-1}, s_{i,\tau}, \tilde{a}_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, a_{i,t-1}, s_{i,t})$ is *faithful* if $a_{i,\tau} = \tilde{a}_{i,\tau}$ for all $\tau \leq t-1$ such that $(m_{i,\tau}^A, m_{i,\tau}^\ominus, m_{i,\tau}^{\leq \tau-1}) = (\tilde{a}_{i,\tau}, \rho, 0)$. When referring to a faithful history for player i , we implicitly assume that player i has reported truthfully so far; thus, a faithful history \mathfrak{h}_i^t may be viewed as a history for player i augmented with the auxiliary messages $(\tilde{a}_{i,\tau})_{\tau=1}^{t-1}$, where player i 's period τ report is assumed to equal $(a_{i,\tau-1}, s_{i,\tau})$ for all $\tau \leq t-1$.

At a faithful history, a player's actions (as well as the mediator's actions) are determined as follows: Player i reports truthfully with probability $1 - f_4(k)$ and send each possible misreport with probability $f_4(k) / (|A_{i,t-1}| |S_{i,t}| - 1)$. After reporting truthfully and receiving message $(m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1})$, player i 's action is defined as follows:

1. If $m_{i,t}^{\leq t-1} = 0$ and $m_{i,t}^\ominus = \rho$, play $a_{i,t} = m_{i,t}^A$ with probability $1 - f_4(k)$ and play each $a'_{i,t} \neq a_{i,t}$ with probability $f_4(k) / (|A_{i,t}| - 1)$.
2. If $m_{i,t}^{\leq t-1} = 0$ and $m_{i,t}^\ominus = \pi$, play each $a_{i,t}$ with probability $\psi_i^k[t](a_{i,t} | \hat{\mathfrak{h}}_i^t, m_{i,t}^A)$.
3. If $m_{i,t}^{\leq t-1} = \tau$ with $0 < \tau < t$ and $m_{i,\tau}^\ominus = \pi$, play each $a_{i,t}$ with probability $\psi_i^k[\tau](a_{i,t} | \hat{\mathfrak{h}}_i^\tau, m_{i,\tau}^A, a_{i,\tau}, s_{i,\tau+1}, \dots, a_{i,t-1}, s_{i,t})$.
4. If $m_{i,t}^{\leq t-1} = \tau$ with $0 < \tau < t$ and $m_{i,\tau}^\ominus = \rho$, play an arbitrary best response.

As we will see, a player who receives message $m_{i,t}^\ominus = \pi$ subsequently assigns probability 0 to the event that another player received message $m_{j,t}^\ominus = \rho$, regardless of the specification of play in Case 4. This implies that the specification of play in Case 4 does not affect players' incentives in Cases 1–3. Similarly, we will see that a player who has been faithful so far believes with probability 1 that her opponents have also been faithful. Hence, the specification of play for non-faithful players also does not affect faithful players' incentives, so we can let non-faithful players play arbitrary best responses. The existence of mutual best responses in Case 4 or after being non-faithful follows from standard results. We thus leave play at such histories unspecified.

6.3.4 Transition Probabilities: $\varepsilon^k(\mathfrak{h}_0^t)$

To complete the description of (σ^k, μ^k) , it remains to specify the state transition probability $\varepsilon^k(\mathfrak{h}_0^t)$ for each fictitious history \mathfrak{h}_0^t . The goal is to define these variables so that, conditional on the event $\hat{\theta}_t = \rho$ (or equivalently $m_t^{\leq t-1} = 0$), a previously faithful player's beliefs about \mathfrak{h}_0^t are close to her beliefs under the original move distribution $\tilde{\alpha}^k$. More precisely, we wish to ensure that, conditional on the event $\hat{\theta}_t = \rho$ and on reaching faithful history $\mathfrak{h}_i^t = ((a_{i,\tau-1}, s_{i,\tau}, \tilde{a}_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, a_{i,t-1}, s_{i,t})$, player i 's belief about $((a_{\tau-1}, s_\tau, \tilde{a}_\tau, m_{-i,\tau})_{\tau=1}^{t-1}, a_{t-1}, s_t)$ is asymptotically (as $k \rightarrow \infty$) equal to

$$F(\tilde{\alpha}^k) \left((a_{\tau-1}, s_\tau, \tilde{a}_\tau)_{\tau=1}^{t-1}, a_{t-1}, s_t \mid (a_{i,\tau-1}, s_{i,\tau}, \tilde{a}_{i,\tau})_{\tau=1}^{t-1}, a_{i,t-1}, s_{i,t} \right);$$

that is,

$$\lim_k \frac{\Pr^{\sigma^k, \mu^k} \left((a_{\tau-1}, s_\tau, \tilde{a}_\tau, m_\tau)_{\tau=1}^{t-1}, a_{t-1}, s_t \mid \mathfrak{h}_i^t, \hat{\theta}_t = \rho \right)}{F(\tilde{\alpha}^k) \left((a_{\tau-1}, s_\tau, \tilde{a}_\tau)_{\tau=1}^{t-1}, a_{t-1}, s_t \mid (a_{i,\tau-1}, s_{i,\tau}, \tilde{a}_{i,\tau})_{\tau=1}^{t-1}, a_{i,t-1}, s_{i,t} \right)} = 1$$

for each faithful history $((a_{\tau-1}, s_\tau, \tilde{a}_\tau, m_\tau)_{\tau=1}^{t-1}, a_{t-1}, s_t)$.

The current section constructs $\varepsilon^k(\mathfrak{h}_0^t)$ —completing the description of (σ^k, μ^k) —and the following section verifies the desired convergence of beliefs.

We again proceed recursively. Note that $\varepsilon^k(\mathfrak{h}_0^t)$ is defined only for $t \geq 2$. We are thus left to define $\varepsilon^k(\mathfrak{h}_0^{t+1})$ for each $t \geq 1$, given $\varepsilon^k(\mathfrak{h}_0^\tau)$ for $\tau \leq t$.

Given \mathfrak{h}_0^{t+1} and \tilde{a}_t , we define $\varepsilon^k(\mathfrak{h}_0^{t+1})$ as follows: Let $I = \{1, \dots, N\}$ and let $I_t^* = \{i \in I : m_{i,t}^\ominus = \rho\}$. Recalling that $\hat{\theta}_t = \rho$ implies $m_t^A = \tilde{a}_t$ and $m_t^{\leq t-1} = 0$, the message $(m_t^A, m_t^\ominus, m_t^{\leq t-1})$ is completely determined by \tilde{a}_t and I_t^* .

If $I_t^* = I$ then $\varepsilon^k(\mathfrak{h}_0^{t+1}) := 0$. In particular, this implies that $I_t^* = I$ and $\hat{\theta}_{t+1} = \rho$ if and only if $\theta_t = \rho$: ignoring terms of order $f_4(k)$,

$$\begin{aligned} \Pr^{\sigma^k, \mu^k} \left(I_t^* = I, a_t = \tilde{a}_t, \hat{\theta}_{t+1} = \rho | \mathfrak{h}_0^t, \tilde{a}_t, \hat{\theta}_t = \rho \right) &= \Pr^{\sigma^k, \mu^k} \left(\theta_t = \rho | \mathfrak{h}_0^t, \tilde{a}_t, \hat{\theta}_t = \rho \right) \\ &= 1 - f_2(k). \end{aligned}$$

For every $I_t^* \subsetneq I$ and every $a_t \in A_t$ such that $a_{i,t} = \tilde{a}_{i,t} = m_{i,t}^A$ for all $i \in I_t^*$ and $a_{i,t} \neq \tilde{a}_{i,t}$ for all $i \notin I_t^*$, we define $\varepsilon^k(\mathfrak{h}_0^{t+1})$ such that

$$\Pr^{\sigma^k, \mu^k} \left(I_t^*, a_t, \hat{\theta}_{t+1} = \rho | \mathfrak{h}_0^t, \tilde{a}_t, \hat{\theta}_t = \rho \right) = F(\tilde{\alpha}^k) \left(a_t | \hat{\mathfrak{h}}_0^t, \tilde{a}_t \right). \quad (4)$$

To see that this is possible, recall that, given \mathfrak{h}_0^t , \tilde{a}_t , and $\hat{\theta}_t = \rho$, (i) the mediator recommends $m_{i,t}^\ominus = \rho$ for each $i \in I_t^*$ and $m_{i,t}^\ominus = \pi$ for each $i \notin I_t^*$ with probability at least $(f_2(k))^{N+1}$, and (ii) each player i with $m_{i,t}^\ominus = \rho$ plays $\tilde{a}_{i,t}$ with probability $1 - f_4(k)$, and each player i with $m_{i,t}^\ominus = \pi$ plays each action with probability at least $f_1(k)$. Hence,

$$\Pr^{\sigma^k, \mu^k} \left(I_t^*, a_t | \mathfrak{h}_0^t, \tilde{a}_t, \hat{\theta}_t = \rho \right) \geq (f_2(k))^{N+1} (f_1(k))^{N+1}.$$

In addition, since $a_t \neq \tilde{a}_t$, we have

$$F(\tilde{\alpha}^k) \left(a_t | \hat{\mathfrak{h}}_0^t, \tilde{a}_t \right) \leq \bar{f}_3(k).$$

We can therefore define

$$\begin{aligned} \varepsilon^k(\mathfrak{h}_0^{t+1}) &= \frac{F(\tilde{\alpha}^k) \left(a_t | \hat{\mathfrak{h}}_0^t, \tilde{a}_t \right)}{\Pr^{\sigma^k, \mu^k} \left(I_t^*, a_t | \mathfrak{h}_0^t, \tilde{a}_t, \hat{\theta}_t = \rho \right)} \\ &\leq \frac{\bar{f}_3(k)}{(f_2(k))^{N+1} (f_1(k))^{N+1}} \rightarrow 0 \text{ by (3)}. \end{aligned} \quad (5)$$

For other action profiles a_t —that is, action profiles where $a_{i,t} \neq \tilde{a}_{i,t}$ for some $i \in I_t^*$, or $a_{i,t} = \tilde{a}_{i,t}$ for some $i \notin I_t^*$ —we let $\varepsilon^k(\mathfrak{h}_0^{t+1}) = 0$. This completes the definition of $\varepsilon^k(\mathfrak{h}_0^{t+1})$, and thus completes the construction of (σ^k, μ^k) .

Given this construction of μ^k , let \tilde{Q} be the mediation range that restricts the mediator to sending recommendations in $\text{supp } \mu^k$. (Note that this does not depend on k .) We henceforth view (σ^k, μ^k) as a strategy profile in game $G|_{\tilde{Q}}$.

6.4 Belief Convergence

We introduce key two lemmas. First, after receiving message $(m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \rho, 0)$, player i 's beliefs are the same (in the limit as $k \rightarrow \infty$) as in the original equilibrium:

Lemma 4 *For any t , any faithful \mathfrak{h}_i^t , and any $(m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1})$ with $m_{i,t}^A \in A_i$, $m_{i,t}^\Theta = \rho$, and $m_{i,t}^{\leq t-1} = 0$,*

1. *Player i believes that $\theta_t = \rho$ with probability 1, and her belief about $\tilde{a}_{-i,t}$ is the same as in the original equilibrium: For any faithful \mathfrak{h}_{-i}^t and \tilde{a}_t ,*

$$\begin{aligned} & \lim_k \Pr^{\sigma^k, \mu^k} \left(\begin{array}{l} \mathfrak{h}_{-i}^t, \theta_t = \rho, (m_{-i,t}^A, m_{-i,t}^\Theta, m_{-i,t}^{\leq t-1}) = (\tilde{a}_{-i,t}, \rho, 0) \\ \mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \rho, 0) \end{array} \right) \\ &= \lim_k F(\tilde{\alpha}^k) \left(\hat{\mathfrak{h}}_{-i}^t, \tilde{a}_{-i,t} | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t} \right). \end{aligned} \quad (6)$$

2. *For any $\tilde{a}_{i,t}$, $a_{i,t}$, and $s_{i,t+1}$, if player i receives $(m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \rho, 0)$, plays $a_{i,t}$, and observes $s_{i,t+1}$, then player i 's belief about θ_t is degenerate: For any faithful \mathfrak{h}_{-i}^t and $\tilde{a}_{i,t}, a_{i,t}, s_{i,t+1}$,*

$$\lim_k \Pr^{\sigma^k, \mu^k} (\theta_t = \rho | \mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \rho, 0), a_{i,t}, s_{i,t+1}) \in \{0, 1\}. \quad (7)$$

3. *For any $\tilde{a}_{i,t}$, $a_{i,t}$, and $s_{i,t+1}$, if player i receives $(m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \rho, 0)$, plays $a_{i,t}$, observes $s_{i,t+1}$, reports truthfully, and receives $(m_{i,t+1}^\Theta, m_{i,t+1}^{\leq t}) = (\rho, 0)$, then player*

i 's belief is the same as in the original equilibrium: For any faithful \mathfrak{h}_{-i}^t and $\tilde{a}_t, a_t, s_{t+1}$,

$$\begin{aligned} & \lim_k \Pr^{\sigma^k, \mu^k} \left(\begin{array}{c} \mathfrak{h}_{-i}^t, m_{-i,t}, a_{-i,t}, s_{-i,t+1} \\ |\mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \rho, 0), a_{i,t}, s_{i,t+1}, (m_{i,t+1}^\ominus, m_{i,t+1}^{\leq t}) = (\rho, 0) \end{array} \right) \\ &= \lim_k F(\tilde{\alpha}^k) \left(\hat{\mathfrak{h}}_{-i}^t, a_{-i,t}, s_{-i,t+1} | \hat{\mathfrak{h}}_i^t, a_{i,t}, s_{i,t+1} \right). \end{aligned} \quad (8)$$

