The Revelation Principle in Multistage Games*

Takuo Sugaya and Alexander Wolitzky

Stanford GSB and MIT

June 18, 2018

Abstract

The *communication revelation principle* of mechanism design states that any outcome that can be implemented using any communication system can also be implemented by a canonical mechanism. In multistage games, we prove that the communication revelation principle holds for *conditional probability perfect Bayesian equilibrium* (CPPBE), but fails for sequential equilibrium. Our main result is that, nonetheless, the following *implementation revelation principle* holds: an outcome is implementable in sequential equilibrium if and only if it is implementable in (canonical) CPPBE. The implementation revelation principle holds only if the mediator is allowed to tremble—otherwise, the set of implementable outcomes is strictly smaller. In the special case of games with adverse selection but no moral hazard, Nash and sequential equilibrium are essentially equivalent, and a virtual-implementation version of the communication revelation principle holds for any standard solution concept.

*For helpful comments, we thank Dirk Bergemann, Drew Fudenberg, Dino Gerardi, Shengwu Li, George Mailath, Roger Myerson, Alessandro Pavan, Harry Pei, Ludovic Renou, Juuso Toikka, and Bob Wilson; seminar participants at Austin, Cowles, Harvard-MIT, LSE, Princeton, Queen Mary, and UCL; and the anonymous referees. Wolitzky acknowledges financial support from the NSF.

1 Introduction

What we will call the *communication 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. While the communication revelation principle is a result about the communication systems required to implement any outcome, its importance mostly comes from its usefulness in characterizing the set of implementable outcomes. The revelation principle thus does two things at once: it characterizes implementable outcomes, and it characterizes the minimal communication system required for implementation.

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 (PBE). And it is unknown 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 contribution of the current paper is to resolve this question by establishing revelation principles for PBE and sequential equilibrium in multistage games. In particular, we will

[1] For dynamic mechanism design, see for example Courty and Li (2000), Battaglini (2005), Eső 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).

show that the communication revelation principal holds for PBE but fails for sequential equilibrium. Nonetheless, our main result establishes that the implementation side of the revelation principle remains valid even for sequential equilibrium.

The key prior paper on the revelation principle ("RP" henceforth) in multistage games is Myerson (1986). In this beautiful paper, Myerson introduces the concept of *sequential communication equilibrium (SCE)*, which is a kind of PBE—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 information, and the mediator recommends actions. Myerson discusses how the logic of the RP suggests that restricting attention to canonical communication structures is without loss of generality, but he does not state a formal RP 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.[2] Myerson's paper also proves that there is an equivalence between conditional probability systems—the key objects 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.[3]

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 in CPPBE and implementation in sequential equilibrium, so that Myerson's characterization still applies under the more restrictive consistency requirement of Kreps and Wilson?

We answer both of these questions in the affirmative: we prove the communication RP

[2] We review Myerson's characterization and the definition of codomination in Section 6.

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

for CPPBE, and we prove that implementation in sequential equilibrium is equivalent to implementation in CPPBE. The first of these results may be viewed as a formalization of ideas implicit in Myerson (1986). The second—which we refer to as the *implementation revelation principle* for sequential equilibrium—is quite subtle and (to us) unexpected. Indeed, we show that the *communication* RP *fails* for sequential equilibrium. We thus show

$$\text{SCE} := \text{canonical CPPBE} = \text{CPPBE} = \text{sequential eqm} \supsetneq \text{canonical sequential eqm.} \quad (1)$$

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—which we call *player sequential equilibrium (PSE)*—that we show is outcome-equivalent to SCE. If instead the mediator cannot tremble—that is, if we insist on treating the mechanism as part of the immutable physical environment, leading to the concept of *machine sequential equilibrium (MSE)*—then it has been shown by Gerardi and Myerson (2007) that there are SCE that cannot be implemented in sequential equilibrium. Thus, if we want even the implementation RP, we have to let the mediator tremble. We will discuss whether letting the mediator tremble is "reasonable" once the solution concepts have been precisely defined. In any case, while equation (1) summarizes our main results (where "sequential eqm" means "PSE"), Figure 1 presents a more complete picture of the relationship among the sets of implementable outcomes for different solution concepts, once the distinction between PSE and MSE is taken into account.[4]

The various failures of the RP for sequential equilibrium (i.e., failure of the communication version when the mediator can tremble, and failure of both versions when he cannot) 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,

[4]For the non-nested relationship between MSE and canonical PSE: Gerardi and Myerson's Example 3 shows canonical PSE $\not\subseteq$ MSE. An example in Section 6 of the current paper shows MSE $\not\subseteq$ canonical PSE.

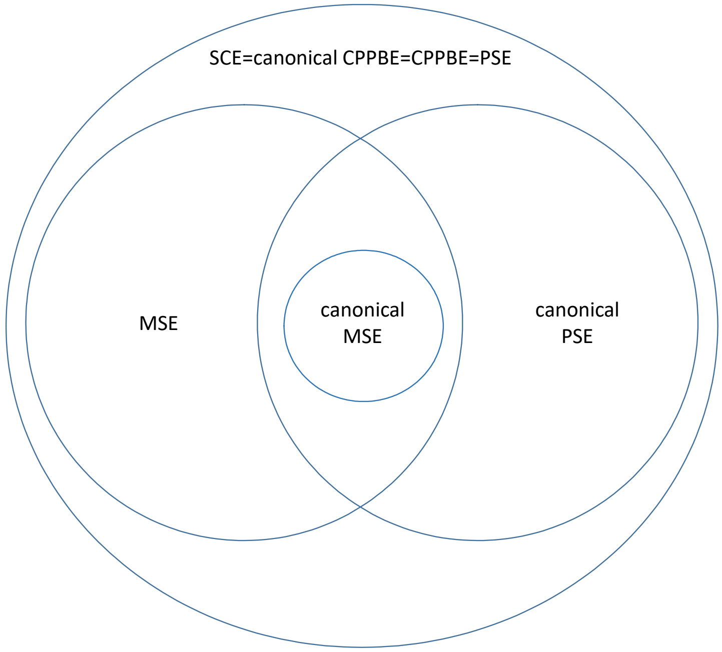

Figure 1: Relationship among implementable outcomes for different solution concepts and communication systems.

we show that Nash, perfect Bayesian, and sequential equilibrium are essentially equivalent, and a virtual-implementation version of the RP always holds—that is, all the sets depicted in Figure 1 coincide up to ε slack. As this class of games encompasses most of the recent literature on dynamic mechanism design, this *virtual revelation principle* may be viewed as another main contribution of our paper.

By way of further motivation for the paper, we note that there seems to be some uncertainty in the literature as to what is known about the RP in multistage 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) are 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 equilib-

rium (PBE) as our solution concept."

Our results provide a foundation for this approach, while also showing that Nash and PBE are essentially outcome-equivalent in pure adverse selection settings like this one.[5]

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

2 Illustrative Examples

Our model and results are at times somewhat complicated and notation-heavy. We therefore start with two examples that give many of the main ideas.

Our first example shows that the communication RP for sequential equilibrium can fail when the mediator cannot tremble.[6] Gerardi and Myerson (2007) have already shown that, without mediator trembles, both the communication and implementation RPs can fail in one-shot games with both adverse selection and moral hazard when some type profiles have 0 probability.[7] We instead consider a two-period game with pure adverse selection. The point is just to set the stage by illustrating a very simple setting where the RP fails.

[5] A 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 believe that our results for PBE concepts should extend to continuous type spaces.

[6] This 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).

[7] In 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.

Example 1: Failure of the Communication RP without Mediator Trembles

The game is one of pure adverse selection, in that 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 equal probability. 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, she 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 her signal in period 2; and player 3 does not communicate with the mediator. Denote player i's report by r_i.

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 falsely reports $r_1 = a$ when $s_1 = b$. 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 her equilibrium strategy at this history, the mediator will take $a_2 = D$ with probability 1. On the other hand, if player 2 misreports her 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 $r_3 \in \{0, 1\}$ at the beginning of the game. Suppose player 3 sends $r_3 = 0$ with equilibrium probability 1, but trembles to sending $r_3 = 1$ 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 $r_1 = a$ *or* $r_3 = 1$, and takes $a_1 = B$ if $r_1 = b$ and $r_3 = 0$. Finally, the mediator takes $a_2 = D$ if and only if $r_1 = a$ and $r_2 = b$.

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 $r_1 = s_1 = b$ but player 3 sent $r_3 = 1$. Player 2 is therefore willing to truthfully report her signal. This behavior in turn deters player 1 from misreporting, as if she misreports 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.

The intuition behind this construction is simply that player 3 trembles "on behalf" of the mediator, who was forbidden from trembling on his own.

Our second example shows that the communication RP can still fail even when the mediator can tremble. This is a novel result. We also describe the logic of the implementation RP in the context of the example. The proof of our main result follows roughly similar lines.

Example 2: Failure of the Communication RP with Mediator Trembles

Again, there are three players (in addition to the mediator) and two periods.

In period 1, players 1 and 2 take actions $(a_1, a_2) \in \{A_1, B_1, C_1\} \times \{A_2, B_2, C_2\}$.

In period 2, player 3 observes signal $s_3 \in \{0, 1\}$, where $s_3 = 0$ if and only if $(a_1, a_2) \in \{(A_1, A_2), (B_1, B_2)\}$. Player 3 then takes an action $a_3 \in \{N, P\}$.

Players 1 and 2's payoffs are given by $u_i(a_1, a_2) - 2 \times \mathbf{1}_{\{a_3 = P\}}$, where u_i $(i = 1, 2)$ is summarized in the following matrix:

	A_2	B_2	C_2
A_1	$0, 0$	$1, 1$	$0, -3$
B_1	$1, 1$	$0, 0$	$0, -3$
C_1	$-3, 0$	$-3, 0$	$-3, -3$

Player 3's payoff $u_3(a_1, a_2, a_3)$ is given by

	$(a_1, a_2) \neq (C_1, C_2)$	$(a_1, a_2) = (C_1, C_2)$
N	1	0
P	0	1

Note that players do not observe their own payoffs (equivalently, they observe them only at the end of the game).

Consider the outcome distribution where players 1 and 2 play (A_1, A_2) and (B_1, B_2) with probability 1/2 each, and player 3 plays A_3. We claim that this outcome is not implementable in a canonical sequential equilibrium (even if the mediator can tremble), by which we mean an equilibrium of the game played with the canonical communication system in which the players report their information truthfully and obey the mediator's recommendations.[8] Yet it can be implemented in a non-canonical equilibrium, using a subtle construction that is generalized by our main result. We sketch the proof of these results here, deferring the details to later in the paper.

Impossibility for canonical sequential equilibrium: Under the target outcome distribution, players 1 and 2 each face a deviation gain of 1 in period 1. Hence, to implement the target outcome, players 1 and 2 must be punished by the play of P in period 2 (with probability at least 1/2) whenever either of them deviates. In turn, for player 3 to be willing to carry out this punishment, she must believe that $(a_1, a_2) = (C_1, C_2)$ with probability at least 1/2 conditional on observing signal $s_3 = 1$ and receiving recommendation $m_3 = P$ from the mediator. The question is, is such a belief consistent?

Clearly, this belief would be consistent if player 3 could believe that players 1 and 2 trembled from their recommended actions to C_i in a correlated manner. But sequential equilibrium requires that players tremble independently conditional on their information—indeed, this is a key difference between sequential equilibrium and PBE. Thus, while the mediator can of course send correlated messages to players 1 and 2, the players must then

[8] As we discuss below, a subtlety in the definition of a canonical equilibrium is whether players must be truthful and obedient even if they have previously lied to the mediator. This issue does not arise in the current example, since each player moves only once.

tremble independently conditional on their messages.

It follows immediately that, in any equilibrium that implements the target distribution (whether canonical or not), player 3 believes that $(a_1, a_2) = (C_1, C_2)$ with probability 0 conditional on observing signal $s_3 = 1$, as C_i is strictly dominated for player $i \in \{1, 2\}$ and (C_1, C_2) requires simultaneous trembles. To complete the argument, we must show that this belief cannot subsequently switch to 1/2 (or more) after receiving recommendation $m_3 = P$. To see this, recall that $m_3 = P$ must be recommended with probability at least 1/2 following a deviation by either player 1 or player 2. Moreover, this probability cannot depend on which action a player deviated to, as the signal s_3 does not distinguish between a deviation to whichever of $\{A_i, B_i\}$ was not recommended and a deviation to C_i, and the deviator cannot be relied on to incriminate herself. Thus, the event $\{s_3 = 1\} \cap \{m_3 = P\}$ must occur with probability at least 1/2 following any deviation by either player 1 or player 2, and a deviation by exactly one of the two players is infinitely more likely than a simultaneous deviation by both of them. Hence, player 3 continues to believe that $(a_1, a_2) = (C_1, C_2)$ with probability 0 even after receiving recommendation $m_3 = P$. So player 3 will not follow the recommendation, a contradiction.

Possibility for non-canonical sequential equilibrium: Why does enriching the communication system overturn this negative result? Suppose the mediator can tremble by handing out a "free pass" to players 1 and 2 at the beginning of the game. If the players get a free pass they are free to take any action in period 1, in that the mediator will always recommend $m_3 = N$ with probability 1. (In particular, this implies the players are always willing to report their actions truthfully.) However, if the players get a free pass and play and report (C_1, C_2) (necessarily as the result of trembles), the mediator may tremble again and recommend $m_3 = P$. Finally, let us assume that, along a sequence of strategy profiles σ^k converging to equilibrium, each mediator tremble (the tremble to handing out the free passes and the tremble to recommending $m_3 = P$) occurs with probability $1/k$ and each player $i = 1, 2$ trembles with probability $1/k$ after receiving a free pass, but each player $i = 1, 2$ trembles with probability only $1/k^5$ when she does not receive a free pass.

Now consider the critical information set where player 3 has received signal $s_3 = 1$ and recommendation $m_3 = P$. There are two possible explanations for how this information set was reached. One is that the mediator behaved "as usual" (i.e., did not hand out free passes) and at least one player $i \in \{1, 2\}$ trembled: this event occurs with probability at most $1/k^5$. The other is that the mediator trembled by handing out free passes (probability $1/k$), players $i = 1, 2$ simultaneously trembled to C_i (probability $1/k^2$), and the mediator trembled again by recommending $m_3 = P$ (probability $1/k$): in total, this event occurs with probability $1/k^4$. Player 3 thus believes that $(a_1, a_2) = (C_1, C_2)$ with probability 1, and she is therefore willing to follow the recommendation. This in turn deters players 1 and 2 from deviating from the target distribution when they do not receive free passes, completing the construction.

(Note however that, after receiving signal $s_3 = 1$ but *before* receiving recommendation $m_3 = P$, player 3 believes that $(a_1, a_2) = (C_1, C_2)$ with probability 0. Thus, while we have shown that it is possible to support the PBE *outcome* $\frac{1}{2}(A_1, A_2) + \frac{1}{2}(B_1, B_2)$ as a sequential equilibrium, not all corresponding PBE *belief systems* can arise in sequential equilibrium.)

For the interested reader, we also note that Example 2 can be given an "economic" interpretation as follows: In the spirit of Kremer, Mansour, and Perry (2014), Che and Hörner (2018), and Hörner and Skrzypacz (2017), suppose the players are drivers choosing commuting routes, and a social planner wants to facilitate social learning. Due to congestion effects, players 1 and 2 want to take different routes in period 1 (i.e., play (A_1, B_2) or (B_1, A_2)), but the planner wants them to take the same route as this generates more information about how travel times respond to congestion. The dominated actions C_1, C_2 correspond to players 1 and 2 quitting their jobs, in which case they have no need to commute in period 2. As for player 3, she can either commute in period 2 (play P) or stay home (play N), where due to congestion effects she is willing to commute only if she believes players 1 and 2 quit in period 1. Finally, again due to congestion, players 1 and 2 are harmed if they do not quit in period 1 but player 3 commutes in period 2. With this interpretation, the analysis above shows that the planner can induce period 1 actions $\frac{1}{2}(A_1, A_2) + \frac{1}{2}(B_1, B_2)$ only if she uses a non-canonical communication system.