Second, after receiving message $(m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0)$, player i 's beliefs are the same as in the auxiliary equilibrium:

Lemma 5 For any t , any faithful \mathfrak{h}_i^t , and any $(m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1})$ with $m_{i,t}^A \in A_i$, $m_{i,t}^\ominus = \pi$, and $m_{i,t}^{\leq t-1} = 0$,

1. Player i believes that $\theta_t = \pi$ with probability 1, and her belief about $\tilde{a}_{-i,t}$ is the same as in the auxiliary equilibrium: For any faithful \mathfrak{h}_{-i}^t and \tilde{a}_t ,

$$\begin{aligned} & \lim_k \Pr^{\sigma^k, \mu^k} \left(\begin{array}{c} \mathfrak{h}_{-i}^t, \theta_t = \pi, (m_{-i,t}^A, m_{-i,t}^\ominus, m_{-i,t}^{\leq t-1}) = (\tilde{a}_{-i,t}, \pi, 0) \\ |\mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0) \end{array} \right) \\ &= \lim_k F(\tilde{\alpha}^k) \left(\hat{\mathfrak{h}}_{-i}^t, \tilde{a}_{-i,t} | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t} \right) = \beta^\psi[t] \left(\hat{\mathfrak{h}}_{-i}^t, \tilde{a}_{-i,t} | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t} \right). \end{aligned} \quad (9)$$

2. For $\tau \geq t+1$, let $m_{i,t:\tau}(\tilde{a}_{i,t})$ denote the sequence of messages given by $(m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0)$ and $(m_{i,t+1}^A, m_{i,t+1}^\ominus, m_{i,t+1}^{\leq t}) = \dots = (m_{i,\tau}^A, m_{i,\tau}^\ominus, m_{i,\tau}^{\leq \tau-1}) = (\star, \pi, t)$. For any faithful \mathfrak{h}_{-i}^τ and $\tau \geq t+1$, $a_{i,t}, \dots, a_{i,\tau}, s_{i,t+1}, \dots, s_{i,\tau+1}$,

$$\begin{aligned} & \lim_k \Pr^{\sigma^k, \mu^k} \left(\begin{array}{c} \mathfrak{h}_{-i}^\tau, (m_{-i,t}^A, m_{-i,t}^\ominus, m_{-i,t}^{\leq t-1}) = (\tilde{a}_{-i,t}, \pi, 0) \\ |\mathfrak{h}_i^t, m_{i,t:\tau}(\tilde{a}_{i,t}), a_{i,t}, s_{i,t+1}, \dots, a_{i,\tau}, s_{i,\tau+1} \end{array} \right) \\ &= \beta^\psi[t] \left(\hat{\mathfrak{h}}_{-i}^\tau | \hat{\mathfrak{h}}_i^t, a_{i,t}, s_{i,t+1}, \dots, a_{i,\tau}, s_{i,\tau+1} \right). \end{aligned} \quad (10)$$

3. For any $\tilde{a}_{i,t}, a_{i,t}$, and $s_{i,t+1}$, if player i receives $(m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0)$, plays $a_{i,t}$, observes $s_{i,t+1}$, reports truthfully, and receives $(m_{i,t+1}^\ominus, m_{i,t+1}^{\leq t}) = (\rho, 0)$, then player

i 's belief is the same as in the original equilibrium: For any faithful \mathfrak{h}_{-i}^t and $\tilde{a}_t, a_t, s_{t+1}$,

$$\begin{aligned} & \lim_k \Pr^{\sigma^k, \mu^k} \left(\begin{array}{c} \mathfrak{h}_{-i}^t, m_{-i,t}, a_{-i,t}, s_{-i,t+1} \\ |\mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), a_{i,t}, s_{i,t+1}, (m_{i,t+1}^\Theta, m_{i,t+1}^{\leq t}) = (\rho, 0) \end{array} \right) \\ &= \lim_k F(\tilde{\alpha}^k) \left(\hat{\mathfrak{h}}_{-i}^t, a_{-i,t}, s_{-i,t+1} | \hat{\mathfrak{h}}_i^t, a_{i,t}, s_{i,t+1} \right). \end{aligned} \quad (11)$$

The proofs of these lemmas are the most technical parts of the proof and are deferred to the appendix. Note that parts 2 and 3 of Lemma 4 expresses the key idea that a player's belief can switch from a point mass on $\theta_t = \rho$ to a point mass on $\theta_t = \pi[\tau]$ for some $\tau \leq t-1$ after receiving a 0-probability signal, but then switches back to a point mass on $\theta_t = \rho$ after receiving a "normal" message of the form $(m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \rho, 0)$. We also note that it is shown in the course of proving Lemmas 4 and 5 that a previously faithful player always believes with probability 1 that the other players have been faithful, which justifies the implicit specification of strategies in Case 4 of Section 6.3.3.

6.5 Sequential Rationality

Let $(\sigma^*, \mu^*) = \lim_k (\sigma^k, \mu^k)$, taking a convergent subsequence if necessary. Clearly, (σ^*, μ^*) induces outcome distribution ρ . It remains to show that it is sequentially rational for a previously faithful player to be truthful and obedient under (σ^*, μ^*) . There are three cases, depending on the form of the player's most recent message from the mediator.

If the most recent message takes the form $(m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \rho, 0)$ then, since $\varepsilon^k(\mathfrak{h}^\tau) \rightarrow 0$ for all τ , player i believes with probability 1 that the mediator's state will equal ρ in all future periods. Hence, (8), combined with the fact that playing $\tilde{a}_{i,t}$ is optimal when $(\hat{\mathfrak{h}}_{-i}^t, \tilde{a}_{-i,t})$ is distributed according to $\lim_k F(\tilde{\alpha}^k) \left(\hat{\mathfrak{h}}_{-i}^t, \tilde{a}_{-i,t} | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t} \right)$, implies that playing $\tilde{a}_{i,t}$ is optimal after receiving $(m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \rho, 0)$ at faithful history \mathfrak{h}_i^t . Moreover, after observing $s_{i,t+1}$ in the following period, (7) implies that either player i 's belief is the same as in the original equilibrium or she believes with probability 1 that $(m_{i,\tau}^A, m_{i,\tau}^\Theta, m_{i,\tau}^{\leq \tau-1})$ will equal (\star, π, t) for all $\tau \geq t+1$ regardless of her own report. So truthfulness is also sequentially rational.

If the most recent message takes the form $(m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0)$ then, again

because $\varepsilon^k (\mathfrak{h}^{t+1}) \rightarrow 0$, player i believes with probability 1 that $(m_{i,\tau}^A, m_{i,\tau}^\Theta, m_{i,\tau}^{\leq \tau-1})$ will equal (\star, π, t) for all $\tau \geq t + 1$. Hence, (9) implies that playing $\psi_i[t](a_{i,t} | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t})$ is optimal. And, after observing $s_{i,t+1}$, player i again believes that the mediator's future messages will not depend on her report, so it is optimal for her to report truthfully.

Finally, if the most recent message takes the form $(m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t}) = (\star, \pi, t')$ for $t' \leq t-1$, then the mediator's state is absorbed at $\pi[\tau]$, so $(m_{i,\tau}^A, m_{i,\tau}^\Theta, m_{i,\tau}^{\leq \tau-1})$ will equal (\star, π, t') for all $\tau \geq t + 1$. By the specification of the mediator's strategy, either $(m_{i,t'}^A, m_{i,t'}^\Theta, m_{i,t'}^{\leq t'}) = (\tilde{a}_{i,t'}, \pi, 0)$ for $\tilde{a}_{i,t'} \in A_{i,t'}$ and $(m_{i,\tau}^A, m_{i,\tau}^\Theta, m_{i,\tau}^{\leq t'}) = (\star, \pi, t')$ for all $\tau \geq t'+1$, or $(m_{i,t'}^A, m_{i,t'}^\Theta, m_{i,t'}^{\leq t'}) = (\tilde{a}_{i,t'}, \rho, 0)$ for $\tilde{a}_{i,t'} \in A_{i,t'}$. In the former case, (10) implies that it is optimal for player i to play $\psi_i[t](\cdot | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t})$ and report truthfully. In the latter case, Case 4 in the specification of the player's strategy applies, and the player is assumed to play a best response.

7 Summary

In lieu of a conclusion, we provide a brief summary of our results.

- If no player can perfectly detect another's deviation, then Nash, perfect Bayesian, and sequential equilibrium are equivalent, and the revelation principle holds.
- The revelation principle holds in single-agent settings.
- In pure adverse selection games (a class that encompasses much of the literature on dynamic mechanism design), Nash, perfect Bayesian, and sequential equilibrium are essentially equivalent, and a virtual-implementation version of the revelation principle holds.
- The revelation principle holds for weak PBE.
- The revelation principle holds for CPPBE, a stronger notion of PBE introduced by Myerson (1986).
- When the mediator can tremble, sequential equilibrium is outcome-equivalent to CPPBE. This provides a version of the revelation principle for sequential equilibrium.

References

- [1] Athey, S. and I. Segal (2013), “An Efficient Dynamic Mechanism,” *Econometrica*, 81, 2463-2485.
- [2] Battaglini, M. (2005), “Long-Term Contracting with Markovian Consumers,” *American Economic Review*, 95, 637-658.
- [3] Battigalli, P. (1996), “Strategic Independence and Perfect Bayesian Equilibria,” *Journal of Economic Theory*, 70, 201-234.
- [4] Bergemann, D. and S. Morris (2013), “Robust Predictions in Games with Incomplete Information,” *Econometrica*, 81, 1251-1308.
- [5] Bergemann, D. and S. Morris (2016), “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games,” *Theoretical Economics*, 11, 487-522.
- [6] Bergemann, D. and J. Välimäki (2010), “The Dynamic Pivot Mechanism,” *Econometrica*, 78, 771-789.
- [7] Bester, H. and R. Strausz (2001), “Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case.” *Econometrica*, 69, 1077-1098.
- [8] Board, S. and A. Skrzypacz (2016), “Revenue Management with Forward-Looking Buyers,” *Journal of Political Economy*, 124, 1046-1087.
- [9] Che, Y.-K. and J. Hörner (2017), “Optimal Design for Social Learning,” *working paper*.
- [10] Conitzer, V. and T. Sandholm (2004), “Computational Criticisms of the Revelation Principle,” *Proceedings of the 5th ACM conference on Electronic Commerce*.
- [11] Courty, P. and L. Hao (2000), “Sequential Screening,” *Review of Economic Studies*, 67, 697-717.
- [12] Dhillon, A. and J.-F. Mertens (1996), “Perfect Correlated Equilibria,” *Journal of Economic Theory*, 68, 279-302.
- [13] Ely, J.C. (2017), “Beeps,” *American Economic Review*, 107, 31-53.
- [14] Ely, J., A. Frankel, and E. Kamenica (2015), “Suspense and surprise,” *Journal of Political Economy*, 123, 215-260.
- [15] Epstein, L.G. and M. Peters (1999), “A Revelation Principle for Competing Mechanisms,” *Journal of Economic Theory*, 88, 119-160.
- [16] Eső, P. and B. Szentes (2007), “Optimal Information Disclosure in Auctions and the Handicap Auction,” *Review of Economic Studies*, 74, 705-731.
- [17] Forges, F. (1986), “An Approach to Communication Equilibria,” *Econometrica*, 54, 1375-1385.

- [18] Fudenberg, D. and J. Tirole (1991), “Perfect Bayesian Equilibrium and Sequential Equilibrium,” *Journal of Economic Theory*, 53, 236-260.
- [19] Garrett, D.F. and A. Pavan (2012), “Managerial Turnover in a Changing World,” *Journal of Political Economy*, 120, 879-925.
- [20] Gerardi, D. and R.B. Myerson (2007), “Sequential Equilibria in Bayesian Games with Communication,” *Games and Economic Behavior*, 60, 104-134.
- [21] Green, J.R. and J.-J. Laffont (1986), “Partially Verifiable Information and Mechanism Design,” *Review of Economic Studies*, 53, 447-456.
- [22] Kakade, S.M., I. Lobel, and H. Nazerzadeh (2013), “Optimal Dynamic Mechanism Design and the Virtual-Pivot Mechanism,” *Operations Research*, 61, 837-854.
- [23] Kohlberg, E. and P.J. Reny, “Independence on Relative Probability Spaces and Consistent Assessments in Game Trees,” *Journal of Economic Theory*, 75, 280-313.
- [24] Kremer, I., Y. Mansour, and M. Perry (2014), “Implementing the ‘Wisdom of the Crowd’,” *Journal of Political Economy*, 122, 988-1012.
- [25] Kreps, D.M. and R. Wilson (1982), “Sequential Equilibria,” *Econometrica*, 50, 863-894.
- [26] Mailath, G.J. and L. Samuelson (2006), *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press.
- [27] Martimort, D. and L. Stole (2002), “The Revelation and Delegation Principles in Common Agency Games,” *Econometrica*, 70, 1659-1673.
- [28] Myerson, R.B. (1982), “Optimal Coordination Mechanisms in Generalized Principal-Agent Problems,” *Journal of Mathematical Economics*, 10, 67-81.
- [29] Myerson, R.B. (1986), “Multistage Games with Communication,” *Econometrica*, 54, 323-358.
- [30] Pavan, A., I. Segal, and J. Toikka (2014), “Dynamic Mechanism Design: A Myersonian Approach,” *Econometrica*, 82, 601-653.
- [31] Renault, J., E. Solan, and N. Vieille (2017), “Optimal Dynamic Information Provision,” *Games and Economic Behavior*, 104, 329-349.
- [32] Sugaya, T. and A. Wolitzky (2017), “Bounding Equilibrium Payoffs in Repeated Games with Private Monitoring,” *Theoretical Economics*, 12, 691-729.
- [33] Townsend, R.M. (1988), “Information Constrained Insurance: The Revelation Principle Extended,” *Journal of Monetary Economics*, 21, 411-450.

Appendix: Omitted Proofs

8 Details of Example 2

We consider the case where actions are observed only by the players, as the same proof works when they are also observed by the mediator.

Claim 1 *There is a non-canonical MSE in which action profile (A_1, A_2, A_3) is played with positive probability in the first period.*

Proof. Consider the following strategy profile:

The strategy of all players is to follow the mediator's recommendation (defined below) at every history, and to truthfully report the realized period 1 action profile to the mediator in period 2. (In period 1, there is nothing to report.)

The mediator's strategy is as follows:

Period 1:

- Recommend (A_1, A_2, A_3) and (A_1, B_2, A_3) with probability $1/2$ each. When recommending either action profile, inform players 1 and 2 of the full action profile, while sending only the message “ A_3 ” to player 3.

Period 2:

- If the recommended period 1 action profile was (A_1, A_2, A_3) and at least two players report that player 1 played B_1 , then recommend (A_1, C_2, A_3) . Inform each player of only her own recommendation.
- Otherwise, recommend (A_1, B_2, A_3) . Again, inform each player of only her own recommendation.

Note that this strategy profile is not canonical because in period 1 players 1 and 2 are informed of the full action profile rather than only their own actions.

Players beliefs are derived from the following sequence of completely mixed strategies σ^k : In period 1, player 1 trembles to B_1 with probability $1/k^3$ when (A_1, A_2, A_3) is recommended and trembles with probability $1/k$ when (A_1, B_2, A_3) is recommended; player 2 trembles to each of his non-recommended actions with probability $1/k^3$ when (A_1, A_2, A_3) is recommended and trembles with probability $1/k$ when (A_1, B_2, A_3) is recommended; and player 3 trembles to B_3 with probability $1/k$. Period 2 trembles are irrelevant, as the game is over at the end of period 2. All trembles at the reporting stage occur with probability $1/k$.

In the limit as $k \rightarrow \infty$, after observing any action profile other than (A_1, A_2, A_3) or (A_1, B_2, A_3) in period 1, player 3 assigns probability 1 to the event that the recommended action profile was (A_1, B_2, A_3) and players 1 and/or 2 deviated. (In particular, action profile (B_1, A_1, A_3) is infinitely more likely to result from simultaneous trembles by players 1 and 2 from recommendation (A_1, B_2, A_3) than from a unilateral deviation by player 1 from

recommendation (A_1, A_2, A_3) .) Hence, player 3 assigns probability 1 to the event that the recommended period 2 action profile is (A_1, B_2, A_3) .

With these beliefs, it is easy to see that following the mediator's recommendation is sequentially rational. Player 2's incentives are trivial. Player 1's only tempting deviation is to B_1 when (A_1, A_2, A_3) is recommended: this yields a gain of 1 in period 1 but leads to a loss of 1 in period 2, and is thus unprofitable. Player 3 never expects B_1 or C_2 to be played with positive probability, so she is willing to play A_3 . Truthfully reporting actions is also sequentially rational, as a unilateral misreport does not affect the mediator's play. ■

Claim 2 *There is no canonical MSE in which (A_1, A_2, A_3) is played with positive probability in the first period.*

Proof. First, note that with probability 1 player 3's payoff equals 1 in every period in any MSE (σ, μ, β) of the repeated game. This follows because 1 is both player 3's minmax payoff and her maximum feasible payoff. In particular, in period 1, $\Pr^{(\sigma, \mu)}(a_1 = A_1 \cap a_2 \in \{A_2, B_2\} | a_3 = A_3) = 1$.

Next, note that player 1's payoff also equals 1 in every period any MSE of the repeated game. This follows because player 1 receives payoff 1 whenever player 3 does.

Now suppose there were a MSE assessment (σ, μ, β) of the kind described in the claim. Let $\hat{\sigma} = (\hat{\sigma}_1, \sigma_2, \sigma_3)$ denote the strategy profile that results when player 1 modifies her strategy by playing B_1 in period 1 whenever she is recommended A_1 . (Thus, $\hat{\sigma}_1$ denotes the modified strategy for player 1.) Then, conditional on the event that both A_1 and A_3 are recommended in period 1, player 1's expected period 2 payoff under $\hat{\sigma}$ must be strictly less than 1, as otherwise $\hat{\sigma}_1$ would be a profitable deviation, given that (A_1, A_2, A_3) is played with positive probability in period 1. Since player 1's payoff is always at least as large as player 3's, this also implies that, conditional on the same (positive probability) event, player 3's expected period 2 payoff under $\hat{\sigma}$ must be strictly less than 1.

Now, let (h_3^2, s_3^2) denote the history for player 3 where in period 1 she was recommended and played A_3 and the realized action profile was (A_1, B_2, A_3) . Under the original assessment (σ, μ, β) , history (h_3^2, s_3^2) is reached only if player 1 was recommended A_1 but played B_1 in period 1 (because in any equilibrium $\Pr^{(\sigma, \mu)}(a_1 = A_1 | a_3 = A_3) = 1$, and the mediator does not tremble). Consistency of beliefs (and, in particular, the requirement that players tremble independently along a sequence of strategy profiles justifying equilibrium beliefs) now requires that player 3's assessment of the probability that player 2 was recommended A_2 and B_2 coincides with the ex ante distribution under σ .¹⁴ Therefore, because σ and $\hat{\sigma}$ differ only in player 1's play in period 1, player 3's expected period 2 payoff at history (h_3^2, s_3^2) under (σ, μ, β) equals her expected period 2 payoff under $\hat{\sigma}$ conditional on the event that A_1 and A_3 are recommended in period 1. Hence, player 3's expected period 2 payoff at history (h_3^2, s_3^2) under (σ, μ, β) is strictly less than 1. But player 3's minmax payoff is 1, so this contradicts the hypothesis that (σ, μ, β) is a MSE. ■

¹⁴This is the key step in the argument that uses the assumption that (σ, μ, β) is canonical. If players 1 and 2 received additional correlated signals on path (as they do in the equilibrium that yields Claim 1), then seeing player 1 deviate could affect player 3's beliefs about player 2's recommendation.

9 Proof of Lemma 1

NE: Fix a mediation range Q and a NE (σ, μ) in $G|_Q$ that induces outcome distribution $\rho \in \Delta(X)$. Let $(\tilde{\sigma}, \tilde{\mu})$ be any strategy profile in G that agrees with (σ, μ) at information sets containing a node in $G|_Q$ (that is, at histories where a player has never received a message outside Q , or where the mediator has never sent a message outside Q to any player). Then $(\tilde{\sigma}, \tilde{\mu})$ and (σ, μ) induce the same outcome distribution, and $(\tilde{\sigma}, \tilde{\mu})$ remains a NE because deviations by players other than the mediator do not lead to nodes outside $G|_Q$.

WPBE: Fix a mediation range Q and a WPBE (σ, μ, β) in $G|_Q$ that induces outcome distribution $\rho \in \Delta(X)$. For each vector $((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t})$, define an arbitrary mapping $g : M_i^t \setminus Q_i((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t}) \rightarrow Q_i((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t})$. Define an assessment $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta})$ in G by specifying that: (i) The mediator's strategy and players' strategies and beliefs at histories where they have never sent or received messages outside $Q_i((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t})$ are the same as in $G|_Q$. (In particular, players assign probability 0 to nodes outside $G|_Q$.) (ii) The mediator's strategy at any history h_0^t where he has sent messages $m_i^{t'} \in M_i^{t'} \setminus Q_i((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t'-1}, r_{i,t'})$ for some i and $t' \leq t$ is the same as his strategy in $G|_Q$ at the history \tilde{h}_0^t where every such message $m_i^{t'}$ is replaced by $g(m_i^{t'})$. (iii) A player's strategy and belief at any history h_i^t where she has received messages $m_i^{t'} \in M_i^{t'} \setminus Q_i((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t'-1}, r_{i,t'})$ for some $t' \leq t$ is the same as her strategy and belief in $G|_Q$ at the history \tilde{h}_i^t where every such message $m_i^{t'}$ is replaced by $g(m_i^{t'})$. With this construction, sequential rationality of (σ, μ, β) implies sequential rationality of $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta})$, and $\tilde{\beta}$ coincides with β on $(\tilde{\sigma}, \tilde{\mu})$ -positive probability events, so $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta})$ is a WPBE in G .

PSE: Fix a mediation range Q and a PSE (σ, μ, β) in $G|_Q$ that induces outcome distribution $\rho \in \Delta(X)$. Take a sequence of strategy profiles (σ^k, μ^k) in $G|_Q$ converging to (σ, μ) that generates beliefs β . Let $f_1(k) = \min\{\min_{h_i^t} \sigma_i^k(h_i^t), \min_{h_0^t} \mu^k(h_0^t)\}$. Fix a sequence $f_2(k)$ such that $\lim_k f_2(k) (f_1(k))^{-2(N+1)T} = 0$. Consider now the game G with the additional constraints that (i) a player must follow σ_i^k at histories where she has received only messages in Q_i ; (ii) a player must take every action with probability at least $f_1(k)$ at every history; (iii) the mediator follows μ_i^k with probability $1 - f_2(k)$ at histories where he has sent only messages in Q ; and (iv) the mediator sends messages outside of Q with probability $f_2(k)$ at every history. This constrained game has a PSE $(\tilde{\sigma}^k, \tilde{\mu}^k, \tilde{\beta}^k)$ by standard arguments (e.g., Kreps and Wilson, 1982). Let $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta}) = \lim_k (\tilde{\sigma}^k, \tilde{\mu}^k, \tilde{\beta}^k)$, taking a convergent subsequence if necessary. By standard arguments, $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta})$ is consistent and is sequentially rational for a player at any history where she has ever received a message outside Q_i . (Note that the constraint imposed at this history is that the player takes every action with probability at least $f_1(k)$, and $f_1(k) \rightarrow 0$.) In addition, at any history where a player has always received messages in Q_i , her beliefs under $\tilde{\beta}$ coincide with her beliefs under β . (In particular, she assigns probability 1 to the event that the mediator has only sent messages in Q . This follows because the number of moves in G is $2(N+1)T$, and the probability that the mediator ever sends a message outside of Q vanishes relative to the

probability of any $2(N+1)T$ trembles that do not involve such a message.) Thus, $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta})$ is also sequentially rational at these histories.

CPPBE: The fact that the lemma holds for CPPBE is not used elsewhere in the paper, so we freely use results established later on. In particular, Theorem 1 implies that any restricted CPPBE outcome distribution is a restricted PSE outcome distribution. As the lemma holds for PSE, any restricted PSE outcome distribution is a PSE outcome distribution. Finally, Theorem 1 of Myerson (1986) implies that any PSE outcome distribution is a CPPBE outcome distribution.

MSE: It follows immediately from the definition that any MSE (σ, μ, β) of $G|_Q$ is also a MSE of G .

10 Proof of Lemma 2

Fix a game G and a NE (σ, μ) . Construct a strategy profile $(\tilde{\sigma}, \tilde{\mu})$ in $C(G)$ as follows:

Players are truthful and obedient if they have been truthful in the past. Players take arbitrary best responses if they have ever lied to the mediator.¹⁵

The mediator's strategy is constructed as follows: Denote player i 's period t report by $\hat{r}_{i,t} \in A_{i,t-1} \times S_{i,t}$, where $A_{i,0} := \emptyset$. In period 1, given report $\hat{r}_{i,1} = \hat{s}_{i,1}$, the mediator draws a fictitious report $r_{i,1} \in R_{i,1}$ (the set of possible reports in G) according to $\sigma_i^R(\hat{s}_{i,1})$ (player i 's equilibrium strategy in G), independently across players. Given the resulting vector of fictitious reports $r_1 = (r_{i,1})_{i \neq 0}$, the mediator draws a vector of fictitious messages $m_1 \in M_1$ (the set of possible messages in G) according to $\mu(s_{0,1}, r_1)$. Next, given $\hat{s}_{i,1}, r_{i,1}, m_{i,1}$, the mediator draws an action recommendation $b_{i,1} \in A_{i,1}$ according to $\sigma_i^A(\hat{s}_{i,1}, r_{i,1}, m_{i,1})$. Finally, the mediator sends message $b_{i,1}$ to player i .¹⁶

Inductively, for $t = 2, \dots, T$, given $h_{i,0}^t = (\hat{a}_{i,\tau-1}, \hat{s}_{i,\tau}, r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}$ and $\hat{r}_{i,t} = (\hat{a}_{i,t-1}, \hat{s}_{i,t})$, the mediator draws $r_{i,t} \in R_{i,t}$ according to $\sigma_i^R(h_{i,0}^t, \hat{r}_{i,t})$. (Note that the history $h_{i,0}^t$ combines player i 's actual reports $\hat{h}_i^t = (\hat{a}_{i,\tau-1}, \hat{s}_{i,\tau})_{\tau=1}^{t-1}$ and the fictitious reports and messages $(r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}$.) If \hat{h}_i^t is not a possible history (that is, if there is no \hat{h}_{-i}^t with $\Pr(\hat{h}_i^t, \hat{h}_{-i}^t) > 0$), the mediator draws $r_{i,t}$ randomly. Given the resulting vector $r_t = (r_{i,t})_{i \neq 0}$, the mediator draws $m_t \in M_t$ according to $\mu((s_{0,\tau}, r_\tau, m_\tau, a_{0,\tau})_{\tau=1}^{t-1}, s_{0,t}, r_t)$. Then, given $h_{i,0}^t = (h_{i,0}^t, \hat{a}_{i,t-1}, \hat{s}_{i,t}), r_{i,t}, m_{i,t}$, the mediator draws $b_{i,t}$ according to $\sigma_i^A(h_{i,0}^t, r_{i,t}, m_{i,t})$. (Again, if \hat{h}_i^t is not a possible history, the mediator draws $b_{i,t}$ randomly.) The mediator sends message $b_{i,t}$ to player i .