The remainder of the paper is organized as follows. Section 3 defines the model and solution concepts. Section 4 presents the (positive and negative) communication RP results, including the virtual RP for pure adverse selection games. Section 5 presents our main result: the implementation RP for PSE. Section 6 reviews Myerson's characterization of SCE. Section 7 summarizes our results. Proofs are deferred to the appendix.

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 $i = 0, 1, \ldots, N$) over T periods (indexed by $t = 1, \ldots, T$). 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.[9] 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 $H^t = \prod_{\tau=1}^{t-1} \left(\prod_{i=0}^{N} S_{i,\tau} \times \prod_{i=1}^{N} R_{i,\tau} \times \prod_{i=1}^{N} M_{i,\tau} \times \prod_{i=0}^{N} A_{i,\tau} \right)$ 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 $\mathring{H}^t = \prod_{\tau=1}^{t-1} \prod_{i=0}^{N} (S_{i,\tau} \times A_{i,\tau})$ 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 $\mathring{h}^t = \left((s_{i,\tau}, a_{i,\tau})_{i=0}^{N} \right)_{\tau=1}^{t-1}$ denote the projection of h^t onto $\mathring{H}^t$. Let $X = \mathring{H}^{T+1} = \prod_{\tau=1}^{T} \prod_{i=0}^{N} (S_{i,\tau} \times A_{i,\tau})$ denote the set of final, payoff-relevant outcomes of the game. Let $u_i : X \to \mathbb{R}$ 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 $p\left(s_t | \mathring{h}^t\right)$, where $\mathring{h}^t \in \mathring{H}^t$ is the

[9] We also use male pronouns for the mediator and female pronouns for the players.

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 $a_{i,t} \in A_{i,t}$.

We refer to the tuple $(N, T, S, A, u, p) = \left(N, T, \prod_{\tau=1}^{T} \prod_{i=0}^{N} S_{i,t}, \prod_{\tau=1}^{T} \prod_{i=0}^{N} A_{i,t}, \prod_{i=0}^{N} u_i, p\right)$ 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 $S_{i,t} = \bigcup_{\mathring{h}^t \in \mathring{H}^t} \operatorname{supp} p_i \left(\cdot | \mathring{h}^t\right)$ for all i, t, where p_i denotes the marginal distribution of p.

For $i \neq 0$, let $H_i^t = \prod_{\tau=1}^{t-1} (S_{i,\tau} \times R_{i,\tau} \times M_{i,\tau} \times A_{i,\tau})$ denote the set of player i's possible histories of signals, reports, messages, and actions at the beginning of period t. Let $H_0^t = \prod_{\tau=1}^{t-1} \left(S_{0,\tau} \times \prod_{i=1}^{N} R \times \prod_{i=1}^{N} M_{i,\tau} \times A_{0,\tau}\right)$ 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 $h_i^t = (s_{i,\tau}, r_{i,\tau}, m_{i,\tau}, a_{i,\tau})_{\tau=1}^{t-1}$ denote the projection of h^t onto H_i^t, and let $h_0^t = \left(s_{0,\tau}, (r_{i,\tau})_{i=1}^{N}, (m_{i,\tau})_{i=1}^{N}, a_{0,\tau}\right)_{\tau=1}^{t-1}$ denote the projection of h^t onto H_0^t. We use the notation $\mathring{h}_i^t$ and $\mathring{h}_0^t$ analogously.

A strategy for player $i \neq 0$ is a function $\sigma_i = \left(\sigma_i^R, \sigma_i^A\right) = \left(\sigma_{i,t}^R, \sigma_{i,t}^A\right)_{t=1}^{T}$, where $\sigma_{i,t}^R : H_i^t \times S_{i,t} \to \Delta(R_{i,t})$ and $\sigma_{i,t}^A : H_i^t \times S_{i,t} \times R_{i,t} \times M_{i,t} \to \Delta(A_{i,t})$. 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 $\beta_i = \left(\beta_i^R, \beta_i^A\right) = \left(\beta_{i,t}^R, \beta_{i,t}^A\right)_{t=1}^{T}$, where $\beta_{i,t}^R : H_i^t \times S_{i,t} \to \Delta(H^t \times S_t)$ and $\beta_{i,t}^A : H_i^t \times S_{i,t} \times R_{i,t} \times M_{i,t} \to \Delta(H^t \times S_t \times R_t \times M_t)$. A strategy for the mediator is a function $\mu = \left(\mu^M, \mu^A\right) = \left(\mu_t^M, \mu_t^A\right)_{t=1}^{T}$, where $\mu_t^M : H_0^t \times S_{0,t} \times R_t \to \Delta(M_t)$ and $\mu_t^A : H_0^t \times S_{0,t} \times R_t \times M_t \to \Delta(A_{0,t})$. 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 $\sigma_{i,t}^A$, $\beta_{i,t}^R$, $\beta_{i,t}^A$, μ_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 (and similarly for $\bar{u}_i(\sigma, \mu | h^t, s^t)$,

$\bar{u}_i\left(\sigma,\mu|h^t,s^t,r^t\right)$, and $\bar{u}_i\left(\sigma,\mu|h^t,s^t,r^t,m^t\right)$). Note that $\bar{u}_i\left(\sigma,\mu|h^t\right)$ does not depend on player i's "beliefs," as h^t is a single node in the game tree.

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 $\bar{u}_i\left(\sigma,\mu\right)\geq\bar{u}_i\left(\sigma_i',\sigma_{-i},\mu\right)$ 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 reports]* For all $i\neq 0,t,\sigma_i'\in\Sigma_i$ and all $\mathfrak{h}_i^t\in H_i^t\times S_{i,t}$,

$$\sum_{\mathfrak{h}^t\in H^t\times S_t}\beta_i\left(\mathfrak{h}^t|\mathfrak{h}_i^t\right)\bar{u}_i\left(\sigma,\mu|\mathfrak{h}^t\right)\geq\sum_{\mathfrak{h}^t\in H^t\times S_t}\beta_i\left(\mathfrak{h}^t|\mathfrak{h}_i^t\right)\bar{u}_i\left(\sigma_i',\sigma_{-i},\mu|\mathfrak{h}^t\right). \tag{2}$$

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

$$\begin{aligned}&\sum_{(\mathfrak{h}^t,r_t,m_t)\in H^t\times S_t\times R_t\times M_t}\beta_i\left(\mathfrak{h}^t,r_t,m_t|\mathfrak{h}_i^t,r_{i,t},m_{i,t}\right)\bar{u}_i\left(\sigma,\mu|\mathfrak{h}^t,r_t,m_t\right)\\ \geq&\sum_{(\mathfrak{h}^t,r_t,m_t)\in H^t\times S_t\times R_t\times M_t}\beta_i\left(\mathfrak{h}^t,r_t,m_t|\mathfrak{h}_i^t,r_{i,t},m_{i,t}\right)\bar{u}_i\left(\sigma_i',\sigma_{-i},\mu|\mathfrak{h}^t,r_t,m_t\right).\end{aligned} \tag{3}$$

- *[Bayes' rule]* For all $i\neq 0$, all $\mathfrak{h}_i^t\in(H_i^t,S_{i,t})$ such that $\Pr^{(\sigma,\mu)}\left(\mathfrak{h}_i^t\right)>0$, and all $\mathfrak{h}^t\in H^t\times S_t$ with i-component $\mathfrak{h}_i^t$,

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

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

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

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).[10] 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.[11]

[10] Recall that a CPS on a finite set Ω is a function $f(\cdot|\cdot) : 2^\Omega \times 2^\Omega \to [0,1]$ 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 $Y \neq \emptyset$, $f(X|Y) f(Y|Z) = f(X|Z)$.

[11] Mailath (2018) defines a notion of "almost perfect Bayesian equilibrium," which appears to coincide with CPPBE in multistage games. Most other notions of "perfect Bayesian equilibrium" (e.g., Fudenberg and Tirole (1991), Watson (2017)) impose some form of "no signaling what you don't know," which is not required by CPPBE.

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)*. As we will see, 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 reports]* For all $i \neq 0, t, \sigma_i'$ and all $\mathfrak{h}_i^t \in H_i^t \times S_{i,t}$, (2) holds.

- *[Sequential rationality of actions]* For all $i \neq 0, t, \sigma_i'$ and all $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) \in H_i^t \times S_{i,t} \times R_{i,t} \times M_{i,t}$, (3) holds.

- *[Player Consistency]* There exists a sequence of completely mixed strategy profiles $\left(\sigma^k, \mu^k\right)_{k=1}^{\infty}$ such that $\lim_{k\to\infty} \left(\sigma^k, \mu^k\right) = (\sigma, \mu)$;

$$\beta_i \left(\mathfrak{h}^t | \mathfrak{h}_i^t\right) = \lim_{k\to\infty} \frac{\Pr^{\left(\sigma^k, \mu^k\right)}\left(\mathfrak{h}^t\right)}{\Pr^{\left(\sigma^k, \mu^k\right)}\left(\mathfrak{h}_i^t\right)}$$

  for all $i \neq 0$, all $\mathfrak{h}_i^t \in H_i^t \times S_{i,t}$, and all $\mathfrak{h}^t \in H^t \times S_t$ with i-component $\mathfrak{h}_i^t$; and

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

  for all $i \neq 0$, all $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) \in H_i^t \times S_{i,t} \times R_{i,t} \times M_{i,t}$, and all $(\mathfrak{h}^t, r_t, m_t) \in H^t \times S_t \times R_t \times M_t$ with i-component $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t})$.

Note that the requirement that $\left(\sigma^k, \mu^k\right)$ 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

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

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 reports]* For all $i \neq 0, t, \sigma_i'$ and all $\mathfrak{h}_i^t \in (H \times S)_i^t(\mu)$, (2) holds.
- *[Sequential rationality of actions]* For all $i \neq 0, t, \sigma_i'$ and all $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) \in (H \times S \times R \times M)_i^t(\mu)$, (3) holds.
- *[Machine Consistency]* There exists a sequence of completely mixed player strategy profiles $\left(\sigma^k\right)_{k=1}^{\infty}$ such that $\lim_{k\to\infty} \sigma^k = \sigma$;

  $$\beta_i\left(\mathfrak{h}^t | \mathfrak{h}_i^t\right) = \lim_{k\to\infty} \frac{\Pr^{\left(\sigma^k,\mu\right)}\left(\mathfrak{h}^t\right)}{\Pr^{\left(\sigma^k,\mu\right)}\left(\mathfrak{h}_i^t\right)}$$

  for all $i \neq 0$, all $\mathfrak{h}_i^t \in (H \times S)_i^t(\mu)$, and all $\mathfrak{h}^t \in H^t \times S_t$ with i-component $\mathfrak{h}_i^t$; and

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

  for all $i \neq 0$, all $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t}) \in (H \times S \times R \times M)_i^t(\mu)$, and all $(\mathfrak{h}^t, r_t, m_t) \in H^t \times S_t \times R_t \times M_t$ with i-component $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t})$.

Let us comment on the "reasonableness" of allowing mediator trembles. "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.) While the more restrictive consistency requirement of MSE may have some conceptual appeal, our finding that the implementation RP holds for PSE but not MSE argues in favor of letting the mediator tremble. If one insists on MSE, a pragmatic approach for applications is to use the RP to characterize the set of PSE-implementable outcomes, and then attempt to verify directly that the "optimal" PSE outcome is also implementable in MSE.[12]

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.[13]

There is however one context where the argument for using MSE is stronger. 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. In this context, letting the mediator tremble would be analogous to letting nature tremble in the unmediated game, which is typically not allowed.[14]

Finally, Gerardi and Myerson (2007; Theorem 3) show that the sets of PSE-implementable

[12] A referee suggests the following alternative interpretation of mediator trembles: Suppose there is a "true" mechanism designer with commitment power and a "metaphorical" mediator who attempts to execute the designer's plan. Then the true designer can be assumed to commit to a plan without error, while the metaphorical mediator may tremble.

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

[14] However, Makris and Renou (2017) consider revelation principles in this "information design" setting and show that mediator trembles are required for implementation in canonical equilibrium.

and MSE-implementable outcome distributions coincide in one-shot games if $N \geq 3$. In this case, the players can "replicate" the mediator's trembles. Their proof does not easily extend to multistage games—we have not been able to resolve this question in general, and doing so would likely require a separate paper.[15] If it does extend, then our results imply that the implementation RP is also valid for MSE when $N \neq 2$.[16]

3.4 Restricted Solution Concepts

The equilibrium definitions given above are the natural ones from a game-theoretic perspective. However, towards introducing the RP, it will be helpful to consider additional restrictions on the mediator's possible messages. These restrictions will let us require that players obey the mediator's recommendations even when the mediator is allowed to tremble. For example, a player will never follow a recommendation to play a strictly dominated action, so obedience and mediator trembles are inconsistent without restricting the mediator's messages. The results in this section will also clarify that letting the mediator tremble can only expand the set of implementable outcomes: that is, $PSE \supseteq MSE$.

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,

[15] Replicating mediator trembles by player trembles for sequential equilibrium in multistage games is related to the well-studied problem of entirely replacing the mediator with cheap talk. This literature has made little progress for sequential equilibrium in multistage games. Gerardi (2004) considers sequential equilibrium in one-shot games. Heller, Solan, and Tomala (2012) consider Nash equilibrium in repeated games. Sugaya and Wolitzky (2017; Proposition 2) consider sequential equilibrium in repeated games, but impose an assumption under which Nash and sequential equilibium coincide.

[16] Gerardi and Myerson's counterexample to the implementation revelation principle is for $N = 2$. We also note two cases in which it is easy to prove their result does extend. The first is when there exists a player whose payoff is independent of the outcome of the game and can thus tremble "on behalf" of the mediator, like player 3 in Example 1. The second is when there exist two players who do not observe signals or take actions, in which case a similar argument applies.

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.

Lemma 1 *For each of the five equilibrium concepts introduced above (NE, WPBE, CPPBE, PSE, and MSE) and for any game G, an outcome distribution $\rho \in \Delta(X)$ arises in an equilibrium of G if and only if there exists a mediation range Q such that ρ arises in an equilibrium of $G|_Q$ (for the same solution concept).*

A similar argument implies that every MSE outcome distribution is also a PSE outcome distribution, as every MSE (σ, μ) is a PSE in the restricted game where the mediation range is given by the set of messages that can arise under some strategy profile (σ', μ) in which the mediator follows his equilibrium strategy.

4 The Communication RP

Our treatment of the communication RP is organized as follows. Section 4.1 defines the communication RP. Section 4.2 presents the communication RP for NE, WPBE, and CPPBE. Section 4.3 records the failure of the communication RP for PSE (and MSE, for which this result was already established by Gerardi and Myerson). Section 4.4 gives conditions under which the communication RP holds for PSE and MSE.

4.1 Definition

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 $m_{i,\tau} \in Q_i\left(\left((a_{i,\tau'-1}, s_{i,\tau'}), a_{i,\tau'}\right)_{\tau'=1}^{\tau-1}, (a_{i,\tau-1}, s_{i,\tau})\right)$ 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 $m_{i,\tau} \in Q_i\left(\left((a_{i,\tau'-1}, s_{i,\tau'}), a_{i,\tau'}\right)_{\tau'=1}^{\tau-1}, (a_{i,\tau-1}, s_{i,\tau})\right)$ for all $\tau \leq t$.