We claim that the resulting canonical strategy profile is a NE. To see this, note that: (i) If player i has a profitable misreport at a history $(h_i^t, s_{i,t})$ in $C(G)$ where she has been truthful up to period t , then there is a profitable deviation in G in which she misreports whenever her payoff-relevant history equals $(\hat{h}_i^t, s_{i,t})$. (ii) If player i has a profitable deviation at a history $(h_i^t, s_{i,t}, r_{i,t}, b_{i,t})$ in $C(G)$ where she has been truthful up to and including period t ,

¹⁵As the current proof is for NE, one could alternatively specify that players are truthful and obedient at all histories. The specification in the proof is the one that generalizes to stronger solution concepts.

¹⁶The mediator's own actions can be treated in the same way, by imagining a fictitious message that the mediator sends to herself. We omit the details.

then there is a profitable deviation in G in which she deviates whenever her payoff-relevant history equals $(\hat{h}_i^t, s_{i,t})$ and her equilibrium strategy assigns positive probability to action $b_{i,t}$. Therefore, the fact that σ_i is optimal in G implies that $\tilde{\sigma}_i$ is optimal in $C(G)$.

It remains to show that $\bigcup_{\sigma'_{-i} \in \Sigma_{-i}} \text{supp } \rho_i^{(\sigma_i, \sigma'_{-i}, \mu)} \supseteq \bigcup_{\tilde{\sigma}'_{-i} \in C(\Sigma_{-i})} \text{supp } \rho_i^{(\tilde{\sigma}_i, \tilde{\sigma}'_{-i}, \tilde{\mu})}$. This follows because, for any $\tilde{\sigma}' \in C(\Sigma)$, one can construct a strategy profile $\sigma' \in \Sigma$ such that $\rho^{(\sigma', \mu)} = \rho^{(\tilde{\sigma}', \tilde{\mu})}$ as follows:

In period 1, given signal $s_{i,1}$, player i draws a fictitious ‘‘type report’’ $\hat{s}_{i,1}$ according to $\tilde{\sigma}_i^R(s_{i,1})$. Player i then sends report $r_{i,1} \in R_{i,1}$ according to $\sigma_i^R(\hat{s}_{i,1})$. Next, after receiving message $m_{i,1} \in M_{i,1}$, player i draws a fictitious action recommendation $b_{i,1} \in A_{i,1}$ according to $\sigma_i^A(\hat{s}_{i,1}, r_{i,1}, m_{i,1})$. Finally, player i takes action $a_{i,1} \in A_{i,1}$ according to $\tilde{\sigma}_i^A(s_{i,1}, \hat{s}_{i,1}, b_{i,1})$.

Inductively, given $\hat{h}_i^t = (s_{i,\tau}, \hat{a}_{i,\tau-1}, \hat{s}_{i,\tau}, b_{i,\tau}, a_{i,\tau})_{\tau=1}^{t-1}$, vector of reports and messages $(r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}$, and signal $s_{i,t}$, player i draws a fictitious type report $(\hat{a}_{i,t-1}, \hat{s}_{i,t})$ according to $\tilde{\sigma}_i^R(\hat{h}_i^t, \hat{s}_{i,t})$.

Player i then sends $r_{i,t} \in R_{i,t}$ according to $\sigma_i^R((\hat{s}_{i,\tau}, r_{i,\tau}, m_{i,\tau}, \hat{a}_{i,\tau})_{\tau=1}^{t-1}, \hat{s}_{i,t})$. (If $(\hat{a}_{i,\tau-1}, \hat{s}_{i,\tau})_{\tau=1}^t$ is not a possible history, player i draws $r_{i,t}$ randomly, and similarly for $m_{i,t}$ and $a_{i,t}$ in what follows.) Next, after receiving message $m_{i,t} \in M_{i,t}$, player i draws a fictitious action recommendation $b_{i,t} \in A_{i,t}$ according to $\sigma_i^A((\hat{s}_{i,\tau}, r_{i,\tau}, m_{i,\tau}, \hat{a}_{i,\tau})_{\tau=1}^{t-1}, \hat{s}_{i,t}, r_{i,t}, m_{i,t})$. Finally, player i takes action $a_{i,t} \in A_{i,t}$ according to $\tilde{\sigma}_i^A(\hat{h}_i^t, \hat{a}_{i,t-1}, \hat{s}_{i,t}, b_{i,t})$.

Moreover, in this construction the truthful and obedient strategy $\tilde{\sigma}_i$ is mapped to the original equilibrium strategy σ_i , so $\bigcup_{\sigma'_{-i} \in \Sigma_{-i}} \rho_i^{(\sigma_i, \sigma'_{-i}, \mu)} \supseteq \bigcup_{\tilde{\sigma}'_{-i} \in C(\Sigma_{-i})} \rho_i^{(\tilde{\sigma}_i, \tilde{\sigma}'_{-i}, \tilde{\mu})}$ and hence $\bigcup_{\sigma'_{-i} \in \Sigma_{-i}} \text{supp } \rho_i^{(\sigma_i, \sigma'_{-i}, \mu)} \supseteq \bigcup_{\tilde{\sigma}'_{-i} \in C(\Sigma_{-i})} \text{supp } \rho_i^{(\tilde{\sigma}_i, \tilde{\sigma}'_{-i}, \tilde{\mu})}$.

11 Proof of Proposition 2

Fix a game G and a WPBE (σ, μ, β) . Construct a canonical strategy profile $(\tilde{\sigma}, \tilde{\mu})$ from (σ, μ) as in the proof of Lemma 2. Let \tilde{Q} be the mediation range that restricts the mediator to sending recommendations in $\text{supp } \tilde{\mu}$: that is, for all $i \neq 0, t$, $(r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t}$,

$$\tilde{Q}_i((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t}) = \bigcup_{\substack{(\hat{h}_0^t, \tilde{s}_{0,t}, \tilde{r}_t) \\ ((\tilde{r}_{i,\tau}, \tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \tilde{r}_{i,t}) = ((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t})}} \text{supp } \tilde{\mu}_i(\hat{h}_0^t, \tilde{s}_{0,t}, \tilde{r}_t).$$

We thus view $(\tilde{\sigma}, \tilde{\mu})$ as a strategy profile in game $C(G)|_{\tilde{Q}}$. To construct a belief system $\tilde{\beta}$ in $C(G)|_{\tilde{Q}}$, for each history $(h_i^t, s_{i,t})$ in $C(G)$ with $\hat{r}_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$ for all $\tau < t$, if $\text{Pr}^{(\sigma, \mu)}(\hat{h}_i^t, s_{i,t}) > 0$ then (i) let

$$\tilde{\beta}_i(h_i^t, s_{i,t})|_{\hat{H}^t \times S_t} = \sum_{\hat{h}_i^t \in H_i^t: \hat{h}_i^t = \hat{h}_i^t} \frac{\text{Pr}^{(\sigma, \mu)}(\bar{h}_i^t, s_{i,t})}{\text{Pr}^{(\sigma, \mu)}(\hat{h}_i^t, s_{i,t})} \beta_i(\bar{h}_i^t, s_{i,t})|_{\hat{H}^t \times S_t}$$

(thus, $\tilde{\beta}_i(h_i^t, s_{i,t})|_{\hat{H}^t \times S_t}$ corresponds to the conditional belief about (\hat{h}_i^t, s_t) given $(\hat{h}_i^t, s_{i,t})$ in the original equilibrium), and (ii) let $\tilde{\beta}_i(h_i^t, s_{i,t})$ assign probability 1 to the event that players $-i$ have always been truthful and obedient. If instead $\Pr^{(\sigma, \mu)}(\hat{h}_i^t, s_{i,t}) = 0$, then (i) let $\tilde{\beta}_i(h_i^t, s_{i,t})|_{\hat{H}^t \times S_t} = \beta_i(\bar{h}_i^t, s_{i,t})|_{\hat{H}^t \times S_t}$ for an arbitrary history \bar{h}_i^t in G with $\hat{h}_i^t = \bar{h}_i^t$ and $r_{i,\tau} \in \text{supp } \sigma_i^R(\bar{h}_i^\tau, s_{i,\tau})$ for all $\tau < t$, and (ii) let $\tilde{\beta}_i(h_i^t, s_{i,t})$ assign probability 1 to the event that players $-i$ have always been truthful (but not necessarily obedient). Construct beliefs for histories of the form $(h_i^t, s_{i,t}, \hat{r}_{i,t}, b_{i,t})$ analogously (with $\hat{r}_{i,t} = (\hat{a}_{i,t-1}, \hat{s}_{i,t})$ as in the proof of Lemma 2), where if $\Pr^{(\sigma, \mu)}(\hat{h}_i^t, s_{i,t}, b_{i,t}) = 0$ then $\tilde{\beta}_i(h_i^t, s_{i,t}, \hat{r}_{i,t}, b_{i,t})|_{\hat{H}^t \times S_t} = \beta_i(\bar{h}_i^t, s_{i,t}, r_{i,t}, m_{i,t})|_{\hat{H}^t \times S_t}$ for an arbitrary history $(\bar{h}_i^t, s_{i,t}, r_{i,t}, m_{i,t})$ with $\hat{h}_i^t = \bar{h}_i^t$ and $r_{i,\tau} \in \text{supp } \sigma_i^R(\bar{h}_i^\tau, s_{i,\tau})$ for all $\tau \leq t$ and $b_{i,t} \in \text{supp } \sigma_i^A(\bar{h}_i^t, s_{i,t}, r_{i,t}, m_{i,t})$. Beliefs at histories where a player has lied to the mediator may be specified arbitrarily.

We claim that any assessment $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta})$ so constructed is a WPBE. To see this, it is straightforward to check the following claims: (i) If a previously truthful player i has a profitable misreport in $C(G)$ at history $(h_i^t, s_{i,t})$ with belief $\tilde{\beta}_i(h_i^t, s_{i,t})$, then she has a profitable deviation in G at some history $(\bar{h}_i^t, s_{i,t})$ with belief $\beta_i(\bar{h}_i^t, s_{i,t})$. (ii) If a previously truthful player i has a profitable deviation in $C(G)$ at history $(h_i^t, s_{i,t}, \hat{r}_{i,t}, b_{i,t})$ with belief $\tilde{\beta}_i(h_i^t, s_{i,t}, \hat{r}_{i,t}, b_{i,t})$, then she has a profitable deviation in G at some history $(\bar{h}_i^t, s_{i,t}, r_{i,t}, m_{i,t})$ with belief $\beta_i(\bar{h}_i^t, s_{i,t}, r_{i,t}, m_{i,t})$. Given this, the fact that (σ, μ, β) is sequentially rational in G implies that $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta})$ is sequentially rational in $C(G)$. Finally, $\tilde{\beta}|_{\hat{H}^{T+1}}$ coincides with $\beta|_{\hat{H}^{T+1}}$ on positive-probability histories, so $\tilde{\beta}$ is consistent with Bayes rule.

12 Proofs of Lemmas 4 and 5

In this appendix, we reduce notation by writing \Pr^k for \Pr^{σ^k, μ^k} and F^k for $F(\tilde{\alpha}^k)$.

12.1 A Preliminary Lemma

We first show that, conditional on the event $\hat{\theta}_{t+1} = \rho$, a faithful player's beliefs about $((a_{\tau-1}, s_\tau, \tilde{a}_\tau)_{\tau=1}^t, a_t, s_{t+1})$ are the same as in the original equilibrium.

Lemma 6 *For every faithful history $\mathfrak{h}_i^{t+1} = ((a_{i,\tau-1}, s_{i,\tau}, \tilde{a}_{i,\tau}, m_{i,\tau})_{\tau=1}^t, a_{i,t}, s_{i,t+1})$ consistent with $\hat{\theta}_{t+1} = \rho$,*

$$\lim_k \frac{\Pr^k \left((a_{\tau-1}, s_\tau, \tilde{a}_\tau, m_\tau)_{\tau=1}^t, a_t, s_{t+1} | \mathfrak{h}_i^{t+1}, \hat{\theta}_{t+1} = \rho \right)}{F^k \left((a_{\tau-1}, s_\tau, \tilde{a}_\tau)_{\tau=1}^t, a_t, s_{t+1} | (a_{i,\tau-1}, s_{i,\tau}, \tilde{a}_{i,\tau})_{\tau=1}^t, a_{i,t}, s_{i,t+1} \right)} = 1 \quad (12)$$

for all $(a_{\tau-1}, s_\tau, \tilde{a}_\tau, m_\tau)_{\tau=1}^t, a_t, s_{t+1}$ such that \mathfrak{h}_j^{t+1} is faithful for all $j \neq i$.