As assessment (σ, μ, β) of a canonical game is *canonical* if the strategy profile (σ, μ) is canonical and in addition each player believes her opponents have been truthful as long as she herself has been truthful (but possibly not obedient): if $(h^t, s_t) \in \operatorname{supp} \beta_{i,t}^R(h_i^t, s_{i,t})$ and, for all $\tau < t$, $r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$ and $m_{i,\tau} \in Q_i\left(\left((a_{i,\tau'-1}, s_{i,\tau'}), a_{i,\tau'}\right)_{\tau'=1}^{\tau-1}, (a_{i,\tau-1}, s_{i,\tau})\right)$, then $r_{j,\tau} = (a_{j,\tau-1}, s_{j,\tau})$ for all $j \neq i$ and all $\tau < t$ (and similarly for histories $(h^t, s_t, r_t, m_t) \in \operatorname{supp} \beta_{i,t}^A(h_i^t, s_{i,t}, r_{i,t}, m_{i,t})$). We will 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.

A classical statement of the RP is as follows:

Communication 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 $C(G)$.

Two remarks are in order:

First, Townsend (1988) extends the RP by requiring a player to be truthful and obedient even if she has previously lied to the mediator, and correspondingly lets a player report her entire history of actions and signals every period (thus giving players opportunities to "confess" any lie). This distinction is irrelevant for our main results.

Second, one could alternatively define a canonical assessment without the requirement that truthful players believe their opponents have been truthful.[17] Using this weaker definition would weaken our positive results and strengthen our negative results. However, our

[17] This alternative definition may be preferrable when considering MSE, as for MSE there are examples where, for some mediator strategies and some histories, a truthful player cannot believe her opponents have been truthful.

negative results (and their proofs) already apply equally for the weaker definition. We use the stronger definition because—as we will see—it is the one that leads to an equivalence with Myerson's approach.

4.2 NE, WPBE, and CPPBE

Forges (1986) established the communication RP for NE.

Proposition 1 (Forges (1986; Proposition 1)) *The communication RP holds for NE.*

As we build on this result, we reprise the proof in our notation in the appendix.

Another preliminary result is the RP for WPBE.

Proposition 2 *The communication RP holds for WPBE.*

This proposition is a fairly straightforward extension of Proposition 1. The required construction combines the strategy profile constructed in the proof of Proposition 1 with the belief system that results from projecting beliefs in the original equilibrium onto the set of payoff-relevant histories.

Our first main result formalizes the communication RP 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 communication RP, 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 communication RP holds for CPPBE.*

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.[18] 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. We review Myerson's characterization in Section 6.

4.3 PSE and MSE

In contrast to the situation for NE or PBE, the communication RP can fail for PSE or MSE.

Proposition 4 *The communication RP does not hold for PSE or MSE.*

As noted in the Introduction, the negative result for MSE is due to Gerardi and Myerson (2007), and applies even in one-shot games (with both adverse selection and moral hazard). Our Example 1 shows that in multistage games the same conclusion applies even without moral hazard.[19]

For PSE, Example 2 proves the result. The details are in the appendix. In addition, the analysis of this example is robust to perturbations of the payoff functions and target outcome distribution. Thus, the failure of the communication RP for PSE is generic.

[18] Myerson's notation is quite different from ours, so superficially his definition of an SCE looks somewhat different from our definition of a canonical CPPBE. The main difference is that he omits reports and messages from the CPS. However, reports can be recovered from players' histories (given that a truthful player believes her opponents have been truthful), and messages can be recovered as a function of the mediator's pure strategy (for this reason, Myerson considers mixed rather than behavior strategies for the mediator; this makes no difference by Kuhn's theorem).

[19] We have also constructed an example showing the communication RP can fail for MSE without adverse selection (i.e., with pure moral hazard).

4.4 Sufficient Conditions for the Communication RP for PSE and MSE

In light of Proposition 4, it is natural to ask when the communication RP for PSE and MSE does hold. We give three simple sufficient conditions.

First, the communication RP holds under a full support condition: any NE outcome distribution under which no player can perfectly detect another's deviation is a canonical MSE outcome distribution. This result may be "folk knowledge" among game theorists, but we are not aware of a reference.

Second, the communication RP holds in single-agent settings. This follows as a trivial corollary of the full-support result. It is however applicable to many models of dynamic moral hazard (e.g., Garrett and Pavan, 2012) and dynamic information design (e.g., Ely, 2017).

Third, 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 communication RP holds for pure adverse selection games. In particular, the distinction between Nash, perfect Bayesian, and sequential equilibrium is essentially immaterial in these games.

To formally state the result, 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 $\mathring{H}_i^{T+1}$. In addition, let $\|\cdot\|$ denote the sup norm on $\Delta(X)$: for distributions $\rho, \rho' \in \Delta(X)$, $\|\rho - \rho'\| = \max_{\mathring{h}^{T+1} \in X} \left|\rho\left(\mathring{h}^{T+1}\right) - \rho'\left(\mathring{h}^{T+1}\right)\right|$.

Proposition 5 *The following hold:*

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

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

3. [Virtual RP for Pure Adverse Selection Games] *If G is a game of pure adverse selection then, 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 $\|\rho - \rho'\| < \varepsilon$.*

The first two parts of the proposition are fairly straightforward. For the third part, the intuition is that perturbing the mediator's strategy in a pure adverse selection game yields an equilibrium satisfying the full support condition of the first part.

The reader may ask why perturbing the mediator's strategy does not always yield a nearby canonical sequential equilibrium even in games with moral hazard. There are two issues. One is that the players' best response correspondence may not be lower hemi-continuous. The other is that perturbing the mediator's strategy may not induce a full support distribution over payoff-relevant histories. For example, no matter how one perturbs the mediator's strategy in Example 2, (C_1, C_2) is played with probability 0.

5 The Implementation RP

Our main result is that, rather surprisingly, the failure of the communication RP for PSE poses no obstacle to the characterization of PSE-implementable outcomes.

Theorem 1 (Implementation Revelation Principle for PSE) *For any base game, an outcome distribution arises in PSE for some message set if and only if it is a canonical CPPBE (equivalently, an 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. Theorem 1 shows it is possible to combine the tractability of Myerson's characterization with Kreps-Wilson consistency.

As every PSE is a CPPBE and the communication RP holds for CPPBE (Proposition 3), one direction of the theorem is immediate. The substance of the theorem is showing that every canonical CPPBE outcome arises as a PSE for some message set. Interestingly, the message set used in our construction is "almost canonical": for all players $i \neq 0$ and periods t, $R_{i,t} = A_{i,t-1} \times S_{i,t}$ and $M_{i,t} = A_{i,t} \cup \{\star\}$. Thus, the only difference from canonical communication is that the mediator has an extra message $\star$ to send to the players. This message corresponds to the "free pass" in Example 2. As in that example, the role of message $\star$ is to signal to players that they should tremble with higher probability. (When a player instead receives a message $m_{i,t} \neq \star$, she is supposed to play $a_{i,t} = m_{i,t}$ as in a canonical equilibrium.) Moreover, the proof proceeds by "replicating" players' beliefs in the original CPPBE *conditional* on their receiving messages other than $\star$. Thus, while it can be impossible to replicate CPPBE beliefs in PSE at every history (as we saw in Example 2), by introducing an extra message we can replicate the desired beliefs at all histories where players receive action recommendations.

Somewhat more precisely, the proof constructs an automaton strategy with two states for the mediator, ρ and π. In state ρ, the mediator plays "as usual": in the context of Example 2, this corresponds to recommending $\frac{1}{2}(A_1, A_2) + \frac{1}{2}(B_1, B_2)$ and N on path, while recommending P if someone deviates in period 1. In state π (which arises only as a result of the mediator's tremble), the mediator sends message $\star$: this corresponds to sending "free passes" and always recommending N. The "double deviation" in Example 2 where the mediator sends free passes but nonetheless recommends P following a deviation by player 1 or 2 corresponds to the mediator trembling to state π in period 1 but then trembling "back" to state ρ in period 2. Finally, the transition probabilities between the states are chosen so that players' beliefs conditional on state ρ (equivalently, conditional on receiving action recommendations) coincide with those in the original CPPBE.

The reader may ask whether the enlargement of the message set in the proof of Theorem 1 is strictly necessary—while Example 2 showed that PSE-implementable outcomes are not always implementable in a canonical equilibria, it did not show that the "problem" is the size of the message set rather than the requirement that players are truthful and obedient. We show in the appendix that enlarging the message set can indeed be necessary.

Proposition 6 *There exist a game and a CPPBE outcome that cannot be implemented in any PSE with message set $M_{i,t} = A_{i,t}$ for all $i \neq 0, t$.*

However, an immediate implication of the preceding discussion is that expanding the message set beyond $M_{i,t} = A_{i,t}$ is unnecessary when each player has at least one codominated action at every information set. In that case, the mediator can use the "recommendation" of the codominated action to mean $\star$.

We end this section with two remarks on alternative solution concepts. First, the equivalence between sequential equilibrium and SCE established in Theorem 1 is lost if one strengthens the PSE concept to MSE. This is shown by Gerardi and Myerson (2007; Example 3). Thus, the implementation RP for sequential equilibrium holds if and only if the mediator can tremble.

Second, Proposition 3 and Theorem 1 jointly imply that the implementation RP holds for any notion of PBE which is stronger than CPPBE but weaker than PSE. Many notions of PBE that impose some form of "no signaling what you don't know" fall into this category, such as PBE satisfying Battigalli's (1996) "independence property" or Watson's (2017) "plain PBE."

6 Myerson's Characterization of SCE

Much of Theorem 1's importance comes from the tractability of Myerson's characterization of SCE. For the reader's convenience, we review his characterization and translate it to our notation.

In this section, message sets are assumed canonical. We denote a typical report $r_{i,t} = (r^A_{i,t-1}, r^S_{i,t}) \in A_{i,t-1} \times S_{i,t}$. We also restrict attention to feasible outcome histories: when we write $\mathring{\mathfrak{h}}^t = ((s_\tau, a_\tau)_{\tau=1}^{t-1}, s_t) \in \mathring{H}^t$, it is assumed that $p(s_{t'}|(s_\tau, a_\tau)_{\tau=1}^{t'-1}) > 0$ for each $t' \le t$.

6.1 Codominated Actions

Intuitively, an action is codominated in period t if, whenever players are truthful and obedient prior to taking actions in period t, it is never optimal for any player to play a codominated action in period t if she believes other players avoid codominated actions in future periods.

Fix $t \in \{1, \ldots, T\}$. Suppose $\mathring{\mathfrak{h}}^t = ((s_\tau, a_\tau)_{\tau=1}^{t-1}, s_t) \in \mathring{H}^t$ is distributed according to $\beta \in \Delta(\mathring{H}^t)$. Let $\mathfrak{h}_0^{t:\tau} = ((s_{0,t'}, r_{t'}, m_{t'}, a_{0,t'})_{t'=t}^{\tau-1}, s_{0,\tau}, r_\tau)$ denote a history for the mediator between periods t and $\tau \ge t+1$. Given $\mathring{\mathfrak{h}}^t$ and $\mathfrak{h}_0^{t:\tau}$, let $(\mathring{\mathfrak{h}}_i^t, \mathfrak{h}_{i,0}^{t:\tau}) = ((r_{i,t'}, m_{i,t'})_{t'=1}^{t-1}, r_{i,t}), (r_{i,t'}, m_{i,t'})_{t'=t}^{\tau-1}, r_{i,\tau}))$, where $r_{i,t'} = (a_{i,t'-1}, s_{i,t'})$ for $t' \le t$ and $m_{i,t'} = a_{i,t'}$ for $t' < t$. In particular, this notation assumes players are truthful and obedient prior to taking actions in period t. Let $\Sigma_i^{t,*}$ denote the set of strategies for player i that are truthful for all $t' \leq t$ and obedient for all $t' < t$.

Let B denote a correspondence that specifies sets $B(\mathring{\mathfrak{h}}_i^t) \subset A_{i,t}$ for each $i \neq 0$, t, and $\mathring{\mathfrak{h}}_i^t$. We say B is a *codomination correspondence* if the following condition is satisfied: whenever (i)

$$\operatorname{supp}\left(\mu_i(\cdot|\mathring{\mathfrak{h}}^t, \mathfrak{h}_0^{t:\tau})\right) \cap B(\mathring{\mathfrak{h}}_i^t, \mathfrak{h}_{i,0}^{t:\tau}) = \emptyset$$

for each $\mathring{\mathfrak{h}}^t$, $i \neq 0$, $\tau \geq t+1$, and $\mathfrak{h}_0^\tau$, and (ii)

$$\operatorname{supp}\left(\mu_i(\cdot|\mathring{\mathfrak{h}}^t)\right) \cap B(\mathring{\mathfrak{h}}_i^t) \neq \emptyset$$

for some $i \neq 0$ and $\mathring{\mathfrak{h}}^t$ with $\beta(\mathring{\mathfrak{h}}^t) > 0$, then there exist $i \neq 0$, a truthful and obedient history $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t})$ with $m_{i,t} \in B(\mathring{\mathfrak{h}}_i^t)$ and $\beta(\mathring{\mathfrak{h}}^t)\mu_i(m_{i,t}|\mathring{\mathfrak{h}}^t) > 0$ for some $\mathring{\mathfrak{h}}^t \in \mathring{H}^t$, and $\sigma_i \in \Sigma_i^{t,*}$ such that

$$\sum_{\mathring{\mathfrak{h}}^t \in \mathring{H}^t, m_t \in A_t} \beta(\mathring{\mathfrak{h}}^t)\mu(m_t|\mathring{\mathfrak{h}}^t)\left(\bar{u}_i(\mu|\mathring{\mathfrak{h}}^t) - \bar{u}_i(\sigma_i, \mu|\mathring{\mathfrak{h}}^t)\right) < 0,$$

where $\bar{u}_i(\mu|\mathring{\mathfrak{h}}^t)$ is i's continuation payoff at $\mathring{\mathfrak{h}}^t$ when she is truthful and obedient and $\bar{u}_i(\sigma_i, \mu|\mathring{\mathfrak{h}}^t)$ is her continuation payoff from σ_i. That is, if the mediator never recommends an action in B after period t but does recommends an action in B in period t with positive probability, then some player has a profitable deviation in period t.

As G is finite, there are finitely many codomination correspondences. One can check that the union of two codomination correspondences is also a codomination correspondence. Let D denote the union of all codomination correspondences. An action $a_i \in A_i^t$ is *codominated* at history $\mathring{\mathfrak{h}}_i^t$ if $a_i \in D(\mathring{\mathfrak{h}}_i^t)$.

Proposition 7 (Myerson (1986; Theorem 2)) *A canonical NE* (σ, μ) *is an SCE if and only if* $\operatorname{supp}(\mu_i(\cdot|\mathring{\mathfrak{h}}^t)) \cap D(\mathring{\mathfrak{h}}_i^t) = \emptyset$ *for each* i.