Proof. We claim that it suffices to show that

$$\lim_k \frac{\Pr^k \left(\tilde{a}_{-i,t}, m_{-i,t}^\ominus, a_{-i,t} | \mathfrak{h}_0^t, \tilde{a}_{i,t}, m_{i,t}^\ominus, a_{i,t}, \hat{\theta}_{t+1} = \rho \right)}{F^k \left(\tilde{a}_{-i,t}, a_{-i,t} | \hat{\mathfrak{h}}_0^t, \tilde{a}_{i,t}, a_{i,t} \right)} = 1 \quad (13)$$

for all $\mathfrak{h}_0^t, \tilde{a}_t, m_t^\ominus, a_t$ such that $(\mathfrak{h}_j^t, (\tilde{a}_{j,t} = m_{j,t}^A), m_{j,t}^\ominus, a_{j,t})$ is faithful for all j . Note that we have omitted the term $m_t^{\leq t-1}$, since $\hat{\theta}_{t+1} = \rho$ implies $m_t^{\leq t-1} = 0$.

To see why (13) implies (12), note the following: First, summing (13) over $\tilde{a}_{-i,t}, m_{-i,t}^\ominus, a_{-i,t}$ such that $(\mathfrak{h}_j^t, \tilde{a}_{j,t}, m_{j,t}^\ominus, a_{j,t})$ is faithful for all $j \neq i$, and noting that this sum includes all pairs $(\tilde{a}_{-i,t}, a_{-i,t})$, we have

$$\sum_{\substack{\tilde{a}_{-i,t}, m_{-i,t}^\ominus, a_{-i,t}: \\ (\mathfrak{h}_j^t, \tilde{a}_{j,t}, m_{j,t}^\ominus, a_{j,t}) \text{ is faithful } \forall j \neq i}} \Pr^k \left(\begin{array}{c} \tilde{a}_{-i,t}, m_{-i,t}^\ominus, a_{-i,t} \\ | \mathfrak{h}_0^t, \tilde{a}_{i,t}, m_{i,t}^\ominus, a_{i,t}, \hat{\theta}_{t+1} = \rho \end{array} \right) = 1. \quad (14)$$

This implies that a faithful player believes that all other players have been faithful. Second, the distribution of s_{t+1} is determined by $\hat{\mathfrak{h}}_0^t$ and a_t , and $\varepsilon^k(\mathfrak{h}_0^{t+1})$ does not depend on s_{t+1} .

We now prove (13). Note that if $\hat{\theta}_{t+1} = \rho$ then, for each i and $\tau \leq t$, the event that either $[m_{i,\tau}^\ominus = \rho \text{ and } \tilde{a}_{i,\tau} \neq a_{i,\tau}]$ or $[m_{i,\tau}^\ominus = \pi \text{ and } \tilde{a}_{i,\tau} = a_{i,\tau}]$ occurs only if some player is unfaithful, and hence occurs with probability at most $f_4(k)$. This follows because (i) if $\theta_\tau = \pi[\tau']$ for $\tau' \leq \tau - 1$ then $\hat{\theta}_{t+1}$ would equal $\pi[\tau']$ rather than ρ , (ii) if $\theta_\tau = \pi[\tau]$ and either $[m_{i,\tau}^\ominus = \rho \text{ and } \tilde{a}_{i,\tau} \neq a_{i,\tau}]$ or $[m_{i,\tau}^\ominus = \pi \text{ and } \tilde{a}_{i,\tau} = a_{i,\tau}]$ then $\varepsilon^k(\mathfrak{h}_0^{\tau+1}) = 0$ when players are truthful, and hence $\hat{\theta}_{t+1}$ would equal $\pi[\tau]$ rather than ρ , and (iii) if $\theta_\tau = \rho$ then the event $[m_{i,\tau}^\ominus = \rho \text{ and } \tilde{a}_{i,\tau} \neq a_{i,\tau}]$ occurs only if player i takes an $f_4(k)$ action. So, with probability at least $1 - f_4(k)$, if $\tilde{a}_{i,\tau} = a_{i,\tau}$ then $m_{i,\tau}^\ominus = \rho$, and if $\tilde{a}_{i,\tau} \neq a_{i,\tau}$ then $m_{i,\tau}^\ominus = \pi$. Hence,

$$\left| \begin{array}{c} \Pr^k \left(\tilde{a}_{-i,t}, m_{-i,t}^\ominus, a_{-i,t} | \mathfrak{h}_0^t, \tilde{a}_{i,t}, m_{i,t}^\ominus, a_{i,t}, \hat{\theta}_{t+1} = \rho \right) \\ - \Pr^k \left(\tilde{a}_{-i,t}, a_{-i,t} | \mathfrak{h}_0^t, \tilde{a}_{i,t}, a_{i,t}, \hat{\theta}_{t+1} = \rho \right) \end{array} \right| \leq f_4(k)$$

for $\tilde{a}_t, m_t^\ominus, a_t$ such that $(\mathfrak{h}_j^t, (\tilde{a}_{j,t} = m_{j,t}^A), m_{j,t}^\ominus)$ is faithful for all j .

As $F^k \left(\tilde{a}_{-i,t}, a_{-i,t} | \hat{\mathfrak{h}}_0^t, \tilde{a}_{i,t}, a_{i,t} \right) \geq \underline{f}_3(k)$, to establish (13) it suffices to show

$$1 - f_2(k) + O\left(\underline{f}_3(k)\right) \leq \frac{\Pr^k \left(\tilde{a}_{-i,t}, a_{-i,t} | \mathfrak{h}_0^t, \tilde{a}_{i,t}, a_{i,t}, \hat{\theta}_{t+1} = \rho \right)}{F^k \left(\tilde{a}_{-i,t}, a_{-i,t} | \hat{\mathfrak{h}}_0^t, \tilde{a}_{i,t}, a_{i,t} \right)} \leq \frac{1}{1 - f_2(k)} + O\left(\underline{f}_3(k)\right), \quad (15)$$

where $O(f_n(k))$ denotes a function $g(k)$ with $\limsup |g(k)|/f_n(k) < \infty$.

For any $\tilde{a}_t, a_t \in A_t$, let $I^{**}(\tilde{a}_t, a_t) = \{i \in I : \tilde{a}_{i,t} = a_{i,t}\}$. We claim that

$$\begin{aligned} & \Pr^k \left(\tilde{a}_{-i,t}, a_{-i,t} | \mathfrak{h}_0^t, \tilde{a}_{i,t}, a_{i,t}, \hat{\theta}_{t+1} = \rho \right) \\ & \Pr^k \left(\tilde{a}_t | \mathfrak{h}_0^t, \hat{\theta}_t = \rho \right) \left[\frac{\mathbf{1}_{\{I^{**}(\tilde{a}_t, a_t) = I\}} (1 - f_2(k))}{+\mathbf{1}_{\{I^{**}(\tilde{a}_t, a_t) \neq I\}} F^k \left(a_t | \mathfrak{h}_0^t, \tilde{a}_t \right)} \right] + O(f_4(k)) \\ = & \frac{\Pr^k \left(\tilde{a}_{-i,t}, a_{-i,t} | \mathfrak{h}_0^t, \hat{\theta}_{t+1} = \rho \right)}{\sum_{\tilde{a}'_{-i,t}, a'_{-i,t}} \Pr^k \left((\tilde{a}_{i,t}, \tilde{a}'_{-i,t}), (a_{i,t}, a'_{-i,t}) | \mathfrak{h}_0^t, \hat{\theta}_{t+1} = \rho \right)}. \end{aligned} \quad (16)$$

In particular, this equation asserts (i) if all players take actions equal to their provisional recommendations, then $\Pr^k \left(\tilde{a}_t, a_t | \mathfrak{h}_0^t, \hat{\theta}_{t+1} = \rho \right)$ equals $\Pr^k \left(\tilde{a}_t | \mathfrak{h}_0^t, \hat{\theta}_t = \rho \right) (1 - f_2(k)) + O(f_4(k))$, and (ii) if some players' actions differ from their provisional recommendations, then $\Pr^k \left(\tilde{a}_t, a_t | \mathfrak{h}_0^t, \hat{\theta}_{t+1} = \rho \right)$ equals $\Pr^k \left(\tilde{a}_t | \mathfrak{h}_0^t, \hat{\theta}_t = \rho \right) \left(F^k \left(a_t | \mathfrak{h}_0^t, \tilde{a}_t \right) \right) + O(f_4(k))$. To see this, consider the $I^{**}(\tilde{a}_t, a_t) = I$ and $I^{**}(\tilde{a}_t, a_t) \neq I$ cases in turn:

1. $I^{**}(\tilde{a}_t, a_t) = I$. In this case, if $\theta_t = \pi[t]$ then $\varepsilon^k(\mathfrak{h}_0^{\tau+1}) = 0$, and therefore $\hat{\theta}_{t+1}$ would equal $\pi[t]$ rather than ρ . Hence, $\theta_t = \rho$.¹⁷ As $\Pr^k(a_t | \tilde{a}_t, \theta_t = \rho) = 1 - O(f_4(k))$,

$$\begin{aligned} & \Pr^k \left(\tilde{a}_t, a_t | \mathfrak{h}_0^t, \hat{\theta}_{t+1} = \rho \right) \\ = & \Pr^k \left(\tilde{a}_t | \mathfrak{h}_0^t, \hat{\theta}_t = \rho \right) \Pr^k \left(\theta_t = \rho | \mathfrak{h}_0^t, \tilde{a}_t, \hat{\theta}_t = \rho \right) + O(f_4(k)) \\ = & \Pr^k \left(\tilde{a}_t | \mathfrak{h}_0^t, \hat{\theta}_t = \rho \right) (1 - f_2(k)) + O(f_4(k)). \end{aligned}$$

2. $I^{**}(\tilde{a}_t, a_t) \neq I$. In this case, the possibility that $\theta_t = \rho$ contributes a term of order $f_4(k)$, as $\Pr^k(a_t | \tilde{a}_t, \theta_t = \rho) = O(f_4(k))$. If instead $\theta_t = \pi[t]$ then, for each player j , $m_{j,t}^\ominus = \rho$ (or equivalently $j \in I_t^*$) if and only if $a_{j,t} = \tilde{a}_{j,t}$, as otherwise $\varepsilon^k(\mathfrak{h}_0^{t+1}) = 0$ and $\hat{\theta}_{t+1}$ would equal $\pi[t]$. Hence,

$$\begin{aligned} & \Pr^k \left(\tilde{a}_t, a_t | \mathfrak{h}_0^t, \hat{\theta}_{t+1} = \rho \right) \\ = & \Pr^k \left(\tilde{a}_t | \mathfrak{h}_0^t, \hat{\theta}_t = \rho \right) \Pr^k \left(I^* = I^{**}(\tilde{a}_t, a_t), a_t, \hat{\theta}_{t+1} = \rho | \mathfrak{h}_0^t, \tilde{a}_t, \hat{\theta}_t = \rho \right) + O(f_4(k)) \\ = & \Pr^k \left(\tilde{a}_t | \mathfrak{h}_0^t, \hat{\theta}_t = \rho \right) F^k \left(a_t | \mathfrak{h}_0^t, \tilde{a}_t \right) + O(f_4(k)), \end{aligned}$$

where the last line follows by (4).

Combining the cases yields (16).

Next, recall that

$$\Pr^k \left(\tilde{a}_t | \mathfrak{h}_0^t, \hat{\theta}_t = \rho \right) = \tilde{\mu}^k \left(\tilde{a}_t | \mathfrak{h}_0^t \right).$$

¹⁷Again, if $\theta_t = \pi[\tau]$ for $\tau \leq t-1$, then $\hat{\theta}_{t+1}$ would equal $\pi[\tau]$.

Hence, (16) implies

$$\begin{aligned}
& \Pr^k \left(\tilde{a}_{-i,t}, a_{-i,t} | \mathfrak{h}_0^t, \tilde{a}_{i,t}, a_{i,t}, \hat{\theta}_{t+1} = \rho \right) \\
&= \frac{\tilde{\mu}^k \left(\tilde{a}_t | \widehat{\mathfrak{h}}_0^t \right) \left[\begin{array}{l} \mathbf{1}_{\{I^{**}(\tilde{a}_t, a_t) = I\}} (1 - f_2(k)) \\ + \mathbf{1}_{\{I^{**}(\tilde{a}_t, a_t) \neq I\}} F^k \left(a_t | \widehat{\mathfrak{h}}_0^t, \tilde{a}_t \right) \end{array} \right] + O(f_4(k))}{\sum_{\tilde{a}'_{-i,t}, a'_{-i,t}} \tilde{\mu}^k \left(\tilde{a}_{i,t}, \tilde{a}'_{-i,t} | \widehat{\mathfrak{h}}_0^t \right) \left[\begin{array}{l} \mathbf{1}_{\{I^{**}((\tilde{a}_{i,t}, \tilde{a}'_{-i,t}), (a_{i,t}, a'_{-i,t})) = I\}} (1 - f_2(k)) \\ + \mathbf{1}_{\{I^{**}((\tilde{a}_{i,t}, \tilde{a}'_{-i,t}), (a_{i,t}, a'_{-i,t})) \neq I\}} F^k \left(a_{i,t}, a'_{-i,t} | \widehat{\mathfrak{h}}_0^t, \tilde{a}_{i,t}, \tilde{a}'_{-i,t} \right) \\ + O(f_4(k)) \end{array} \right]}.
\end{aligned}$$

Since

$$\tilde{\mu}^k \left(\tilde{a}_t | \widehat{\mathfrak{h}}_0^t \right) \mathbf{1}_{\{I^{**}(\tilde{a}_t, a_t) = I\}} (1 - f_2(k)) = \tilde{\mu}^k \left(\tilde{a}_t | \widehat{\mathfrak{h}}_0^t \right) \mathbf{1}_{\{I^{**}(\tilde{a}_t, a_t) = I\}} F^k \left(a_t | \widehat{\mathfrak{h}}_0^t, \tilde{a}_t \right) \frac{1 - f_2(k)}{F^k \left(a_t | \widehat{\mathfrak{h}}_0^t, \tilde{a}_t \right)}$$

and

$$\begin{aligned}
\frac{1 - f_2(k)}{F^k \left(a_t | \widehat{\mathfrak{h}}_0^t, \tilde{a}_t \right)} &\in \left[1 - f_2(k), \frac{1 - f_2(k)}{1 - \underline{f}_3(k)} \right] \leq 1 \text{ if } I^{**}(\tilde{a}_t, a_t) = I \\
&\text{(since } I^{**}(\tilde{a}_t, a_t) = I \text{ implies } a_t = \tilde{a}_t),
\end{aligned}$$

we have

$$\begin{aligned}
& \Pr^{\sigma^k, \mu^k} \left(\tilde{a}_{-i,t}, a_{-i,t} | \mathfrak{h}_0^t, \tilde{a}_{i,t}, \hat{\theta}_{t+1} = \rho \right) \\
&\leq \frac{1}{1 - f_2(k)} \frac{\tilde{\mu}^k \left(\tilde{a}_t | \widehat{\mathfrak{h}}_0^t \right) F^k \left(a_t | \widehat{\mathfrak{h}}_0^t, \tilde{a}_t \right) + O(f_4(k))}{\sum_{\tilde{a}'_{-i,t}, a'_{-i,t}} \tilde{\mu}^k \left(\tilde{a}_{i,t}, \tilde{a}'_{-i,t} | \widehat{\mathfrak{h}}_0^t \right) F^k \left(a_{i,t}, a'_{-i,t} | \widehat{\mathfrak{h}}_0^t, \tilde{a}_{i,t}, \tilde{a}'_{-i,t} \right) + O(f_4(k))} \\
&\leq \frac{1}{1 - f_2(k)} F^k \left(\tilde{a}_{-i,t}, a_{-i,t} | \widehat{\mathfrak{h}}_0^t, \tilde{a}_{i,t}, a_{i,t} \right) + O \left(\frac{f_4(k)}{\underline{f}_3(k)} \right)
\end{aligned}$$

and

$$\begin{aligned}
& \Pr^k \left(\tilde{a}_{-i,t}, a_{-i,t} | \mathfrak{h}_0^t, \tilde{a}_{i,t}, \hat{\theta}_{t+1} = \rho \right) \\
&\geq (1 - f_2(k)) \frac{\tilde{\mu}^k \left(\tilde{a}_t | \widehat{\mathfrak{h}}_0^t \right) F^k \left(a_t | \widehat{\mathfrak{h}}_0^t, \tilde{a}_t \right) + O(f_4(k))}{\sum_{\tilde{a}'_{-i,t}, a'_{-i,t}} \tilde{\mu}^k \left(\tilde{a}_{i,t}, \tilde{a}'_{-i,t} | \widehat{\mathfrak{h}}_0^t \right) F^k \left(a_{i,t}, a'_{-i,t} | \widehat{\mathfrak{h}}_0^t, \tilde{a}_{i,t}, \tilde{a}'_{-i,t} \right) + O(f_4(k))} \\
&= (1 - f_2(k)) F^k \left(\tilde{a}_{-i,t}, a_{-i,t} | \widehat{\mathfrak{h}}_0^t, \tilde{a}_{i,t}, a_{i,t} \right) + O \left(\frac{f_4(k)}{\underline{f}_3(k)} \right).
\end{aligned}$$

As $f_4(k) / \underline{f}_3(k) < \underline{f}_3(k)$, this yields (15). ■

12.2 Inductive Hypothesis

We prove Lemmas 4 and 5 by induction. If the conclusions of Lemmas 4 and 5 hold for period $t - 1$ then, for every faithful \mathfrak{h}_i^t and $\hat{\mathfrak{h}}_i^t$, we have

$$\lim_k \Pr^k (\mathfrak{h}_{-i}^t | \mathfrak{h}_i^t, (m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\rho, 0)) = \lim_k F^k (\hat{\mathfrak{h}}_{-i}^t | \hat{\mathfrak{h}}_i^t). \quad (17)$$

(This follows because part 3 of Lemmas 4 and 5 cover all cases with $(m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\rho, 0)$, as $m_{i,t+1}^{\leq t} = 0$ implies $m_{i,t}^{\leq t-1} = 0$). Moreover, (17) holds for $t = 1$. Hence, by induction, it suffices to prove the claims in Lemmas 4 and 5 for period t , given (17).

Note also that (17) implies that, for non-faithful \mathfrak{h}_{-i}^t ,

$$\lim_k \Pr^k (\mathfrak{h}_{-i}^t | \mathfrak{h}_i^t, (m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\rho, 0)) = 0. \quad (18)$$

This follows by summing (17) over faithful \mathfrak{h}_{-i}^t and recalling that, under F^k , truthful players believe their opponents have been truthful.

12.3 Proof of Lemma 4

Since $m_{i,t}^{\leq t-1} = 0$ implies $\hat{\theta}_t = \rho$ and $m_{-i,t}^{\leq t-1} = 0$, we have

$$\begin{aligned} & \Pr^k (\mathfrak{h}_{-i}^t, \theta_t = \rho, (m_{-i,t}^A, m_{-i,t}^\ominus, m_{-i,t}^{\leq t-1}) = (\tilde{a}_{-i,t}, \rho, 0) | \mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \rho, 0)) \\ & \Pr^k (\mathfrak{h}_{-i}^t | \mathfrak{h}_i^t, \hat{\theta}_t = \rho) \Pr^k (\theta_t = \rho | \hat{\theta}_t = \rho) \mu^k (m_t^A = (\tilde{a}_{i,t}, \tilde{a}_{-i,t}) | \mathfrak{h}^t) \\ = & \frac{\Pr^k (\mathfrak{h}_{-i}^t | \mathfrak{h}_i^t, \hat{\theta}_t = \rho) \Pr^k (\theta_t = \rho | \hat{\theta}_t = \rho) \mu^k (m_t^A = (\tilde{a}_{i,t}, \tilde{a}_{-i,t}) | \mathfrak{h}^t)}{\left(\begin{aligned} & \sum_{\tilde{\mathfrak{h}}_{-i}^t, \hat{a}_{-i,t}} \Pr^k (\tilde{\mathfrak{h}}_{-i}^t | \mathfrak{h}_i^t, \hat{\theta}_t = \rho) \Pr^k (\theta_t = \rho | \hat{\theta}_t = \rho) \mu^k (m_t^A = (\tilde{a}_{i,t}, \hat{a}_{-i,t}) | \mathfrak{h}_i^t, \tilde{\mathfrak{h}}_{-i}^t) \\ & + \sum_{\tilde{\mathfrak{h}}_{-i}^t, \hat{a}_{-i,t}} \Pr^k (\tilde{\mathfrak{h}}_{-i}^t | \mathfrak{h}_i^t, \hat{\theta}_t = \rho) \Pr^k (\theta_t \neq \rho | \hat{\theta}_t = \rho) \mu^k (m_{i,t}^A = (\tilde{a}_{i,t}, \hat{a}_{-i,t}) | \mathfrak{h}_i^t, \tilde{\mathfrak{h}}_{-i}^t) \end{aligned} \right)}. \end{aligned}$$

For each $\tilde{\mathfrak{h}}_{-i}^t, \hat{a}_{-i,t}$, we have

$$\frac{\Pr^k (\tilde{\mathfrak{h}}_{-i}^t | \mathfrak{h}_i^t, \hat{\theta}_t = \rho) \Pr^k (\theta_t = \rho | \hat{\theta}_t = \rho) \mu^k (m_t^A = (\tilde{a}_{i,t}, \hat{a}_{-i,t}) | \mathfrak{h}_i^t, \tilde{\mathfrak{h}}_{-i}^t)}{\Pr^k (\tilde{\mathfrak{h}}_{-i}^t | \mathfrak{h}_i^t, \hat{\theta}_t = \rho) \Pr^k (\theta_t \neq \rho | \hat{\theta}_t = \rho) \mu^k (m_t^A = (\tilde{a}_{i,t}, \hat{a}_{-i,t}) | \mathfrak{h}_i^t, \tilde{\mathfrak{h}}_{-i}^t)} = \frac{1 - f_2(k)}{f_2(k)} \rightarrow \infty.$$

Hence,

$$\begin{aligned} & \lim_k \Pr^k (\mathfrak{h}_{-i}^t, \theta_t = \rho, (m_{-i,t}^A, m_{-i,t}^\ominus, m_{-i,t}^{\leq t-1}) = (\tilde{a}_{-i,t}, \rho, 0) | \mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \rho, 0)) \\ = & \lim_k \frac{\Pr^k (\mathfrak{h}_{-i}^t | \mathfrak{h}_i^t, \hat{\theta}_t = \rho) \mu^k (m_t^A = (\tilde{a}_{i,t}, \tilde{a}_{-i,t}) | \mathfrak{h}^t)}{\sum_{\tilde{\mathfrak{h}}_{-i}^t, \hat{a}_{-i,t}} \Pr^k (\tilde{\mathfrak{h}}_{-i}^t | \mathfrak{h}_i^t, \hat{\theta}_t = \rho) \mu^k (m_t^A = (\tilde{a}_{i,t}, \hat{a}_{-i,t}) | \mathfrak{h}_i^t, \tilde{\mathfrak{h}}_{-i}^t)} \\ = & \lim_k F^k (\hat{\mathfrak{h}}_{-i}^t, \tilde{a}_{-i,t} | \hat{\mathfrak{h}}_i^t, a_{i,t}), \end{aligned}$$

by (12). This gives (6).

We now prove (7). Since $m_{i,t}^{\leq t-1} = 0$ implies $\hat{\theta}_t = \rho$ and $m_{-i,t}^{\leq t-1} = 0$, and $\theta_t = \rho$ implies $m_{i,t}^\Theta = \rho$, we have

$$\begin{aligned}
& \lim_k \Pr^k (\theta_t = \rho | \mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\Theta, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \rho, 0), a_{i,t}, s_{i,t+1}) \\
&= \lim_k \frac{\sum_{\hat{\mathfrak{h}}_0^t \left(\Pr^k \left(\hat{\mathfrak{h}}_0^t | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t} \right) \right.}{\sum_{\hat{\mathfrak{h}}_0^t \left(\Pr^k \left(\hat{\mathfrak{h}}_0^t | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t} \right) \right.} \left(\tilde{\mu}^k \left(\tilde{a}_{-i,t} | \tilde{a}_{i,t}, \hat{\mathfrak{h}}_0^t \right) \underbrace{(1 - f_2(k))}_{\Pr^k(\theta_t = \rho | \hat{\theta}_t = \rho)} p \left(s_{i,t+1} | \hat{\mathfrak{h}}_0^t, a_{i,t}, \tilde{a}_{-i,t} \right) \right)}{\sum_{\hat{\mathfrak{h}}_0^t \left(\Pr^k \left(\hat{\mathfrak{h}}_0^t | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t} \right) \right.} \left(\tilde{\mu}^k \left(\tilde{a}_{-i,t} | \tilde{a}_{i,t}, \hat{\mathfrak{h}}_0^t \right) (1 - f_2(k)) p \left(s_{i,t+1} | \hat{\mathfrak{h}}_0^t, a_{i,t}, \tilde{a}_{-i,t} \right) \right)} \left(\Pr^k \left(\hat{\mathfrak{h}}_0^t | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t} \right) \right. \\
& \quad \left. + \sum_{\hat{\mathfrak{h}}_0^t \left(\Pr^k \left(\hat{\mathfrak{h}}_0^t | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t} \right) \right.} \left(\tilde{\mu}^k \left(\tilde{a}_{-i,t} | \tilde{a}_{i,t}, \hat{\mathfrak{h}}_0^t \right) \underbrace{(f_2(k))^2}_{\Pr^k(\theta_t \neq \rho | \hat{\theta}_t = \rho)} \right) \left(\sum_{a_{-i,t}} \Pr^k \left(a_{-i,t} | \mathfrak{h}_i^t, \hat{\mathfrak{h}}_0^t, \tilde{a}_{i,t}, \theta_t \neq \rho \right) p \left(s_{i,t+1} | \hat{\mathfrak{h}}_0^t, a_t \right) \right) \right) \quad (19)
\end{aligned}$$

The numerator represents the probability that $\theta_t = \rho$, where we have neglected the $O(f_4(k))$ event that players disobey their recommended actions \tilde{a}_t . (One could instead keep track of this term as we did in the proof of Lemma 6; it makes no difference.) The denominator adds the probability that $\theta_t \neq \rho$, where the $O(f_1(k))$ event that players disobey their recommendations must be taken into account.

Given $\mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t}$, let us partition the set of canonical histories $\mathfrak{H}^t \ni \hat{\mathfrak{h}}_0^t$ into sets $\mathfrak{H}^t[1], \dots, \mathfrak{H}^t[L] \subseteq \mathfrak{H}^t$ according to the following criteria:

$$\begin{aligned}
\lim_k \frac{\Pr^k \left(\mathfrak{h} | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t} \right)}{\Pr^k \left(\mathfrak{h}' | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t} \right)} &\in (0, \infty) \text{ for each } \mathfrak{h}, \mathfrak{h}' \in \mathfrak{H}^t[l] \text{ with } l = 1, \dots, L; \\
\lim_k \frac{\Pr^k \left(\mathfrak{h} | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t} \right)}{\Pr^k \left(\mathfrak{h}' | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t} \right)} &= 0 \text{ for each } \mathfrak{h} \in \mathfrak{H}^t[l], \mathfrak{h}' \in \mathfrak{H}^t[l'] \text{ with } l' < l.
\end{aligned}$$

That is, each $\mathfrak{H}^t[l]$ is a set of canonical histories with mutually bounded likelihood ratios, and each $\mathfrak{H}^t[l']$ with $l' < l$ is a set of qualitatively less likely histories. Let l^* be the smallest number l such that there exists $\hat{\mathfrak{h}}_0^t \in \mathfrak{H}^t[l]$ with $p \left(s_{i,t+1} | \hat{\mathfrak{h}}_0^t, a_{i,t}, \tilde{a}_{-i,t} \right) > 0$ for some $a_{-i,t}$.