6.2 Codomination Systems

Myerson also derives a linear programming characterization of codominated actions. Let $A^* = \prod_{i=1}^{N} \prod_{t=1}^{T} \mathbb{R}_+^{\Sigma_i^{t,*} \times A_{i,t} \times \mathring{H}_i^t}$ be the set of all functions $\alpha(\sigma_i|m_{i,t}, \mathring{\mathfrak{h}}_i^t)$ that specify a nonnegative number for every $i \neq 0$, $\sigma_i \in \Sigma_i^{t,*}$, $m_{i,t} \in A_{i,t}$, and $\mathring{\mathfrak{h}}_i^t \in \mathring{H}_i^t$. Given a pure strategy for the mediator, let $f_\tau(\mathring{\mathfrak{h}}^t, \mathfrak{h}_0^{t:\tau}) \in A_\tau$ denote the recommendation profile at history $(\mathring{\mathfrak{h}}^t, \mathfrak{h}_0^{t:\tau})$. Given f, $\mathring{\mathfrak{h}}^t$, and α, define

$$V^t(f, \mathring{\mathfrak{h}}^t, \alpha) = \sum_{i=1}^{N} \sum_{\sigma_i \in \Sigma_i^{t,*}} \alpha(\sigma_i|f_{i,t}(\mathring{\mathfrak{h}}^t), \mathring{\mathfrak{h}}_i^t) \left(\bar{u}_i(f|\mathring{\mathfrak{h}}^t) - \bar{u}_i(\sigma_i, f|\mathring{\mathfrak{h}}^t) \right).$$

The pair (B, α) is a *codomination system* if $\alpha \in A^*$, $B(\mathring{\mathfrak{h}}_i^t) \subset A_{i,t}$ $\forall i, t, \mathring{\mathfrak{h}}_i^t$,

$$\forall i, \forall t, \forall \mathring{\mathfrak{h}}_i^t \in \mathring{H}_i^t, \forall m_{i,t} \in A_{i,t}, \forall \sigma_i \in \Sigma_i^{t,*}, \text{ if } m_{i,t} \notin B(\mathring{\mathfrak{h}}_i^t) \text{ then } \alpha(\sigma_i|f_{i,t}(\mathring{\mathfrak{h}}^t), \mathring{\mathfrak{h}}_i^t) = 0; \text{ and}$$

$$\forall i, \forall t, \forall \mathring{\mathfrak{h}}^t \in \mathring{H}^t, \text{ if } f_\tau(\mathring{\mathfrak{h}}^t, \mathfrak{h}_0^{t:\tau}) \notin D(\mathring{\mathfrak{h}}^t, \mathfrak{h}_0^{t:\tau}) \ \forall \tau \geq t+1 \text{ and } f_{i,t}(\mathring{\mathfrak{h}}^t) \in B(\mathring{\mathfrak{h}}_i^t) \text{ then } V^t(f, \mathring{\mathfrak{h}}^t, \alpha) < 0.$$

Proposition 8 (Myerson (1986; Theorem 4)) *B is a codomination correspondence if and only if there exists some $a \in A^*$ such that (B, α) is a codomination system.*

6.3 An Example

We illustrate Myerson's characterization by combining our Example 2 with his Example 5 (pp. 346-347).

There are 5 players. In period 1, players $i = 1, 2$ take actions $a_i \in \{A_i, B_i, C_i\}$. In period 2, player 3 observes $s_3 \in \{0, 1\}$, where $s_3 = 0$ if and only if $(a_1, a_2) \in \{(A_1, A_2), (B_1, B_2)\}$, and then takes action $a_3 \in \{A_3, B_3, C_3\}$. In period 3, players $i = 4, 5$ take actions $a_i \in \{A_i, B_i, C_i\}$ (without receiving any signals).

Players 1 and 2's payoffs are given by $u_i(a_1, a_2) - 2 \times \mathbf{1}_{\{a_3 = B_3\}} - 10 \times \mathbf{1}_{\{a_3 = C_3\}}$, where $u_i(a_1, a_2)$ is defined as in Example 2.

Player 3's payoff $u_3(a_1, a_2, a_3, a_4, a_5)$ is summarized in the following table:

	$(a_1, a_2) \neq (C_1, C_2)$	$(a_1, a_2) = (C_1, C_2)$
A_3	1	0
B_3	0	1
C_3	-1	-1
	$(a_4, a_5) = (A_4, A_5)$	

A_3	0
B_3	0
C_3	2
$a_4 \neq (A_4, A_5)$	

Finally, player 4's and 5's payoffs are given by $u_i(a_4, a_5)$, defined as follows:

$$\begin{array}{cccc} & A_5 & B_5 & C_5 \\ A_4 & 1,1 & 2,0 & 2,0 \\ B_4 & 0,2 & 3,0 & 0,3 \\ C_4 & 0,2 & 0,3 & 3,0 \end{array} \tag{4}$$

This payoff matrix for players 4 and 5 is the same as in Myerson's Example 5.

We show that an outcome distribution is an SCE (and hence is implementable in CPPBE or PSE) if and only if it assigns positive probability only to actions A_1, B_1, A_2, B_2, A_3, A_4, and A_5. Only a subset of these distributions are implementable in canonical PSE or MSE. Finally, in this example, the set of outcomes implementable in MSE (for some message set) coincides with the entire set of SCE.

SCE: As Myerson noted, B_4, C_4, B_5, and C_5 are codominated. To see why, let $\alpha(A_i|B_i) = \alpha(A_i|C_i) = 1$ for $i = 4, 5$. Then the aggregate incentive values V^3 equal -1 at each of the eight cells other than (A_4, A_5) in the payoff matrix (4). For example, at (B_4, B_5) we have $V^3 = 0 + (3-2) + (0-2) = -1$, and at (A_4, C_5) we have $V^3 = 0 + 0 + (0-1) = -1$.

Given that B_4, C_4, B_5, and C_5 are codominated, C_3 is clearly codominated. After eliminating all of these actions, the remaining game is equivalent to Example 2.

Moreover, after eliminating C_3, C_1 and C_2 are also codominated. This implies that B_3 is not played with positive probability in equilibrium. Note however that B_3 is not codominated. In fact, it is straightforward to check that no other actions are codominated.

It remains to show that any distribution over $\{A_1, B_1\} \times \{A_2, B_2\}$ is implementable. To see this, let players 1 and 2 follow their recommendations, while player 3 play B_3 if and only if $s_3 = 1 \cap (m_1, m_2) \in (A_1, A_2), (B_1, B_2)$. This is clearly a NE. As no player takes a codominated action, by Myerson's theorem it is also an SCE.

Canonical PSE: A small extension of the proof of Proposition 4 shows that not all distributions over $\{A_1, B_1\} \times \{A_2, B_2\}$ are implementable in canonical PSE. (In particular, $\frac{1}{2}(A_1, A_2) + \frac{1}{2}(B_1, B_2)$ is not implementable.)

MSE: In this example, the set of outcome distribution implementable in MSE and PSE coincide. The reason is that, since players 4 and 5 always play (A_4, A_5) in period 3, they are essentially "dummy players" who can tremble on the mediator's behalf in periods 1 and 2.

7 Summary

In lieu of a conclusion, we provide a brief summary of the results and their potential implications for applied mechanism design. In what follows, the negative results for MSE are due to Gerardi and Myerson (2007), while all other results are original to the current paper.

First and foremost, the *implementation RP* holds for PSE (but not for MSE). Therefore, when characterizing implementable outcomes, it is without loss of generality to restrict attention to canonical games and use the CPPBE solution concept, even if one insists on Kreps-Wilson consistency. One must however accept the possibility of trembles by the mediator.

Second, the stronger *communication RP* holds for weak PBE and CPPBE, but not for PSE (or MSE). In particular, while one can characterize the set of PSE-implementable outcomes as CPPBE of the canonical game, it may be necessary to allow one extra message to implement these outcomes as PSE.

Third, there are however some important settings where the communication RP does hold for PSE and MSE. If no player can perfectly detect another's deviation, then Nash, perfect Bayesian, and sequential equilibrium are equivalent. Games with a single agent are a special case. Finally, if the game is one of pure adverse selection (a class that includes much of the dynamic mechanism design literature), then Nash, perfect Bayesian, and sequential equilibrium are essentially equivalent, and the *virtual RP* holds.

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 (2000), "Imperfect Commitment and the Revelation Principle: The Multi-Agent Case," *Economics Letters*, 69, 165-171.

[8] Bester, H. and R. Strausz (2001), "Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case." *Econometrica*, 69, 1077-1098.

[9] Board, S. and A. Skrzypacz (2016), "Revenue Management with Forward-Looking Buyers," *Journal of Political Economy*, 124, 1046-1087.

[10] Che, Y.-K. and J. Hörner (2018), "Recommender Systems as Mechanisms for Social Learning," *Quarterly Journal of Economics*, 133, 871-925.

[11] Conitzer, V. and T. Sandholm (2004), "Computational Criticisms of the Revelation Principle," *Proceedings of the 5th ACM conference on Electronic Commerce*.

[12] Courty, P. and L. Hao (2000), "Sequential Screening," *Review of Economic Studies*, 67, 697-717.

[13] Dhillon, A. and J.-F. Mertens (1996), "Perfect Correlated Equilibria," *Journal of Economic Theory*, 68, 279-302.

[14] Ely, J.C. (2017), "Beeps," *American Economic Review*, 107, 31-53.

[15] Ely, J., A. Frankel, and E. Kamenica (2015), "Suspense and Surprise," *Journal of Political Economy*, 123, 215-260.

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

[17] Eső, P. and B. Szentes (2007), "Optimal Information Disclosure in Auctions and the Handicap Auction," *Review of Economic Studies*, 74, 705-731.

[18] Forges, F. (1986), "An Approach to Communication Equilibria," *Econometrica*, 54, 1375-1385.

[19] Fudenberg, D. and J. Tirole (1991), "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory*, 53, 236-260.

[20] Garrett, D.F. and A. Pavan (2012), "Managerial Turnover in a Changing World," *Journal of Political Economy*, 120, 879-925.

[21] Gerardi, D. (2004), "Unmediated Communication in Games with Complete and Incomplete Information," *Journal of Economic Theory*, 114, 104-131.

[22] Gerardi, D. and R.B. Myerson (2007), "Sequential Equilibria in Bayesian Games with Communication," *Games and Economic Behavior*, 60, 104-134.

[23] Green, J.R. and J.-J. Laffont (1986), "Partially Verifiable Information and Mechanism Design," *Review of Economic Studies*, 53, 447-456.

[24] Heller, Y., E. Solan, and T. Tomala (2012), "Communication, Correlation and Cheap-Talk in Games with Public Information," *Games and Economic Behavior*, 74, 222-234.

[25] Hörner, J. and A. Skrzypacz (2017), "Learning, Experimentation, and Information Design," in *Advances in Economics and Econometrics: Eleventh World Congress: Vol. 1*, B. Honoré, A. Pakes, M. Piazzesi, and L. Samuelson, eds., Cambridge University Press.

[26] Kakade, S.M., I. Lobel, and H. Nazerzadeh (2013), "Optimal Dynamic Mechanism Design and the Virtual-Pivot Mechanism," *Operations Research*, 61, 837-854.

[27] Kohlberg, E. and P.J. Reny (1997), "Independence on Relative Probability Spaces and Consistent Assessments in Game Trees," *Journal of Economic Theory*, 75, 280-313.

[28] Kremer, I., Y. Mansour, and M. Perry (2014), "Implementing the 'Wisdom of the Crowd'," *Journal of Political Economy*, 122, 988-1012.

[29] Kreps, D.M. and R. Wilson (1982), "Sequential Equilibria," *Econometrica*, 50, 863-894.

[30] Mailath, G.J. and L. Samuelson (2006), *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press.

[31] Mailath, G.J. (2018), *Modeling Strategic Behavior*, to be published by World Scientific Press.

[32] Makris, M. and L. Renou (2017), "Information Design in Multi-Stage Games," *working paper*.

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

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

[35] Myerson, R.B. (1986), "Multistage Games with Communication," *Econometrica*, 54, 323-358.

[36] Pavan, A., I. Segal, and J. Toikka (2014), "Dynamic Mechanism Design: A Myersonian Approach," *Econometrica*, 82, 601-653.

[37] Renault, J., E. Solan, and N. Vieille (2017), "Optimal Dynamic Information Provision," *Games and Economic Behavior*, 104, 329-349.

[38] Sugaya, T. and A. Wolitzky (2017), "Bounding Equilibrium Payoffs in Repeated Games with Private Monitoring," *Theoretical Economics*, 12, 691-729.

[39] Townsend, R.M. (1988), "Information Constrained Insurance: The Revelation Principle Extended," *Journal of Monetary Economics*, 21, 411-450.

[40] Watson, J. (2017), "A General, Practicable Definition of Perfect Bayesian Equilibrium," *working paper*.

Appendix: Omitted Proofs

8 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 $\left((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t}\right)$, define an arbitrary mapping $g : M_{i,t} \backslash Q_i \left((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t}\right) \to Q_i \left((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t}\right)$. (Intuitively, when player i receives a message outside Q_i, she will reinterpret it according to g.) Define an assessment $\left(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta}\right)$ 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 $Q_i \left((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t'-1}, r_{i,t'}\right)$ 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 $m_{i,t'} \in M_{i,t'} \backslash Q_i \left((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t'-1}, r_{i,t'}\right)$ 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 $m_{i,t'} \in M_{i,t'} \backslash Q_i \left((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t'-1}, r_{i,t'}\right)$ 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 $\left(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta}\right)$, and $\tilde{\beta}$ coincides with β on $(\tilde{\sigma}, \tilde{\mu})$-positive probability events, so $\left(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta}\right)$ 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 $\left(\sigma^k, \mu^k\right)$ in $G|_Q$ converging to (σ, μ) that generates beliefs β. Let $f_1(k) = \min\left\{\min_{h_i^t} \sigma_i^k\left(h_i^t\right), \min_{h_0^t} \mu^k\left(h_0^t\right)\right\}$. To show that ρ is also a PSE outcome distribution in G, we construct a sequence of trembles in G where the mediator sends messages outside Q_i with probability much less than $f_1(k)$. Each player then believes all messages lie in Q_i $\forall i$ unless she receives a message outside her own mediation range, in which case she may play any best response to her beliefs.

Formally, fix a sequence $f_2(k)$ such that $\lim_k f_2(k)\left(f_1(k)\right)^{-2(N+1)T} = 0$. 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 Q with probability $f_2(k)$ at every history. This constrained game has a PSE $\left(\tilde{\sigma}^k, \tilde{\mu}^k, \tilde{\beta}^k\right)$ by standard arguments (e.g., Kreps and Wilson, 1982). Let $\left(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta}\right) = \lim_k \left(\tilde{\sigma}^k, \tilde{\mu}^k, \tilde{\beta}^k\right)$, taking a convergent subsequence if necessary. By standard arguments, $\left(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta}\right)$ 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 $f_1(k) \to 0$.) 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 Q vanishes relative to the probability of any $2(N+1)T$ trembles that do not involve such a message.) Thus, $\left(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta}\right)$ 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 an MSE of G.

9 Proofs of Propositions 1–3

9.1 Proof of Proposition 1

The proof follows the same logic as the static revelation principle. Forges (1986; Proposition 1) gives essentially the same proof, although somewhat less formally.

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 at every history (even if they have lied to the mediator in the past).