Consider the following two cases:

Case 1: There exists $\hat{\mathfrak{h}}_0^t \in \mathfrak{H}^t[l^*]$ and $\tilde{a}_{-i,t}$ such that

$$\lim_k \tilde{\mu}^k \left(\tilde{a}_{-i,t} | \tilde{a}_{i,t}, \hat{\mathfrak{h}}_0^t \right) > 0 \text{ and } p \left(s_{i,t+1} | \hat{\mathfrak{h}}_0^t, a_{i,t}, \tilde{a}_{-i,t} \right) > 0. \quad (20)$$

In this case, we claim that player i believes $\theta_t = \rho$:

$$\lim_k \Pr^k (\theta_t = \rho | \mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \rho, 0), a_{i,t}, s_{i,t+1}) = 1.$$

To see why, first note that if $l < l^*$ and $\mathfrak{h} \in \mathfrak{H}^t[l]$ then

$$\lim_k \Pr^k \left(\mathfrak{h} | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t} \right) \sum_{\tilde{a}_{-i,t}} \left(\rho^k (\tilde{a}_{-i,t} | \tilde{a}_{i,t}, \mathfrak{h}) (1 - f_2(k)) p \left(s_{i,t+1} | \hat{\mathfrak{h}}, a_{i,t}, \tilde{a}_{-i,t} \right) \right) = 0.$$

Next, for each $\mathfrak{h} \in \mathfrak{H}^t[l^*]$ and $\mathfrak{h}' \in \mathfrak{H}^t[l]$ with $l > l^*$, we have

$$\begin{aligned} \frac{\Pr^k \left(\mathfrak{h}' | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t} \right)}{\Pr^k \left(\mathfrak{h} | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t} \right)} &= \frac{\Pr^k(\mathfrak{h}' | \mathfrak{h}_i^t, \hat{\theta}_t = \rho) \tilde{\mu}^k(\tilde{a}_{i,t} | \mathfrak{h}')}{\sum_{\tilde{\mathfrak{h}}} \Pr^k(\tilde{\mathfrak{h}} | \mathfrak{h}_i^t, \hat{\theta}_t = \rho) \tilde{\mu}^k(\tilde{a}_{i,t} | \tilde{\mathfrak{h}})} \\ &= \frac{\Pr^k(\mathfrak{h} | \mathfrak{h}_i^t, \hat{\theta}_t = \rho) \tilde{\mu}^k(\tilde{a}_{i,t} | \mathfrak{h})}{\sum_{\tilde{\mathfrak{h}}} \Pr^k(\tilde{\mathfrak{h}} | \mathfrak{h}_i^t, \hat{\theta}_t = \rho) \tilde{\mu}^k(\tilde{a}_{i,t} | \tilde{\mathfrak{h}})} \\ &\leq \frac{1}{(1 - f_2(k))^2} \frac{F^k \left(\mathfrak{h}' | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t} \right)}{F^k \left(\mathfrak{h} | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t} \right)} \leq \frac{\bar{f}_3(k)}{(1 - f_2(k))^2}, \end{aligned}$$

by (15) and the fact that $\tilde{a}_{i,t}$ follows $\tilde{\mu}^k(\tilde{a}_{i,t} | \hat{\mathfrak{h}}_0^t)$. Note that, for each $a_{-i,t}$,

$$\begin{aligned} \Pr^k \left(a_{-i,t} | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t} \right) &\geq \\ &\underbrace{f_2(k)}_{\Pr^k(\theta_t = \pi[t] | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t})} \times \underbrace{f_2(k) (1 - f_2(k))^N}_{\leq \Pr^k(m_i^\ominus = \rho \text{ but } m_j^\ominus = \pi \text{ for all } j \neq i | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t}, \theta_t = \pi[t])} \times \underbrace{(f_1(k))^N}_{\leq \Pr^k(a_{-i,t} | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t}, \theta_t = \pi[t], m_j^\ominus = \pi \text{ for all } j \neq i)} \end{aligned}$$

which asymptotically dominates $\bar{f}_3(k) / (1 - f_2(k))^2$.¹⁸ Hence, for any \mathfrak{h}^t such that $\hat{\mathfrak{h}}_0^t$ lies in $\mathfrak{H}^t[l^*]$ and satisfies (20), player i 's belief on each $\mathfrak{h}' \in \mathfrak{H}^t[l]$ with $l > l^*$ converges to 0.

Therefore, in computing (19), we may restrict attention to $\hat{\mathfrak{h}}_0^t \in \mathfrak{H}^t[l^*]$. As the likelihood ratio of any two such histories is bounded and $f_2(k) \rightarrow 0$, it follows that (19) equals 1.

Case 2: There does not exist $\hat{\mathfrak{h}}_0^t \in \mathfrak{H}^t[l^*]$ and $\tilde{a}_{-i,t}$ such that

$$\lim_k \tilde{\mu}^k \left(\tilde{a}_{-i,t} | \tilde{a}_{i,t}, \hat{\mathfrak{h}}_0^t \right) > 0 \text{ and } p \left(s_{i,t+1} | \hat{\mathfrak{h}}_0^t, a_{i,t}, \tilde{a}_{-i,t} \right) > 0. \quad (21)$$

In this case, we claim that player i believes $\theta_t \neq \rho$:

$$\lim_k \Pr^k (\theta_t = \rho | \mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \rho, 0), a_{i,t}, s_{i,t+1}) = 0.$$

To see why, first recall that player i 's belief on each $\hat{\mathfrak{h}}_0^t \in \mathfrak{H}^t[l]$ for $l \neq l^*$ converges to 0. For each $\hat{\mathfrak{h}}_0^t \in \mathfrak{H}^t[l^*]$, $s_{i,t+1}$ occurs with probability at most $\bar{f}_3(k)$ when $\theta_t = \rho$, while $s_{i,t+1}$

¹⁸Note that the exponent on the $1 - f_2(k)$ and $f_1(k)$ terms equals N rather than $N - 1$ because we are implicitly treating the mediator like any other player at the action stage of each period. As in footnote 16, we omit the details.

occurs with total probability no less than

$$f_2(k) \times f_2(k) (1 - f_2(k))^N \times (f_1(k))^N \times \min_{s_{t+1}, \hat{\mathfrak{h}}^t, a_t: p(s_{t+1} | \hat{\mathfrak{h}}^t, a_t) > 0} p(s_{t+1} | \hat{\mathfrak{h}}^t, a_t),$$

which asymptotically dominates $\bar{f}_3(k) / (1 - f_2(k))^2$. Hence, (19) equals 0.

In total, we have proven (7).

Finally, since $m_{i,t+1}^\ominus = \rho$ implies $\hat{\theta}_{t+1} = \rho$ and $\Pr^k(m_{i,t+1}^\ominus = \rho | \mathfrak{h}^{t+1}, \hat{\theta}_{t+1} = \rho) = 1 - f_2(k) \rightarrow 1$ for any \mathfrak{h}^{t+1} , we have

$$\begin{aligned} & \lim_k \Pr^k(\mathfrak{h}_{-i}^t, m_{-i,t}, a_{-i,t}, s_{-i,t+1} | \mathfrak{h}_i^t, m_{i,t}, a_{i,t}, s_{i,t+1}, (m_{i,t+1}^\ominus, m_{i,t+1}^{\leq t}) = (\rho, 0)) \\ &= \lim_k \Pr^k(\mathfrak{h}_{-i}^t, m_{-i,t}, a_{-i,t}, s_{-i,t+1} | \mathfrak{h}_i^t, m_{i,t}, a_{i,t}, s_{i,t+1}, \hat{\theta}_{t+1} = \rho). \end{aligned}$$

Hence, (12) implies (8).

12.4 Proof of Lemma 5

If player i receives $(m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0)$, she knows that $\theta_t = \pi[t]$, $\hat{\theta}_t = \rho$, and $m_{-i,t}^{\leq t-1} = 0$. She therefore believes that all other players $j \neq i$ received $m_{j,t}^\ominus = \pi$ with probability $(1 - f_2(k))^N$. Moreover, given that $m_{-i,t}^\ominus = \pi$, the conditional distribution of history \mathfrak{h}_{-i}^t is equal to $\lim_k F^k(\hat{\mathfrak{h}}_{-i}^t, \tilde{a}_{-i,t} | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t})$: for faithful $\mathfrak{h}_i^t, \mathfrak{h}_{-i}^t$,

$$\begin{aligned} & \lim_k \Pr^k(\mathfrak{h}_{-i}^t, \tilde{a}_{-i,t} | \mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), m_{-i,t}^\ominus = \pi) \\ &= \lim_k \frac{\Pr^k(\mathfrak{h}_{-i}^t, \tilde{a}_{-i,t} | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t})}{\sum_{\tilde{\mathfrak{h}}_{-i}^t} \Pr^k(\tilde{\mathfrak{h}}_{-i}^t, \tilde{a}_{-i,t} | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t})} \\ &= \lim_k \frac{\Pr^k(\mathfrak{h}_{-i}^t | \mathfrak{h}_i^t, \hat{\theta}_t = \rho) \tilde{\mu}^k(\tilde{a}_{i,t} | \hat{\mathfrak{h}}_i^t)}{\sum_{\tilde{\mathfrak{h}}_{-i}^t} \Pr^k(\tilde{\mathfrak{h}}_{-i}^t | \mathfrak{h}_i^t, \hat{\theta}_t = \rho, \tilde{a}_{i,t}) \tilde{\mu}^k(\tilde{a}_{i,t} | \hat{\mathfrak{h}}_i^t)} \\ &= \lim_k F^k(\hat{\mathfrak{h}}_{-i}^t, \tilde{a}_{-i,t} | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t}) \text{ by (12)}. \end{aligned}$$

Hence,

$$\begin{aligned} & \lim_k \Pr^k(\mathfrak{h}_{-i}^t, (m_{-i,t}^A, m_{-i,t}^\ominus, m_{-i,t}^{\leq t-1}) = (\tilde{a}_{-i,t}, \pi, 0) | \mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0)) \\ &= \lim_k \Pr^k(\mathfrak{h}_{-i}^t, \tilde{a}_{-i,t} | \mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), m_{-i,t}^\ominus = \pi) \\ &= \lim_k F^k(\hat{\mathfrak{h}}_{-i}^t, \tilde{a}_{-i,t} | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t}) = \beta^\psi[t](\hat{\mathfrak{h}}_{-i}^t, \tilde{a}_{-i,t} | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t}), \end{aligned}$$

and (9) holds.

We next prove (10). We first show that a player who receives message $(m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) =$

$(\tilde{a}_{i,t}, \pi, 0)$ continues to believe that her opponents received $m_{-i,t}^\ominus = \pi$ even after observing any sequence of signals $(s_{i,t+1}, \dots, s_{i,\tau+1})$. The intuition is that $m_{-i,t}^\ominus = \pi$ with probability $1 - O(f_2(k))$ given $m_{i,t}^\ominus = \pi$, and when $m_{-i,t}^\ominus = \pi$ each action profile is played with probability $O(f_1(k))$. As $f_1(k) \gg f_2(k)$, no sequence of signals is sufficiently informative to convince player i that $m_{-i,t}^\ominus = \rho$. We actually prove the following slightly stronger claim.

Claim 3 *For any faithful histories \mathfrak{h}^t and \mathfrak{h}_{-i}^τ , actions $a_{i,t}, \dots, a_{i,\tau}$, and signals $s_{i,t+1}, \dots, s_{i,\tau+1}$,*

$$\lim_k \frac{\Pr^k \left(\begin{array}{l} \mathfrak{h}_{-i}^\tau, m_{-i,t}^\ominus = \pi, m_{i,t+1}^{\leq t} = t, s_{i,t+1}, \dots, s_{i,\tau+1} \\ |\mathfrak{h}^t, \tilde{a}_t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), a_{i,t}, \dots, a_{i,\tau} \end{array} \right)}{\Pr^k \left(\begin{array}{l} \mathfrak{h}_{-i}^\tau, m_{i,t+1}^{\leq t} = t, s_{i,t+1}, \dots, s_{i,\tau+1} \\ |\mathfrak{h}^t, \tilde{a}_t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), a_{i,t}, \dots, a_{i,\tau} \end{array} \right)} = 1.$$

The proof of Claim 3 uses the following notation: for each \mathfrak{h}_{-i}^τ that succeeds \mathfrak{h}_{-i}^t , let $\mathfrak{h}_{-i}^{t+1:\tau}$ be the history from period $t+1$ to period τ such that the concatenation of \mathfrak{h}_{-i}^t and $\mathfrak{h}_{-i}^{t+1:\tau}$ is equal to \mathfrak{h}_{-i}^τ . Note that the specification of the mediation range implies that, for all possible \mathfrak{h}_{-i}^τ with $m_{i,t+1}^{\leq t} = t$, we have $(m_{-i,t+1}^A, m_{-i,t+1}^\ominus, m_{-i,t+1}^{\leq t}) = \dots = (m_{-i,\tau}^A, m_{-i,\tau}^\ominus, m_{-i,\tau}^{\leq \tau-1}) = (\star, \pi, t)$, and there is no meaningful communication between players and the mediator from period $t+1$ on. Hence, for each faithful \mathfrak{h}_{-i}^τ , we can identify $\mathfrak{h}_{-i}^{t+1:\tau}$ with a canonical history $\hat{\mathfrak{h}}_{-i}^{t+1:\tau}$. With this notation, $\mathfrak{h}_{-i}^\tau = \left(\mathfrak{h}_{-i}^t, \hat{\mathfrak{h}}_{-i}^{t+1:\tau} \right)$.