The mediator's strategy is constructed as follows: Denote player i's period t report by $\tilde{r}_{i,t} = (\tilde{a}_{i,t-1}, \tilde{s}_{i,t}) \in A_{i,t-1} \times S_{i,t}$, where $A_{i,0} := \emptyset$. In period 1, given report $\tilde{r}_{i,1}$, the mediator draws a fictitious report $r_{i,1} \in R_{i,1}$ (the set of possible reports in G) according to $\sigma_i^R(\tilde{r}_{i,1})$ (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 $(\tilde{r}_{i,1}, r_{i,1}, m_{i,1})$, the mediator draws an action recommendation $\tilde{m}_{i,1} \in A_{i,1}$ according to $\sigma_i^A(\tilde{r}_{i,1}, r_{i,1}, m_{i,1})$. Finally, the mediator sends message $\tilde{m}_{i,1}$ to player i.[20]

Recursively, for $t = 2, \ldots, T$, let $h_{i,0}^t = (\tilde{r}_{i,\tau}, r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}$ (note that the history $h_{i,0}^t$ combines player i's actual reports $(\tilde{r}_{i,\tau})_{\tau=1}^{t-1}$ and the fictitious reports and messages $(r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}$). Given $h_{i,0}^t$ and player i's period t report $\tilde{r}_{i,t} = (\tilde{a}_{i,t-1}, \tilde{s}_{i,t})$, let

$$h_i^t(h_{i,0}^t, \tilde{r}_{i,t}) = \left((\tilde{s}_{i,\tau}, r_{i,\tau}, m_{i,\tau}, \tilde{a}_{i,\tau})_{\tau=1}^{t-1}, \tilde{s}_{i,t}\right) \tag{5}$$

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

be the corresponding fictitious history for player i. Similarly, given $\mathfrak{h}_{i,0}^t = (\tilde{r}_{i,\tau}, r_{i,\tau}, m_{i,\tau})_{\tau=1}^t$,[21] let

$$\mathfrak{h}_i^t(\mathfrak{h}_{i,0}^t) = ((\tilde{s}_{i,\tau}, r_{i,\tau}, m_{i,\tau}, \tilde{a}_{i,\tau})_{\tau=1}^{t-1}, \tilde{s}_{i,t}, r_{i,t}, m_{i,t}). \tag{6}$$

The mediator draws $r_{i,t} \in R_{i,t}$ according to $\sigma_i^R(h_i^t(h_{i,0}^t, \tilde{r}_{i,t}))$. If $(\tilde{r}_{i,\tau})_{\tau=1}^t$ is not a possible history for player i (that is, if there is no $\mathring{\mathfrak{h}}_{-i}^t$ such that, for $\mathring{\mathfrak{h}}_i^t = (\tilde{r}_{i,\tau})_{\tau=1}^t$, we have $\Pr^{(\sigma',\mu')}(\mathring{\mathfrak{h}}_i^t, \mathring{\mathfrak{h}}_{-i}^t) > 0$ for some (σ', μ')), 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 $\mathfrak{h}_{i,0}^t = (\tilde{r}_{i,\tau}, r_{i,\tau}, m_{i,\tau})_{\tau=1}^t$, the mediator draws $\tilde{m}_{i,t} \in A_{i,t}$ according to $\sigma_i^A(\mathfrak{h}_i^t(\mathfrak{h}_{i,0}^t))$. (Again, if $(\tilde{r}_{i,\tau})_{\tau=1}^t$ is not a possible history, the mediator draws $\tilde{m}_{i,t}$ randomly.) The mediator sends message $\tilde{m}_{i,t}$ 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 $(\mathring{h}_i^t, s_{i,t})$. (ii) If player i has a profitable deviation at a history $(h_i^t, s_{i,t}, \tilde{r}_{i,t}, \tilde{m}_{i,t})$ in $C(G)$ where she has been truthful up to and including period t, then there is a profitable deviation in G in which she deviates whenever her payoff-relevant history equals $(\mathring{h}_i^t, s_{i,t})$ and her equilibrium strategy assigns positive probability to action $\tilde{m}_{i,t}$. Therefore, the fact that σ_i is optimal in G implies that $\tilde{\sigma}_i$ is optimal in $C(G)$.

9.2 Proof of Proposition 2

The construction of strategies is the same as in Proposition 1, except for a player's behavior after she has lied to the mediator. Beliefs are given by Bayes' rule on path, and off path are specified so that (i) a player always believes her opponents tell the mediator the truth, and (ii) beliefs about the payoff-relevant history are taken to equal beliefs in the original, non-canonical game at any history with the same payoff-relevant component. Given this specification, players who have previously lied to the mediator may take any best responses to their beliefs without affecting the incentives of players who so far have always told the truth.

Formally, 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 Proposition 1, leaving unspecified for the moment a player's strategy at histories where she has previously lied to the mediator. Let $\tilde{Q}$ be the mediation range that restricts the mediator to sending recommendations in $\operatorname{supp} \tilde{\mu}_i$: that is, for all $i \neq 0, t, (r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, r_{i,t}$,

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

We thus view $(\tilde{\sigma}, \tilde{\mu})$ as a strategy profile in game $C(G)|_{\tilde{Q}}$. To construct a belief system $\tilde{\beta}$

[21] Note that $\mathfrak{h}_{i,0}^t = h_{i,0}^{t+1}$. We use distinct notation to clarify that this history arises prior to the play of period t actions.

in $C(G)|_{\tilde{Q}}$, if $\Pr^{(\sigma,\mu)}(\mathring{h}_i^t, s_{i,t}) > 0$ then (i) let

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

(thus, $\tilde{\beta}_i\left(h_i^t, s_{i,t}\right)|_{\mathring{H}^t\times S_t}$ is the conditional belief about $(\mathring{h}^t, s_t)$ given $(\mathring{h}_i^t, s_{i,t})$ in the original equilibrium), and (ii) let $\tilde{\beta}_i\left(h_i^t, s_{i,t}\right)$ assign probability 1 to the event that players $-i$ have always been truthful and obedient. If instead $\Pr^{(\sigma,\mu)}(\mathring{h}_i^t, s_{i,t}) = 0$, then (i) let $\tilde{\beta}_i\left(h_i^t, s_{i,t}\right)|_{\mathring{H}^t\times S_t} = \beta_i\left(\bar{h}_i^t, s_{i,t}\right)|_{\mathring{H}^t\times S_t}$ for an arbitrary history $\bar{h}_i^t$ in G with $\mathring{h}_i^t = \mathring{\bar{h}}_i^t$ and $r_{i,\tau} \in \operatorname{supp}\sigma_i^R\left(\bar{h}_i^\tau, s_{i,\tau}\right)$ for all $\tau < t$, and (ii) let $\tilde{\beta}_i\left(h_i^t, s_{i,t}\right)$ assign probability 1 to the event that players $-i$ have always been truthful (but not necessarily obedient). Construct beliefs for histories of the form $(h_i^t, s_{i,t}, \tilde{r}_{i,t}, \tilde{m}_{i,t})$ analogously, where if $\Pr^{(\sigma,\mu)}(\mathring{h}_i^t, s_{i,t}, \tilde{r}_{i,t}, \tilde{m}_{i,t}) = 0$ then $\tilde{\beta}_i\left(h_i^t, s_{i,t}, \tilde{r}_{i,t}, \tilde{m}_{i,t}\right)|_{\mathring{H}^t\times S_t} = \beta_i\left(\bar{h}_i^t, s_{i,t}, r_{i,t}, m_{i,t}\right)|_{\mathring{H}^t\times S_t}$ for an arbitrary history $\left(\bar{h}_i^t, s_{i,t}, r_{i,t}, m_{i,t}\right)$ with $\mathring{h}_i^t = \mathring{\bar{h}}^t$ and $r_{i,\tau} \in \operatorname{supp}\sigma_i^R\left(\bar{h}_i^\tau, s_{i,\tau}\right)$ for all $\tau \le t$ and $\tilde{m}_{i,t} \in \operatorname{supp}\sigma_i^A\left(\bar{h}_i^t, s_{i,t}, r_{i,t}, m_{i,t}\right)$. Beliefs at histories where a player has lied to the mediator may be specified arbitrarily. Finally, specify that a player plays an arbitrary best responses to her belief at any history where she has previously lied to the mediator.

We claim that any assessment $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta})$ so constructed is a restricted WPBE. To see this, it is straightforward to check the following claims: (i) If a previously truthful player i has a profitable misreport in $C(G)|_{\tilde{Q}}$ at history $(h_i^t, s_{i,t})$ with belief $\tilde{\beta}_i\left(h_i^t, s_{i,t}\right)$, then she has a profitable deviation in G at some history $\left(\bar{h}_i^t, s_{i,t}\right)$ with belief $\beta_i\left(\bar{h}_i^t, s_{i,t}\right)$. (ii) If a previously truthful player i has a profitable deviation in $C(G)|_{\tilde{Q}}$ at history $(h_i^t, s_{i,t}, \tilde{r}_{i,t}, \tilde{m}_{i,t})$ with belief $\tilde{\beta}_i\left(h_i^t, s_{i,t}, \tilde{r}_{i,t}, \tilde{m}_{i,t}\right)$, then she has a profitable deviation in G at some history $\left(\bar{h}_i^t, s_{i,t}, r_{i,t}, m_{i,t}\right)$ with belief $\beta_i\left(\bar{h}_i^t, s_{i,t}, r_{i,t}, m_{i,t}\right)$. 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 $C(G)|_{\tilde{Q}}$. Finally, $\tilde{\beta}|_{\mathring{H}^{T+1}}$ coincides with $\beta|_{\mathring{H}^{T+1}}$ on positive-probability histories, so $\tilde{\beta}$ is consistent with Bayes' rule.

9.3 Proof of Proposition 3

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 $\alpha = \left(\alpha_t^R, \alpha_t^M, \alpha_t^A\right)_{t=1}^T$, where $\alpha_t^R : H^t \times S_t \to \Delta(R_t)$, $\alpha_t^M : H^t \times S_t \times R_t \to \Delta(M_t)$, and $\alpha_t^A : H^t \times S_t \times R_t \times M_t \to \Delta(A_t)$, 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 reported signals $\tilde{r}_1^S := (s_{0,1}, \left(\tilde{r}_{i,1}^S\right)_{i\neq 0}) \in S_1$ is in $\operatorname{supp} p(\cdot|\emptyset)$, then the mediator draws a vector of fictitious reports $r_1 \in R_1$ according to $\alpha_1^R(\tilde{s}_1)$ with $\tilde{s}_1 = \tilde{r}_1^S$. Given the resulting vector r_1, the mediator draws a vector of fictitious messages

$m_1 \in M_1$ according to $\alpha_1^M(\tilde{s}_1, r_1)$ with $\tilde{s}_1 = \tilde{r}_1^S$.

If instead $\tilde{r}_1^S \notin \operatorname{supp} p(\cdot|\emptyset)$, then for each $i \neq 0$ the mediator draws $r_{i,1}$ according to $\sigma_i^R(\tilde{r}_{i,1}^S)$. 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 $\tilde{r}_{i,1}^S, r_{i,1}, m_{i,1}$, the mediator draws $\tilde{m}_{i,1}$ according to $\sigma_i^A(\tilde{r}_{i,1}^S, r_{i,1}, m_{i,1})$. Finally, the mediator sends message $\tilde{m}_{i,1}$ to player i.

Recursively, given $h_{i,0}^t = (\tilde{r}_{i,\tau}, r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}$ and $\tilde{r}_{i,t} \in A_{i,t-1} \times S_{i,t}$ for all $i \neq 0$, if $(\tilde{r}_\tau)_{\tau=1}^t$ is a possible history (that is, if $\Pr^{(\sigma',\mu')}((\tilde{r}_\tau)_{\tau=1}^t) > 0$ for some (σ', μ')) then the mediator draws $r_t \in R_t$ according to $\alpha_t^R(h_{I,0}^t, a_{0,t-1}, s_{0,t}, \tilde{r}_t^S)$, where

$$h_{I,0}^t := \left(\left(a_{0,\tau-1}, (\tilde{r}_{i,\tau-1}^A)_{i=1}^N\right), \left(s_{0,\tau}, (\tilde{r}_{i,\tau}^S)_{i=1}^N\right), (r_{i,\tau})_{i=1}^N, (m_{i,\tau})_{i=1}^N\right)_{\tau=1}^{t-1}.$$

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

If instead $(\tilde{r}_\tau)_{\tau=1}^t$ is not a possible history, then for each $i \neq 0$, given $h_{i,0}^t = (\tilde{r}_{i,\tau}, r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}$ and $\tilde{r}_{i,t} \in A_{i,t-1} \times S_{i,t}$, the mediator draws $r_{i,t} \in R_{i,t}$ according to $\sigma_i^R(h_i^t(h_{i,0}^t, \tilde{r}_{i,t}))$ (where $h_i^t(h_{i,0}^t, \tilde{r}_{i,t})$ is defined by (5)). 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 $\mathfrak{h}_{i,0}^t = (h_{i,0}^t, \tilde{r}_{i,t-1}^A, \tilde{r}_{i,t}^S, r_{i,t}, m_{i,t})$, the mediator draws $\tilde{m}_{i,t}$ according to $\sigma_i^A(\mathfrak{h}_i^t(\mathfrak{h}_{i,0}^t))$ (where $\mathfrak{h}_i^t(\mathfrak{h}_{i,0}^t)$ is defined by (6)). (If $(\tilde{r}_{i,\tau})_{\tau=1}^t$ is not a possible history, the mediator draws $\tilde{m}_{i,t}$ randomly.) The mediator sends message $\tilde{m}_{i,t}$ to player i.

Given this construction of $\tilde{\mu}$, let $\tilde{Q}$ be the mediation range that restricts the mediator to sending recommendations in $\operatorname{supp}\tilde{\mu}$ (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 probability distribution $F(\alpha)$ 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 probability distributions induced by completely mixed move distributions. Let $(\alpha^k)_k$ be a sequence of completely mixed move distributions in G such that $\lim_k F(\alpha^k) = f$.

To construct a belief system $\tilde{\beta}$, we begin by constructing a sequence of move distributions $(\tilde{\alpha}^k)_k$ in $C(G)|_{\tilde{Q}}$ as follows:

First, construct a strategy for the mediator $\tilde{\mu}^k$ in $C(G)|_{\tilde{Q}}$ by following the construction of $\tilde{\mu}$ 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 $h_0^t = (s_{0,\tau}, r_\tau, m_\tau, a_{0,\tau})_{\tau=1}^{t-1}$, $r_\tau \in R_\tau$, and $m_\tau \in M_\tau$), if $m_t \in \tilde{Q}(h_0^t, s_{0,t}, r_t)$ then $\tilde{\mu}^k(m_t|h_0^t, s_{0,t}, r_t) > 0$. That is, every message in the mediation range is received with positive probability at every fictitious history.

Second, for each fictitious history $(h_0^t, s_{0,t}, r_t, m_t)$, let $\phi^k(\cdot|h_0^t, s_{0,t}, r_t, m_t, \tilde{a}_t)_{\tilde{a}_t\in A_t, k\in\mathbb{N}}$ be a family of full-support probability distributions on A_t such that (i) $\lim_k \phi^k(a_t|h_0^t, s_{0,t}, r_t, m_t, \tilde{a}_t) =$

1 whenever $a_t = \tilde{a}_t$, and (ii)

$$\sum_{\tilde{a}_t \in A_t} \phi^k \left(a_t | h_0^t, s_{0,t}, r_t, m_t, \tilde{a}_t\right) \alpha^A \left(\tilde{a}_t | h_0^t, s_{0,t}, r_t, m_t\right) = \alpha^{A,k} \left(a_t | h_0^t, s_{0,t}, r_t, m_t\right) \text{ for all } a_t \in A_t. \tag{7}$$

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

(One way to see that such a family ϕ^k exists is to construct it explicitly. To do so, let

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

Note that $d^k > 0$ for all k, $\lim_k d^k = 0$, and $d^k < \alpha^{A,k} (\tilde{a}_t | h_0^t, s_{0,t}, r_t, m_t)$ for all $\tilde{a}_t \in A_t$. For every profile $\tilde{a}_t \in A_t$ with $\alpha^A (\tilde{a}_t | h_0^t, s_{0,t}, r_t, m_t) > 0$, let

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

Then, for every k such that $d^k < \min_{\tilde{a}_t \in A_t : \alpha^A\left(\tilde{a}_t | h_0^t, s_{0,t}, r_t, m_t\right) > 0} \alpha^A (\tilde{a}_t | h_0^t, s_{0,t}, r_t, m_t)$, we see that $\phi^k (\cdot | h_0^t, s_{0,t}, r_t, m_t, \tilde{a}_t)$ is a full-support distribution on A_t. Moreover, it is straightforward to check that $\lim_k \phi^k (a_t | h_0^t, s_{0,t}, r_t, m_t, \tilde{a}_t) = 1$ whenever $a_t = \tilde{a}_t$ and (7) holds.)

Third, define players' action trembles in $C(G)|_{\tilde{Q}}$ by

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

Fourth, let $\tilde{\alpha}^k$ be the move distribution in $C(G)|_{\tilde{Q}}$ that results when the mediator follows strategy $\tilde{\mu}^k$, players report truthfully, and players take actions according to $\tilde{\alpha}^{A,k}(\mathring{h}^t, \mathring{s}_t, \tilde{m}_t)$. Note that $\tilde{\alpha}^k$ is completely mixed in the "truthful game" in which players are required to report truthfully. Fifth, define a CPS $\tilde{f}$ in the truthful game by $\tilde{f} := \lim_k F\left(\tilde{\alpha}^k\right)$. Finally, extend $\tilde{f}$ to a belief system in $C(G)|_{\tilde{Q}}$ by taking any sequence of trembles in which misreports are sufficiently unlikely compared to the lowest probability move in $\tilde{\alpha}^k$. Let $\tilde{\beta}$ be the corresponding belief system.

We claim that $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta})$ is a restricted CPPBE. We first argue that honesty (resp., obedience) is optimal in $C(G)|_{\tilde{Q}}$ for any fictitious history $\left(h_{i,0}^t, \tilde{r}_{i,t}\right))$ (resp., $\mathfrak{h}_{i,0}^t$), 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 $C(G)|_{\tilde{Q}}$ when the mediator's history is $\left(h_{i,0}^t, \tilde{r}_{i,t}\right)$, then there is is a profitable deviation in G in which she misreports at history $h_i^t\left(h_{i,0}^t, \tilde{r}_{i,t}\right)$ (see (5)). (ii) If a previously truthful player i has a profitable deviation in $C(G)|_{\tilde{Q}}$ when the mediator's history is $\mathfrak{h}_{i,0}^t$, then there is a profitable deviation in G in which she deviates at history $h_i^t\left(\mathfrak{h}_{i,0}^t\right)$ (see (6)). 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 $C(G)|_{\tilde{Q}}$ 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 $C(G)|_{\tilde{Q}}$, it follows that $(\tilde{\sigma}, \tilde{\mu}, \tilde{\beta})$ is sequentially rational in $C(G)|_{\tilde{Q}}$. Finally, $\tilde{\beta}$ is derived from a CPS by construction.

10 Proof of Proposition 4

We show $\frac{1}{2}(A_1, A_2) + \frac{1}{2}(B_1, B_2)$ cannot be implemented in a canonical PSE in Example 2.

Suppose towards a contradiction that this distribution is implementable in a canonical PSE (σ, μ, β). Let $\left(\sigma^k, \mu^k\right)_k$ denote the corresponding sequence of trembles. As discussed in the text, player 3 must assign probability at least 1/2 to the event $(a_1, a_2) = (C_1, C_2)$ conditional on $s_3 = 1 \cap m_3 = P$: that is,

$$\lim_k \frac{\sum_{m_1,m_2,r_1,r_2} \left(\begin{array}{c} \mu^k(m_1, m_2) \\ \times \sigma_1^k(C_1|m_1)\sigma_2^k(C_2|m_2)\sigma_1^k(r_1|m_1, C_1)\sigma_2^k(r_2|m_2, C_2) \\ \times \mu^k(m_3 = P|m_1, m_2, r_1, r_2, s_3 = 1) \end{array} \right)}{\sum_{\substack{m_1,m_2,a_1,a_2,r_1,r_2 \\ (a_1,a_2)\notin\{(A_1,A_2),(B_1,B_2)\}}} \left(\begin{array}{c} \mu^k(m_1, m_2) \\ \times \sigma_1^k(a_1|m_1)\sigma_2^k(a_2|m_2)\sigma_1^k(r_1|m_1, a_1)\sigma_2^k(r_2|m_2, a_2) \\ \times \mu^k(m_3 = P|m_1, m_2, r_1, r_2, s_3 = 1) \end{array} \right)}$$

must exceed 1/2. Call this player 3's "critical belief."

Taking $\sigma_1^k(r_1|m_1, C_1)$, $\sigma_2^k(r_2|m_2, C_2)$, and $\mu^k(m_3 = P|m_1, m_2, r_1, r_2, s_3 = 1)$ equal to 1 in the numerator and ignoring some positive terms in the denominator, this belief is no greater than

$$\lim_k \frac{\sum_{m_1,m_2} \mu^k(m_1, m_2)\sigma_1^k(C_1|m_1)\sigma_2^k(C_2|m_2)}{\begin{array}{c} \sum_{m_1,m_2,r_1,r_2} \left(\begin{array}{c} \mu^k(m_1, m_2) \\ \times \sigma_1^k(C_1|m_1)\sigma_2^k(a_2 = m_2|m_2)\sigma_1^k(r_1|m_1, C_1)\sigma_2^k(r_2|m_2, a_2 = m_2) \\ \times \mu^k(m_3 = P|m_1, m_2, r_1, r_2, s_3 = 1) \end{array} \right) \\ + \sum_{m_1,m_2,r_1,r_2} \left(\begin{array}{c} \mu^k(m_1, m_2) \\ \times \sigma_1^k(a_1 = m_1|m_1)\sigma_2^k(C_2|m_2)\sigma_1^k(r_1|m_1, a_1 = m_1)\sigma_2^k(r_2|m_2, C_2) \\ \times \mu^k(m_3 = P|m_1, m_2, r_1, r_2, s_3 = 1) \end{array} \right) \end{array}}.$$

Now, if player 1 plays $a_1 = \{A_1, B_1\} \setminus \{m_1\}$ when $m_1 = m_2 \in \{A_1, B_1\}$, she obtains an expected deviation gain of 1. Hence, conditional on this event, and conditional on any report by player 1, the mediator must recommend P with probability at least 1/2: that is, if $m_1 = m_2 \in \{A_1, B_1\}$ then for sufficiently large k

$$\sum_{r_2} \sigma_2^k(r_2|m_2, a_2 = m_2)\mu^k(m_3 = P|m_1, m_2, r_1, r_2, s_3 = 1) \geq \frac{1}{2} \quad \forall r_1 \in R_1.$$

Similarly, to prevent player 2 from deviating, if $m_1 = m_2 \in \{A_1, B_1\}$ then for sufficiently large k

$$\sum_{r_1} \sigma_1^k(r_1|m_1, a_1 = m_1)\mu^k(m_3 = P|m_1, m_2, r_1, r_2, s_3 = 1) \geq \frac{1}{2} \quad \forall r_2 \in R_2.$$

Finally, for $i = 1, 2$, let $\bar{R}_i := \{r_i : \sigma_i(r_i|m_i, C_i) > 0 \text{ for some } m_i \in \{A_i, B_i\}\}$, and let $\underline{\sigma}_i^R := \min_{r_i, m_i : \sigma_i(r_i|m_i, C_i) > 0} \sigma_i(r_i|m_i, C_i)$.

Therefore, as $\mu^k \to \frac{1}{2}(A_1, A_2) + \frac{1}{2}(B_1, B_2)$, for sufficiently large k

$$\begin{aligned}
& \sum_{m_1, m_2, r_1, r_2} \begin{pmatrix} \mu^k(m_1, m_2) \\ \times \sigma_1^k(C_1|m_1)\sigma_2^k(a_2 = m_2|m_2)\sigma_1^k(r_1|m_1, C_1)\sigma_2^k(r_2|m_2, a_2 = m_2) \\ \times \mu^k(m_3 = P|m_1, m_2, r_1, r_2, s_3 = 1) \end{pmatrix} \\
\geq\ & \frac{1}{2} \sum_{m_1 = m_2} \sum_{r_1 \in \bar{R}_1} \begin{pmatrix} \mu^k(m_1, m_2)\sigma_1^k(C_1|m_1)\sigma_2^k(a_2 = m_2|m_2)\sigma_1^k(r_1|m_1, C_1) \\ \times \sum_{r_2} \sigma_2^k(r_2|m_2, a_2 = m_2)\mu^k(m_3 = P|m_1, m_2, r_1, r_2, s_3 = 1) \end{pmatrix} \\
\geq\ & \frac{1}{4} \sum_{m_1 = m_2} \sum_{r_1 \in \bar{R}_1} \mu^k(m_1, m_2)\sigma_1^k(C_1|m_1)\sigma_1^k(r_1|m_1, C_1) \\
\geq\ & \frac{1}{4} \sum_{m_1 = m_2} \mu^k(m_1, m_2)\sigma_1^k(C_1|m_1)\underline{\sigma}_1^R.
\end{aligned}$$

Similarly,

$$\begin{aligned}
& \sum_{m_1, m_2, r_1, r_2} \begin{pmatrix} \mu^k(m_1, m_2) \\ \times \sigma_1^k(a_1 = m_1|m_1)\sigma_2^k(C_2|m_2)\sigma_1^k(r_1|m_1, a_1 = m_1)\sigma_2^k(r_2|m_2, C_2) \\ \times \mu^k(m_3 = P|m_1, m_2, r_1, r_2, s_3 = 1) \end{pmatrix} \\
\geq\ & \frac{1}{4} \sum_{m_1 = m_2} \mu^k(m_1, m_2)\sigma_2^k(C_2|m_2)\underline{\sigma}_2^R.
\end{aligned}$$

Hence, player 3's critical belief is no greater than

$$\begin{aligned}
& \lim_k \frac{4 \sum_{m_1, m_2} \mu^k(m_1, m_2)\sigma_1^k(C_1|m_1)\sigma_2^k(C_2|m_2)}{\begin{pmatrix} \sum_{m_1 = m_2} \mu^k(m_1, m_2)\sigma_1^k(C_1|m_1)\underline{\sigma}_1^R \\ + \sum_{m_1 = m_2} \mu^k(m_1, m_2)\sigma_2^k(C_2|m_2)\underline{\sigma}_2^R \end{pmatrix}} \\
\leq\ & \frac{4}{\min\{\underline{\sigma}_1^R, \underline{\sigma}_2^R\}} \lim_k \frac{\sum_{m_1, m_2} \mu^k(m_1, m_2)\sigma_1^k(C_1|m_1)\sigma_2^k(C_2|m_2)}{\sum_{m_1 = m_2} \mu^k(m_1, m_2)\sigma_1^k(C_1|m_1) + \sum_{m_1 = m_2} \mu^k(m_1, m_2)\sigma_2^k(C_2|m_2)}.
\end{aligned}$$

Finally, for all (m_1, m_2),

$$\begin{aligned}
& \frac{\mu^k(m_1, m_2)\sigma_1^k(C_1|m_1)\sigma_2^k(C_2|m_2)}{\sum_{\tilde{m}_1 = \tilde{m}_2} \mu^k(\tilde{m}_1, \tilde{m}_2)\sigma_2^k(C_2|\tilde{m}_2) + \sum_{\tilde{m}_1 = \tilde{m}_2} \mu^k(\tilde{m}_1, \tilde{m}_2)\sigma_1^k(C_1|\tilde{m}_1)} \\
\leq\ & \lim_k \frac{\mu^k(m_1, m_2)\sigma_1^k(C_1|m_1)\sigma_2^k(C_2|m_2)}{\mu^k(\tilde{m}_1 = m_2, m_2)\sigma_2^k(C_2|m_2) + \mu^k(m_1, \tilde{m}_2 = m_1)\sigma_1^k(C_1|m_1)} = 0.
\end{aligned}$$

This implies that player 3's critical belief equals 0, a contradiction.

11 Proof of Proposition 5

11.1 Proof of Claim 1

We first slightly strengthen Proposition 1.

Lemma 2 *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 $i \neq 0$, $\bigcup_{\sigma'_{-i}\in\Sigma_{-i}} \operatorname{supp} \rho_i^{\left(\sigma_i,\sigma'_{-i},\mu\right)} \supseteq \bigcup_{\tilde{\sigma}'_{-i}\in C(\Sigma_{-i})} \operatorname{supp} \rho_i^{\left(\tilde{\sigma}_i,\tilde{\sigma}'_{-i},\tilde{\mu}\right)}$.*

This says the canonical equilibrium provides the "least feedback" to the players.

Proof. We show that, for the strategy profile $(\tilde{\sigma}, \tilde{\mu})$ constructed in the proof of Proposition 1, $\bigcup_{\sigma'_{-i}\in\Sigma_{-i}} \operatorname{supp} \rho_i^{\left(\sigma_i,\sigma'_{-i},\mu\right)} \supseteq \bigcup_{\tilde{\sigma}'_{-i}\in C(\Sigma_{-i})} \operatorname{supp} \rho_i^{\left(\tilde{\sigma}_i,\tilde{\sigma}'_{-i},\tilde{\mu}\right)}$. 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" $\tilde{r}_{i,1} \in S_{i,1}$ according to $\tilde{\sigma}_i'^R(s_{i,1})$. Player i then sends report $r_{i,1} \in R_{i,1}$ according to $\sigma_i^R(\tilde{r}_{i,1})$. Next, after receiving message $m_{i,1} \in M_{i,1}$, player i draws a fictitious action recommendation $\tilde{m}_{i,1} \in A_{i,1}$ according to $\sigma_i^A(\tilde{r}_{i,1}, r_{i,1}, m_{i,1})$. Finally, player i takes action $a_{i,1} \in A_{i,1}$ according to $\tilde{\sigma}_i'^A\left(\tilde{r}_{i,1}, \tilde{r}_{i,1}^S, \tilde{m}_{i,1}\right)$.

Recursively, given $\tilde{h}_i^t = (s_{i,\tau}, \tilde{r}_{i,\tau}, \tilde{m}_{i,\tau}, a_{i,\tau})_{\tau=1}^{t-1}$, 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 $\tilde{r}_{i,t} \in A_{i,t-1} \times S_{i,t}$ according to $\tilde{\sigma}_i'^R(\tilde{h}_i^t, s_{i,t})$. Player i then sends $r_{i,t} \in R_{i,t}$ according to $\sigma_i^R((\tilde{r}_{i,\tau}, r_{i,\tau}, m_{i,\tau}), \tilde{r}_{i,t})_{\tau=1}^{t-1}$. (If $(\tilde{r}_{i,\tau})_{\tau=1}^{t}$ 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 $\tilde{m}_{i,t} \in A_{i,t}$ according to $\sigma_i^A((\tilde{r}_{i,\tau}, r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}, \tilde{r}_{i,t}, r_{i,t}, m_{i,t})$. Finally, player i takes action $a_{i,t} \in A_{i,t}$ according to $\tilde{\sigma}_i'^A(\tilde{h}_i^t, s_{i,t}, \tilde{r}_{i,t}, \tilde{m}_{i,t})$.

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

Given Lemma 2, it suffices to show that any NE outcome can be implemented in strategies where each player is sequentially rational following her own deviations.[22] When $\operatorname{supp} \rho_i^{(\sigma,\mu)} = \bigcup_{\sigma'_{-i}\in\Sigma_{-i}} \operatorname{supp} \rho_i^{\left(\sigma_i,\sigma'_{-i},\mu\right)}$, every history for player i is either on-path or follows one of player i's own deviations. Combining the lemmas, every NE outcome is implementable in a canonical NE where players are sequentially rational at all histories—that is, a canonical MSE.

Formally, by a *deviation* from σ_i we mean a report $r_{i,t}$ or action $a_{i,t}$ such that $r_{i,t} \notin \operatorname{supp} \sigma_i(\mathfrak{h}_i^t)$ or $a_{i,t} \notin \operatorname{supp} \sigma_i(\mathfrak{h}_i^t, r_{i,t}, m_{i,t})$, 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.