Proof. For each faithful \mathfrak{h}_{-i}^τ consistent with $m_{i,t+1}^{\leq t} = t$, we write $\mathfrak{h}_{-i}^\tau = \left(\mathfrak{h}_{-i}^t, \hat{\mathfrak{h}}_{-i}^{t+1:\tau} \right)$ for some canonical history $\hat{\mathfrak{h}}_{-i}^{t+1:\tau}$. Let $(a_t, s_{t+1}, \dots, a_\tau, s_{\tau+1})$ be the sequence of actions and signals corresponding to $\hat{\mathfrak{h}}_{-i}^{t+1:\tau}$. Assuming players are faithful, we have

$$\begin{aligned} & \Pr^k \left(\begin{array}{l} \left(\mathfrak{h}_{-i}^t, \hat{\mathfrak{h}}_{-i}^{t+1:\tau} \right), m_{-i,t}^\ominus = \pi, m_{i,t+1}^{\leq t} = t, s_{i,t+1}, \dots, s_{i,\tau+1} \\ |\mathfrak{h}^t, \tilde{a}_t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), a_{i,t}, \dots, a_{i,\tau} \end{array} \right) \\ &= \Pr^k \left(\begin{array}{l} m_{-i,t}^\ominus = \pi, m_{i,t+1}^{\leq t} = t, a_t, s_{t+1}, \dots, a_\tau, s_{\tau+1} \\ |\mathfrak{h}^t, \tilde{a}_t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), a_{i,t}, \dots, a_{i,\tau} \end{array} \right) \\ &\geq \underbrace{(1 - f_2(k))^N}_{\Pr^k(m_{-i,t}^\ominus = \pi | \mathfrak{h}^t, \tilde{a}_t, (m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\pi, 0))} \times \underbrace{f_1(k)^N}_{\leq \Pr^k(a_{-i,t} | \mathfrak{h}^t, \tilde{a}_t, m_{-i,t}^\ominus = \pi)} \times p(s_{t+1} | \hat{\mathfrak{h}}^t, a_t) \\ &\times \underbrace{\left(1 - \frac{\bar{f}_3(k)}{(f_2(k))^{N+1} (f_1(k))^N} \right)}_{\leq \Pr^k((m_{i,t+1}^A, m_{i,t+1}^\ominus, m_{i,t+1}^{\leq t}) = (\star, \pi, t) | \mathfrak{h}^{t+1}), \text{ by (5)}} \times \prod_{\tau=t+1}^{\tau} \left(\underbrace{f_1(k)^N}_{\leq \Pr^k(a_{-i,\tau} | \mathfrak{h}^t, m_{-i,t}^\ominus = \pi, m_{i,t+1}^{\leq t} = t)} \times p(s_{\tau+1} | \hat{\mathfrak{h}}^\tau, a_\tau) \right). \end{aligned}$$

(Taking into account the possibility that players may be non-faithful would only multiply

this lower bound by $1 - O(f_4(k))$, so this can be safely neglected.) On the other hand,

$$\begin{aligned} & \Pr^k \left(\begin{array}{l} m_{-i,t}^\ominus \neq \pi, m_{i,t+1}^{\leq t} = t, a_t, s_{t+1}, \dots, a_\tau, s_{\tau+1} \\ |\mathfrak{h}^t, \tilde{a}_t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), a_{i,t}, \dots, a_{i,\tau} \end{array} \right) \\ & \leq \underbrace{N f_2(k)}_{\geq \Pr^k(m_{-i,t}^\ominus \neq \pi | \mathfrak{h}^t, \tilde{a}_t, (m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\pi, 0))} \prod_{\tau=t}^{\tau} p \left(s_{\tau+1} | \hat{\mathfrak{h}}^\tau, a_\tau \right). \end{aligned}$$

Hence,

$$\begin{aligned} & \frac{\Pr^k \left(\begin{array}{l} m_{-i,t}^\ominus = \pi, m_{i,t+1}^{\leq t} = t, a_t, s_{t+1}, \dots, a_\tau, s_{\tau+1} \\ |\mathfrak{h}^t, \tilde{a}_t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), a_{i,t}, \dots, a_{i,\tau} \end{array} \right)}{\Pr^k \left(\begin{array}{l} m_{i,t+1}^{\leq t} = t, a_t, s_{t+1}, \dots, a_\tau, s_{\tau+1} \\ |\mathfrak{h}^t, \tilde{a}_t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), a_{i,t}, \dots, a_{i,\tau} \end{array} \right)} \\ & \geq \frac{(1 - f_2(k))^N \left(1 - \frac{\bar{f}_3(k)}{(f_2(k))^{N+1} (f_1(k))^N} \right) f_1(k)^{NT}}{(1 - f_2(k))^N \times \left(1 - \frac{\bar{f}_3(k)}{(f_2(k))^{N+1} (f_1(k))^N} \right) f_1(k)^{NT} + N f_2(k)} \\ & \rightarrow 1. \end{aligned}$$

■

Note that Claim 3 also implies

$$\lim_k \frac{\Pr^k \left(\begin{array}{l} \mathfrak{h}_{-i}^\tau, m_{-i,t}^\ominus = \pi, m_{i,t+1}^{\leq t} = t, s_{i,t+1}, \dots, s_{i,\tau+1} \\ |\mathfrak{h}^t, \tilde{a}_t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), a_{i,t}, \dots, a_{i,\tau} \end{array} \right)}{\Pr^k \left(\begin{array}{l} \mathfrak{h}_{-i}^\tau, m_{i,t+1}^{\leq t} = t, s_{i,t+1}, \dots, s_{i,\tau+1} \\ |\mathfrak{h}^t, \tilde{a}_t, m_{-i,t}^\ominus = \pi, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), a_{i,t}, \dots, a_{i,\tau} \end{array} \right)} = 1. \quad (22)$$

This differs from Claim 3 only in that the denominator conditions on $m_{-i,t}^\ominus = \pi$. Since Claim 3 implies that $m_{-i,t}^\ominus = \pi$ with probability 1, this additional conditioning does not change the conclusion.

It follows that player i 's belief converges to $\beta^\psi[\tau]$, establishing (10):

$$\begin{aligned}
& \lim_k \Pr^k (\mathfrak{h}_{-i}^\tau, (m_{-i,t}^A, m_{-i,t}^\ominus, m_{-i,t}^{\leq t-1}) = (\tilde{a}_{-i,t}, \pi, 0) | \mathfrak{h}_i^t, m_{i,t:\tau} (\tilde{a}_{i,t}), a_{i,t}, s_{i,t+1}, \dots, a_{i,\tau}, s_{i,\tau+1}) \\
&= \lim_k \frac{\left(\Pr^k (\mathfrak{h}_{-i}^t, \tilde{a}_{-i,t} | \mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0)) \right. \\
&\quad \left. \times \Pr^k \left(\mathfrak{h}_{-i}^\tau, m_{-i,t}^\ominus = \pi, m_{-i,t+1}^{\leq t} = t, s_{i,t+1}, \dots, s_{i,\tau+1} \right. \right. \\
&\quad \left. \left. | \mathfrak{h}_i^t, \tilde{a}_{-i,t}, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), a_{i,t}, \dots, a_{i,t+1} \right) \right) \\
&= \lim_k \frac{\left(\Pr^k (\tilde{\mathfrak{h}}_{-i}^t, \tilde{\tilde{a}}_{-i,t} | \mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0)) \right. \\
&\quad \left. \sum_{\tilde{\mathfrak{h}}_{-i}^\tau, \tilde{\tilde{a}}_{-i,t}} \left(\times \Pr^k \left(\tilde{\mathfrak{h}}_{-i}^\tau, m_{-i,t+1}^{\leq t} = t, s_{i,t+1}, \dots, s_{i,\tau+1} \right. \right. \right. \\
&\quad \left. \left. | \mathfrak{h}_i^t, \tilde{\mathfrak{h}}_{-i}^t, \tilde{\tilde{a}}_{-i,t}, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), a_{i,t}, \dots, a_{i,\tau} \right) \right) \\
&= \lim_k \frac{\left(\Pr^k (\mathfrak{h}_{-i}^t, \tilde{a}_{-i,t} | \mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0)) \right. \\
&\quad \left. \times \Pr^k \left(\mathfrak{h}_{-i}^\tau, m_{-i,t+1}^{\leq t} = t, s_{i,t+1}, \dots, s_{i,\tau+1} \right. \right. \\
&\quad \left. \left. | \mathfrak{h}_i^t, \tilde{a}_{-i,t}, m_{-i,t}^\ominus = \pi, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), a_{i,t}, \dots, a_{i,\tau} \right) \right) \\
&= \lim_k \frac{\left(\Pr^k (\tilde{\mathfrak{h}}_{-i}^t, \tilde{\tilde{a}}_{-i,t} | \mathfrak{h}_i^t, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0)) \right. \\
&\quad \left. \sum_{\tilde{\mathfrak{h}}_{-i}^\tau, \tilde{\tilde{a}}_{-i,t}} \left(\times \Pr^k \left(\tilde{\mathfrak{h}}_{-i}^\tau, m_{-i,t+1}^{\leq t} = t, s_{i,t+1}, \dots, s_{i,\tau+1} \right. \right. \right. \\
&\quad \left. \left. | \mathfrak{h}_i^t, \tilde{\mathfrak{h}}_{-i}^t, \tilde{\tilde{a}}_{-i,t}, m_{-i,t}^\ominus = \pi, (m_{i,t}^A, m_{i,t}^\ominus, m_{i,t}^{\leq t-1}) = (\tilde{a}_{i,t}, \pi, 0), a_{i,t}, \dots, a_{i,\tau} \right) \right) \\
&\quad \text{(by Claim 3 and (22))} \\
&= \lim_k \frac{\beta^\psi[t] (\hat{\mathfrak{h}}_{-i}^t, \tilde{a}_{-i,t} | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t}) \beta^\psi[t] (\hat{\mathfrak{h}}_{-i}^{t+1:\tau}, s_{i,t+1}, \dots, s_{i,\tau+1} | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t}, a_{i,t}, \dots, a_{i,\tau})}{\sum_{\tilde{\mathfrak{h}}_{-i}^\tau, \tilde{\tilde{a}}_{-i,t}} \beta^\psi[t] (\tilde{\hat{\mathfrak{h}}}_{-i}^t, \tilde{\tilde{a}}_{-i,t} | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t}) \beta^\psi[t] (\tilde{\hat{\mathfrak{h}}}_{-i}^{t+1:\tau}, \tilde{\mathfrak{h}}_{-i}^\tau, s_{i,t+1}, \dots, s_{i,\tau+1} | \hat{\mathfrak{h}}_i^t, \tilde{\hat{\mathfrak{h}}}_{-i}^t, \tilde{a}_{i,t}, a_{i,t}, \dots, a_{i,\tau})} \\
&= \beta^\psi[t] (\hat{\mathfrak{h}}_{-i}^\tau | \hat{\mathfrak{h}}_i^t, \tilde{a}_{i,t}, a_{i,t}, \dots, a_{i,t+1}, s_{i,t+1}, \dots, s_{i,\tau+1}).
\end{aligned}$$

Finally, (11) follows from (12) and the conditional independence of $m_{i,t+1}^\ominus$ and \mathfrak{h}_0^{t+1} given $\hat{\theta}_{t+1} = \rho$:

$$\begin{aligned}
& \lim_k \Pr^k (\mathfrak{h}_{-i}^t, m_{-i,t}, a_{-i,t}, s_{-i,t+1} | \mathfrak{h}_i^t, m_{i,t}, a_{i,t}, s_{i,t+1}, (m_{i,t+1}^\ominus, m_{i,t+1}^{\leq t}) = (\rho, 0)) \\
&= \lim_k \Pr^k ((a_{\tau-1}, s_\tau, \tilde{a}_\tau, m_\tau)_{\tau=1}^t, a_t, s_{-i,t+1} | \mathfrak{h}_i^{t+1}, (m_{i,t+1}^\ominus, m_{i,t+1}^{\leq t}) = (\rho, 0)) \\
&\quad \text{(by definition of } \mathfrak{h}_{-i}^t \text{ and } \mathfrak{h}_i^t) \\
&= \lim_k \Pr^k ((a_{\tau-1}, s_\tau, \tilde{a}_\tau, m_\tau)_{\tau=1}^t, a_t, s_{t+1} | \mathfrak{h}_i^{t+1}, \hat{\theta}_{t+1} = \rho) \\
&\quad \text{(by conditional independence of } m_{i,t+1}^\ominus \text{ and } \mathfrak{h}_0^{t+1}) \\
&= \lim_k F^k ((a_{\tau-1}, s_\tau, \tilde{a}_\tau)_{\tau=1}^t, a_t, s_{t+1} | (a_{i,\tau-1}, s_{i,\tau}, \tilde{a}_{i,\tau})_{\tau=1}^t, a_{i,t}, s_{i,t+1}) \\
&\quad \text{(by (12)).}
\end{aligned}$$