[22] This 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 $\left(\hat{\sigma}, \mu, \hat{\beta}\right)$ such that*

1. *$(\hat{\sigma}, \mu)$ is a NE.*
2. *For all $i \neq 0$, (2) and (3) 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 $\sigma'_{-i} \in \Sigma_{-i}$, $\rho_i^{\left(\hat{\sigma}_i, \sigma'_{-i}, \mu\right)} = \rho_i^{\left(\sigma_i, \sigma'_{-i}, \mu\right)}$.*

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 $\left(\hat{\sigma}^k, \mu, \hat{\beta}^k\right)$. Let $\left(\hat{\sigma}, \mu, \hat{\beta}\right) = \lim_k \left(\hat{\sigma}^k, \mu, \hat{\beta}^k\right)$, 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, $\hat{\sigma}_i$ 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. ■

Proof of Claim 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 $\left(\hat{\sigma}, \hat{\mu}, \hat{\beta}\right)$ 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 $i \neq 0$,

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

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

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

and hence $\operatorname{supp} \rho_i^{(\hat{\sigma},\hat{\mu})} = \bigcup_{\hat{\sigma}'_{-i} \in C(\Sigma_{-i})} \operatorname{supp} \rho_i^{\left(\hat{\sigma}_i, \hat{\sigma}'_{-i}, \hat{\mu}\right)}$. 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. ■

11.2 Proof of Claim 2

If $N = 1$, the condition $\operatorname{supp} \rho_i^{(\sigma,\mu)} = \bigcup_{\sigma'_{-i} \in \Sigma_{-i}} \operatorname{supp} \rho_i^{\left(\sigma_i, \sigma'_{-i}, \mu\right)}$ is vacuous. Hence, the result follows from Claim 1.

11.3 Proof of Claim 3

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 $\tilde{\mu}^M\left(h_0^t, s_{0,t}, r_t\right) = \mu^M\left(h_0^t, s_{0,t}, r_t\right)$. As for actions, at the beginning of the game, with probability ε the mediator selects a path of actions $(a_{0,t})_{t=0}^T \in (A_{0,t})_{t=0}^T$ 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 $\mu^A\left(h_0^t, s_{0,t}, r_t, m_t\right)$. 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, $\left\|\rho - \rho^{(\sigma,\tilde{\mu})}\right\| < \varepsilon$. Finally, since only the mediator takes actions and $S_{i,t} = \bigcup_{\mathring{h}^t \in \mathring{H}^t} \operatorname{supp} p_i\left(\cdot|\mathring{h}^t\right)$ for all i, t, $\operatorname{supp} \rho_i^{(\sigma',\tilde{\mu})} = \mathring{H}_i^{T+1}$ for all $i \neq 0$ and $\sigma' \in \Sigma$. The result now follows from Claim 1.

12 Proof of Theorem 1

12.1 Outline of the Proof

Fix a canonical CPPBE (σ, μ, β) with outcome distribution ρ. Our goal is to construct a sequence of completely mixed strategy profiles $\left(\sigma^k, \mu^k\right)$ with corresponding belief system β^k such that $\left(\sigma^k, \mu^k, \beta^k\right)$ converges to a PSE $(\sigma^*, \mu^*, \beta^*)$ with the same outcome distribution. The proof proceeds as follows:

Section 12.2 constructs $\left(\sigma^k, \mu^k\right)$. As described in the text, (i) we augment the message set with a single extra message $\star$, (ii) roughly speaking, the mediator follows an automaton strategy with two states, ρ and π, where he sends action recommendations in state ρ and sends message $\star$ in state π (in the actual construction, the mediator has both a "provisional state" and a "final state," and whether he recommends an action or sends $\star$ depends on both states), and (iii) players' limit beliefs (as $k \to \infty$) coincide with those in the original CPPBE (σ, μ, β) conditional on receiving action recommendations.

While the construction of $\left(\sigma^k, \mu^k\right)$ is completed in Section 12.2, the verification that players' limit beliefs coincide with those in (σ, μ, β) is postponed to Sections 12.3, 12.4, and 12.6. In particular, Section 12.6 is the longest and most technical part of the proof, but it contains few conceptual novelties and essentially amounts to repeated, careful applications of Bayes' rule.

Once the desired belief properties have been established, verifying sequential rationality of the limit assessment $(\sigma^*, \mu^*, \beta^*)$ is straightforward. This is accomplished in Section 12.5.

12.2 Construction of $\left(\sigma^k, \mu^k\right)$

12.2.1 Message Set

The message set used in the construction is given by $R_{i,t} = A_{i,t-1} \times S_{i,t}$ and $M_{i,t} = A_{i,t} \cup \{\star\}$ 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. We write $r_{i,t} = (r_{i,t-1}^A, r_{i,t}^S) \in A_{i,t-1} \times S_{i,t}$.

12.2.2 Auxiliary Recommendations, Histories, Faithfulness

It will be useful to introduce an auxiliary action recommendation $\tilde{m}_{i,t} \in A_{i,t}$ for each i, t, where $\tilde{m}_{i,t} = m_{i,t}$ whenever $m_{i,t} \in A_{i,t}$. That is, if the mediator recommends an action $m_{i,t} \in A_{i,t}$ then $\tilde{m}_{i,t} = m_{i,t}$, while if the mediator sends message $\star$ then $\tilde{m}_{i,t}$ is some action in $A_{i,t}$, which may be interpreted as the message the mediator "would have sent" had he not trembled to $\star$.

A history for the mediator (including the auxiliary recommendations $\tilde{m}_{i,t}$) is denoted $\mathfrak{h}_0^t = \left((s_{0,\tau}, r_\tau, \tilde{m}_\tau, m_\tau, a_{0,\tau})_{\tau=1}^{t-1}, s_{0,t}, r_t\right)$. Given $\mathfrak{h}_0^t$, let $c(\mathfrak{h}_0^t) = ((s_{0,\tau}, r_\tau, \tilde{m}_\tau, a_{0,\tau})_{\tau=1}^{t-1}, s_{0,t}, r_t)$ be the corresponding canonical history for the mediator assuming the mediator has not trembled to $\star$ and all players have reported truthfully: $r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$ for all $i, \tau \leq t$. Let $\dot{\mathfrak{h}}_0^t$ denote a generic canonical history for the mediator. In addition, given $\dot{\mathfrak{h}}_0^t$, let $\mathring{c}(\dot{\mathfrak{h}}_0^t) = ((s_\tau, a_\tau)_{\tau=1}^{t-1}, s_t)$ with $(a_{i,\tau-1}, s_{i,\tau}) = r_{i,\tau}$ for all $i, \tau \leq t$ be the corresponding payoff-relevant history assuming players are truthful.

Given a truthful history $\mathfrak{h}_i^t = \left((s_{i,\tau}, m_{i,\tau}, a_{i,\tau})_{\tau=1}^{t-1}, s_{i,t}\right)$ for player i and a sequence of auxiliary recommendations $(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}$, let $c(\mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}) = ((s_{i,\tau}, \tilde{m}_{i,\tau}, a_{i,\tau})_{\tau=1}^{t-1}, s_{i,t})$ be the canonical history for player i in which each $m_{i,\tau}$ is replaced by $\tilde{m}_{i,\tau}$. Let $\dot{\mathfrak{h}}_i^t$ denote a generic canonical history for player i. To reduce notation, we omit the reports $(r_{i,\tau})_{\tau=1}^t$ when the history $\mathfrak{h}_i^t$ is assumed truthful.

Finally, it will be useful to generalize the usual notion of "truthfulness and obedience"—players report truthfully and obey action recommendations—to the case where the mediator may send message $\star$, in which case there is no action recommendation to obey. We call this notion "faithfulness." Formally, player i is *unfaithful* in period t if she sends a report $r_{i,t} \neq (a_{i,t-1}, s_{i,t})$, or if $m_{i,t} \neq \star$ but she takes an action $a_{i,t} \neq m_{i,t}$. Otherwise, player i is *faithful* in period t. A history $\mathfrak{h}_i^t$ is *faithful* if player i is faithful in every period $\tau \leq t-1$. Let $\mathcal{F} \subset H^{T+1}$ denote the event that each player is faithful in every period. Since faithfulness implies truthfulness, we omit the reports $(r_{i,\tau})_{\tau=1}^t$ when describing faithful histories.

12.2.3 Canonical Histories and the Canonical CPPBE

It will also be useful to note that, given a canonical history $\dot{\mathfrak{h}}_0^t = ((s_{0,\tau}, r_\tau, \tilde{m}_\tau, a_{0,\tau})_{\tau=1}^{t-1}, s_{0,t}, r_t)$, the corresponding distribution of messages and actions conditional on reaching $\dot{\mathfrak{h}}_0^t$ in the original CPPBE (σ, μ, β) are well-defined. In particular, let f be the CPS on H^{T+1} corresponding to (σ, μ, β). By Proposition 3, there exists a sequence of recommendation distributions $\left(\alpha_t^{M,k}\right)_t$ (where $\alpha_t^{M,k} : H^t \times S_t \times R_t \to \Delta(A_t)$) and a sequence of completely mixed action distributions $(\alpha_t^{A,k})_t$ (where $\alpha_t^{A,k} : H^t \times S_t \times R_t \times M_t \to \Delta(A_t)$) such that, letting α^k be the move distribution that results when the mediator recommends actions according to $\left(\alpha_t^{M,k}\right)_t$ and players report truthfully and take actions according to $(\alpha_t^{A,k})_t$, we have $F\left(\alpha^k\right) \to f$ on all events where players have reported truthfully (where $F\left(\alpha^k\right)$ is the probability distribution on H^{T+1} induced by α^k). The distribution of messages and actions conditional on reaching $\dot{\mathfrak{h}}_0^t$ in (σ, μ, β) are thus given by $F(\alpha^k)(\tilde{m}_t | \dot{\mathfrak{h}}_0^t) = F(\alpha^k)(\tilde{m}_t | (\mathfrak{h}_i^t)_{i=0}^N, r_t)$ and $F(\alpha^k)(a_t | \dot{\mathfrak{h}}_0^t, \tilde{m}_t) = F(\alpha^k)(a_t | (\mathfrak{h}_i^t)_{i=0}^N, r_t, \tilde{m}_t)$, where

$$\mathfrak{h}_i^t = ((s_{i,\tau}, r_{i,\tau}, \tilde{m}_{i,\tau}, a_{i,\tau})_{\tau=1}^{t-1}, s_{i,t}) \text{ with } (a_{i,\tau-1}, s_{i,\tau}) = r_{i,\tau} \ \forall i, \tau \leq t.$$

12.2.4 Overview of the Construction

For each k, we construct $\left(\sigma^k, \mu^k\right)$ by a fixed point argument. After first introducing some notation for tremble sequences, we construct a strategy profile $\bar{\sigma} = (\bar{\sigma}_i)_{i=1}^N$ for the players. We then explicitly perturb $\bar{\sigma}$ to obtain a strategy profile $\bar{\sigma}^k = \left(\bar{\sigma}_i^k\right)_{i=1}^N$ for each k. Given $\bar{\sigma}$ and $\bar{\sigma}^k$, we construct a strategy for the mediator $\bar{\mu}^k$ such that, if the players follows $\bar{\sigma}^k$, then $\bar{\mu}^k$ "replicates" the CPS $F\left(\alpha^k\right)$. We then define $FBR_i^k(\bar{\sigma})$ to be the set of player i's "faithful best responses" (roughly speaking, best responses subject to faithfulness) to opposing strategies $\left(\bar{\sigma}_j^k\right)_{j \neq i}$ and mediator's strategy $\bar{\mu}^k$. Finally, we find a fixed point $\bar{\sigma} \in \prod_{i=1}^N FBR_i^k(\bar{\sigma})$, and we define $\left(\sigma^k, \mu^k\right)$ as the strategies $\left(\bar{\sigma}^k, \bar{\mu}^k\right)$ corresponding to this fixed point.

12.2.5 Tremble Sequences

The construction 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 a PSE. So assume ρ does not have full support, and let

$$\begin{aligned} \underline{f}_3(k) &= \min_{E',E'' \subseteq E \subseteq H^{T+1} : \lim_k \frac{F(\alpha^k)(E'|E)}{F(\alpha^k)(E''|E)} = 0} \frac{F\left(\alpha^k\right)(E'|E)}{F\left(\alpha^k\right)(E''|E)} \\ \bar{f}_3(k) &= \max_{E',E'' \subseteq E \subseteq H^{T+1} : \lim_k \frac{F(\alpha^k)(E'|E)}{F(\alpha^k)(E''|E)} = 0} \frac{F\left(\alpha^k\right)(E'|E)}{F\left(\alpha^k\right)(E''|E)} \end{aligned} \tag{8}$$

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 $k \to \infty$ and satisfy the relationship

$$f_1(k) \gg f_2(k) \gg \bar{f}_3(k) \geq \underline{f}_3(k) \gg f_4(k), \tag{9}$$

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

12.2.6 Construction of $\bar{\sigma}$ and $\bar{\sigma}^k$

For each player i, let $\bar{\sigma}_i$ be an arbitrary strategy under which player i is faithful if she has been faithful in the past: that is, a strategy such that (i) if $\mathfrak{h}_i^t$ is faithful, then $\bar{\sigma}_i^R\left(\mathfrak{h}_i^t\right) = (a_{i,t-1}, s_{i,t})$, and (ii) if $\mathfrak{h}_i^t$ is faithful, $r_{i,t} = (a_{i,t-1}, s_{i,t})$, and $m_{i,t} \neq \star$, then $\bar{\sigma}_i^A\left(\left(\mathfrak{h}_i^t, r_{i,t}, m_{i,t}\right)\right) = m_{i,t}$. In particular, player i's action after $m_{i,t} = \star$ is arbitrary.

Given $\bar{\sigma}_i$, let $\bar{\sigma}_i^k$ denote the following full-support perturbation of $\bar{\sigma}_i$:

1. For each faithful history $\mathfrak{h}_i^t$, player i reports $r_{i,t} = (a_{i,t-1}, s_{i,t})$ with probability $1 - f_4(k)(|R_{i,t}| - 1)$ and sends every other report in $R_{i,t}$ with probability $f_4(k)$.

2. For each faithful history $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t})$ with $m_{i,t} \neq \star$, player i plays $a_{i,t} = m_{i,t}$ with probability $1 - f_4(k)(|A_{i,t}| - 1)$ and plays every other action in $A_{i,t}$ with probability $f_4(k)$.

3. For each faithful history $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t})$ with $m_{i,t} = \star$, player i first draws $b_{i,t}$ from distribution $\bar{\sigma}_i^A(\mathfrak{h}_i^t, r_{i,t}, m_{i,t})$, and then plays $b_{i,t}$ with probability $1 - f_1(k)(|A_{i,t}| - 1)$ and plays every other action in $A_{i,t}$ with probability $f_1(k)$.

4. At any unfaithful history, player i follows $\bar{\sigma}_i$ with probability $1 - f_4(k)$ and randomizes uniformly over her available reports or actions with probability $f_4(k)$.

In particular, $\bar{\sigma}_i^k$ deviates from $\bar{\sigma}_i$ with much higher probability when $m_{i,t} = \star$ than when $m_{i,t} \in A_{i,t}$.

12.2.7 Construction of $\bar{\mu}^k$

We construct an automaton representation of the mediator's strategy $\bar{\mu}^k$. In each period t, the mediator has a "provisional state" $\hat{\theta}_t \in \{\rho, \pi\}$ and a "final state" $\theta_t \in \{\rho, \pi\}$. Intuitively, being in state ρ means the mediator is behaving "as usual," while being in state π means the mediator has switched to sending message $\star$. The mediator begins period t in the provisional state $\hat{\theta}_t$ but may transition to final state $\theta_t \neq \hat{\theta}_t$ as the result of a tremble. We first describe the mediator's play as a function of $\hat{\theta}_t$ and θ_t and then describe how $\hat{\theta}_t$ and θ_t are determined.

Given $\left(\hat{\theta}_t, \theta_t\right) \in \{(\rho,\rho), (\rho,\pi), (\pi,\pi)\}$ and $\mathfrak{h}_0^t$, the mediator plays as follows (as will become clear, it is not possible that $\left(\hat{\theta}_t, \theta_t\right) = (\pi, \rho)$): The mediator first draws $\tilde{m}_t \in A_t$ from distribution $F\left(\alpha^k\right)(\cdot|c(\mathfrak{h}_0^t))$:

$$\tilde{m}_t \sim F\left(\alpha^k\right)(\cdot|c(\mathfrak{h}_0^t)) \quad \forall \mathfrak{h}_0^t. \tag{10}$$

Then,

1. If $\hat{\theta}_t = \theta_t = \rho$, the mediator sends message $m_{i,t} = \tilde{m}_{i,t}$ for all i.

2. If $\hat{\theta}_t = \rho$ and $\theta_t = \pi$, the mediator sends message

$$m_{i,t} = \begin{cases} \star & \text{with probability } 1 - f_2(k) \\ \tilde{m}_{i,t} & \text{with probability } f_2(k) \end{cases}, \tag{11}$$

independently across players i.

3. If $\hat{\theta}_t = \theta_t = \pi$, the mediator sends $m_{i,t} = \star$ for all i.

The provisional state in period 1 is ρ: $\hat{\theta}_1 = \rho$. For $t \geq 2$, the provisional state $\hat{\theta}_t$ is determined recursively below. Given $\hat{\theta}_t$, the period t final state is determined as follows: If

$\hat{\theta}_t = \rho$, then $\theta_t = \rho$ with probability $1 - f_2(k)$ and $\theta_t = \pi$ with probability $f_2(k)$—that is,

$$\Pr\left(\theta_t = \rho | \hat{\theta}_t = \rho, \mathfrak{h}_0^t\right) = 1 - f_2(k), \Pr\left(\theta_t = \pi | \hat{\theta}_t = \rho, \mathfrak{h}_0^t\right) = f_2(k) \quad \forall \mathfrak{h}_0^t. \tag{12}$$

If $\hat{\theta}_t = \pi$, then $\theta_t = \pi$ with probability 1.

Finally, given $\bar{\sigma} = (\bar{\sigma}_i)_{i=1}^N$, $\hat{\theta}_t$, θ_t, and $\mathfrak{h}_0^{t+1}$, the period $t+1$ provisional state $\hat{\theta}_{t+1}$ is determined as follows (this is the only part of the mediator's strategy that depends on $\bar{\sigma}$):

1. If $\theta_t = \rho$, then $\hat{\theta}_{t+1} = \rho$.
2. If $\theta_t = \pi$, then
   (a) If $\hat{\theta}_t = \rho$, then $\hat{\theta}_{t+1} = \rho$ with probability $\varepsilon^k\left(\mathfrak{h}_0^{t+1}\right)$ and $\hat{\theta}_{t+1} = \pi$ with probability $1 - \varepsilon^k\left(\mathfrak{h}_0^{t+1}\right)$, where the transition probability $\varepsilon^k\left(\mathfrak{h}_0^{t+1}\right)$ is determined below (and depends on $\bar{\sigma}$).
   (b) If $\hat{\theta}_t = \pi$, then $\hat{\theta}_{t+1} = \pi$.

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

$$\begin{array}{ccccc}
\hat{\theta}_1 = \rho & & & & \\
\downarrow_{1-f_2(k)} & \searrow_{f_2(k)} & & & \\
\theta_1 = \rho & & \theta_1 = \pi & & \\
\downarrow_1 & \swarrow_{\varepsilon^k\left(\mathfrak{h}_0^2\right)} & & \searrow_{1-\varepsilon^k\left(\mathfrak{h}_0^2\right)} & \\
\hat{\theta}_2 = \rho & & & & \hat{\theta}_2 = \pi \\
\downarrow_{1-f_2(k)} & \searrow_{f_2(k)} & & & \downarrow_1 \\
\theta_2 = \rho & & \theta_2 = \pi & & \theta_2 = \pi \\
\downarrow_1 & \swarrow_{\varepsilon^k\left(\mathfrak{h}_0^3\right)} & & \searrow_{1-\varepsilon^k\left(\mathfrak{h}_0^3\right)} & \downarrow_1 \\
\hat{\theta}_3 = \rho & & & & \hat{\theta}_3 = \pi \\
\downarrow_{1-f_2(k)} & \searrow_{f_2(k)} & & & \downarrow_1 \\
\theta_3 = \rho & & \theta_3 = \pi & & \theta_3 = \pi \\
\vdots & & \vdots & & \vdots
\end{array}$$

(It may be helpful to relate the mediator's strategy to that in Example 2. In the context of the example, $\theta_1 = \pi$ and $m_{i,1} = \star$ correspond to sending "free passes," and $\varepsilon^k\left(\mathfrak{h}_0^2\right)$ is the probability that the mediator trembles a second time and punishes the players after sending free passes. We will define $\varepsilon^k\left(\mathfrak{h}_0^2\right)$ so that, conditional on the event $\hat{\theta}_2 = \rho$ and $m_{3,2} = P$, player 3's beliefs match those in the original CPPBE.)

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

To complete the description of $\bar{\mu}^k$ given $\bar{\sigma}$, it remains to specify the state transition probability $\varepsilon^k\left(\mathfrak{h}_0^t\right)$ for each history $\mathfrak{h}_0^t$. The goal is to define these variables so that, conditional on the event $\hat{\theta}_t = \rho$, a faithful player's beliefs about $\dot{\mathfrak{h}}_0^t$ are close to her beliefs under the original move distribution α^k. More precisely, we wish to ensure that, conditional on $\hat{\theta}_t = \rho$, on $(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}$, on reaching faithful history $\mathfrak{h}_i^t$, and on other players being faithful, player i's belief about $\dot{\mathfrak{h}}_0^t$ is asymptotically (as $k \to \infty$) equal to $F\left(\alpha^k\right)(\dot{\mathfrak{h}}_0^t | c(\mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}))$: that is,

$$\lim_k \frac{\Pr^{\left(\bar{\sigma}^k, \bar{\mu}^k\right)}\left(\dot{\mathfrak{h}}_0^t | \mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho\right)}{F\left(\alpha^k\right)\left(\dot{\mathfrak{h}}_0^t | c(\mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1})\right)} = 1. \tag{13}$$

The current section constructs $\varepsilon^k\left(\mathfrak{h}_0^t\right)$. Sections 12.3, 12.4, and 12.6 verify the desired convergence of beliefs.

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

Given $\mathfrak{h}_0^{t+1}$ and $\tilde{m}_t$, we define $\varepsilon^k\left(\mathfrak{h}_0^{t+1}\right)$ as follows: Let $I = \{1, \ldots, N\}$ and let $I_t^* = \{i \in I : m_{i,t} \neq \star\}$. Note that the message m_t is completely determined by $\tilde{m}_t$ and I_t^*.

If $I_t^* = I$ then $\varepsilon^k\left(\mathfrak{h}_0^{t+1}\right) := 0$. In particular, this implies that $I_t^* = I$ and $\hat{\theta}_{t+1} = \rho$ if and only if $\theta_t = \rho$ (given $\mathcal{F}$): for each $(\mathfrak{h}_0^t, \tilde{m}_t)$,

$$\begin{aligned} & \Pr^{\left(\bar{\sigma}^k, \bar{\mu}^k\right)}\left(I_t^* = I, a_t = \tilde{m}_t, \hat{\theta}_{t+1} = \rho | \mathcal{F}, \mathfrak{h}_0^t, \tilde{m}_t, \hat{\theta}_t = \rho\right) \\ = \; & \Pr^{\left(\bar{\sigma}^k, \bar{\mu}^k\right)}\left(\theta_t = \rho | \mathcal{F}, \mathfrak{h}_0^t, \tilde{m}_t, \hat{\theta}_t = \rho\right) = 1 - f_2(k). \end{aligned} \tag{14}$$

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

$$\Pr^{\left(\bar{\sigma}^k, \bar{\mu}^k\right)}\left(I_t^*, a_t, \hat{\theta}_{t+1} = \rho | \mathfrak{h}_0^t, \tilde{m}_t, \hat{\theta}_t = \rho\right) = F\left(\alpha^k\right)\left(a_t | c(\mathfrak{h}_0^t), \tilde{m}_t\right). \tag{15}$$

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

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

In addition, since $a_t \neq \tilde{m}_t$, (8) implies

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

We can therefore define

$$\begin{aligned} \varepsilon^k\left(\mathfrak{h}_0^{t+1}\right) &= \frac{F\left(\alpha^k\right)\left(a_t | c(\mathfrak{h}_0^t), \tilde{m}_t\right)}{\Pr^{\left(\bar{\sigma}^k, \bar{\mu}^k\right)}\left(I_t^*, a_t | \mathfrak{h}_0^t, \tilde{m}_t, \hat{\theta}_t = \rho\right)} \\ &\leq \frac{\bar{f}_3(k)}{\left(f_2(k)\right)^{N+1}\left(f_1(k)\right)^{N+1}} \to 0 \text{ by (9).} \end{aligned} \tag{16}$$

Note that (15) obtains with this definition.

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

12.2.9 Mediation Range

We now define the mediation range for player i, $Q_i\left(\left(r_{i,\tau}, m_{i,\tau}\right)_{\tau=1}^{t-1}, r_{i,t}\right)$. Let $\bar{Q}_i\left(\left(r_{i,\tau}, m_{i,\tau}\right)_{\tau=1}^{t-1}, r_{i,t}\right)$ be the mediation range in the original CPPBE (σ, μ, β). Intuitively, $Q_i\left(\left(r_{i,\tau}, m_{i,\tau}\right)_{\tau=1}^{t-1}, r_{i,t}\right)$ is defined as

$$\bigcup_{\left(\tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}} \bar{Q}_i\left(\left(r_{i,\tau}, \tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}, r_{i,t}\right) \cup \{\star\},$$

where the union is taken over auxiliary recommendations $\left(\tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}$ that arise with positive probability given communication history $\left(r_{i,\tau}, m_{i,\tau}\right)_{\tau=1}^{t-1}$. The formal construction follows—this may be skipped on a first reading without interfering with the rest of the proof.

The mediation range in period 1 is $Q_i\left(r_{i,1}\right) = \bar{Q}_i(r_{i,1}) \cup \{\star\}$. In period t, given $Q_i\left(\left(r_{i,t'}, m_{i,t'}\right)_{t'=1}^{\tau-1}, r_{i,\tau}\right)$ for each $\tau \leq t-1$, let $R\tilde{M}_i((r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1})$ be the set of pairs $\left(r_{i,\tau}, \tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}$ such that, for each $\tau \leq t-1$, (i) if $m_{i,\tau} \in A_{i,\tau}$ then $\tilde{m}_{i,\tau} = m_{i,\tau}$; and (ii) if $m_{i,\tau} = \star$ then $\tilde{m}_{i,\tau} \in Q_i\left(\left(r_{i,t'}, m_{i,t'}\right)_{t'=1}^{\tau-1}, r_{i,\tau}\right) \cap A_{i,\tau} \setminus \{r_{i,\tau}^A\}$. If $m_{i,t-1} \in A_{i,t-1}$ then the mediation range is

$$Q_i\left(\left(r_{i,\tau}, m_{i,\tau}\right)_{\tau=1}^{t-1}, r_{i,t}\right) := \bigcup_{\left(r_{i,\tau}, \tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1} \in R\tilde{M}\left(\left(r_{i,\tau}, m_{i,\tau}\right)_{\tau=1}^{t-1}\right)} \bar{Q}_i\left(\left(r_{i,\tau}, \tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}, r_{i,t}\right) \cup \{\star\}. \quad (17)$$

If $m_{i,t-1} = \star$ and the previous mediation range $Q_i\left(\left(r_{i,\tau}, m_{i,\tau}\right)_{\tau=1}^{t-2}, r_{i,t-1}\right)$ is singleton then $Q_i\left(\left(r_{i,\tau}, m_{i,\tau}\right)_{\tau=1}^{t-1}, r_{i,t}\right) := \{\star\}$. Otherwise, the mediation range is (17).

We must check that, for any $\bar{\sigma}$ and any $(\mathfrak{h}_i^{t-1}, r_{i,t-1}, m_{i,t-1}, r_{i,t})$, we have

$$\Pr^{\left(\bar{\sigma}_{-i}^k, \bar{\mu}^k(\bar{\sigma})\right)}(m_{i,t} | \mathfrak{h}_i^{t-1}, r_{i,t-1}, m_{i,t-1}, r_{i,t}) > 0$$

if and only if $m_{i,t} \in Q_i\left(\left(r_{i,\tau}, m_{i,\tau}\right)_{\tau=1}^{t-2}, r_{i,t-1}\right)$. The proof is by induction on t.

Suppose first that $m_{i,t-1} \in A_{i,t-1}$. Then, for any $(a_{i,t-1}, s_{i,t})$,

$$\Pr^{\left(\bar{\sigma}_{-i}^k, \bar{\mu}^k(\bar{\sigma})\right)}(\hat{\theta}_{t-1} = \theta_{t-1} = \rho | \mathfrak{h}_i^{t-1}, r_{i,t-1}, m_{i,t-1}, a_{i,t-1}, s_{i,t}) > 0,$$

since players tremble to all actions with positive probability. Hence, for each $r_{i,t}$,

$$\Pr^{\left(\bar{\sigma}_{-i}^k, \bar{\mu}^k(\bar{\sigma})\right)}(\hat{\theta}_t = \rho | \mathfrak{h}_i^{t-1}, r_{i,t-1}, m_{i,t-1}, a_{i,t-1}, s_{i,t}, r_{i,t}) > 0, \quad (18)$$

and so, by the transition rule for θ_t,

$$\min\left\{ \begin{array}{l} \Pr^{\left(\bar{\sigma}_{-i}^k, \bar{\mu}^k(\bar{\sigma})\right)}(\theta_t = \rho | \mathfrak{h}_i^{t-1}, r_{i,t-1}, m_{i,t-1}, a_{i,t-1}, s_{i,t}, r_{i,t}), \\ \Pr^{\left(\bar{\sigma}_{-i}^k, \bar{\mu}^k(\bar{\sigma})\right)}(\theta_t = \pi | \mathfrak{h}_i^{t-1}, r_{i,t-1}, m_{i,t-1}, a_{i,t-1}, s_{i,t}, r_{i,t}) \end{array} \right\} > 0.$$

Given $\theta_t = \rho$, since player i does not observe $\tilde{m}_{i,\tau}$ when $m_{i,\tau} = \star$ and $m_{i,\tau} = \star \cap r_{i,\tau}^A = \tilde{m}_{i,\tau}$

implies $m_{i,\tau'} = \star$ for all $\tau' \geq \tau + 1$, we have

$$\Pr^{\left(\bar{\sigma}^k_{-i}, \bar{\mu}^k(\bar{\sigma})\right)}\left(\left(r_{i,\tau}, \tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1} | \mathfrak{h}_i^{t-1}, r_{i,t-1}, m_{i,t-1}, a_{i,t-1}, s_{i,t}, r_{i,t}\right) > 0$$

if and only if $\left(r_{i,\tau}, \tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1} \in \tilde{RM}(\left(r_{i,\tau}, m_{i,\tau}\right)_{\tau=1}^{t-1})$ by the inductive hypothesis. Hence, any $m_{i,t}$ in

$$\bigcup_{\left(r_{i,\tau}, \tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1} \in \tilde{RM}\left(\left(r_{i,\tau}, m_{i,\tau}\right)_{\tau=1}^{t-1}\right)} \bar{Q}_i\left(\left(r_{i,\tau}, \tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}, r_{i,t}\right)$$

is sent with positive probability. Given $\theta_t = \pi$, the mediator sends $\star$ with positive probability. Hence, the mediation range is (17).

Suppose instead $m_{i,t-1} = \star$. If

$$Q_i\left(\left(r_{i,t'}, m_{i,t'}\right)_{\tau=1}^{t-2}, r_{i,t-1}\right) \cap A_{i,t-1} \setminus \{r_{i,t-1}^A\} = \emptyset$$

then $\varepsilon^k\left(\mathfrak{h}_0^t\right) = 0$ by the inductive hypothesis and the definition of $\varepsilon^k(\mathfrak{h}_0^t)$. Hence, $m_{i,t} = \star$ with probability 1. Inspecting the definition of $\tilde{RM}(\left(r_{i,\tau}, m_{i,\tau}\right)_{\tau=1}^{t-1})$, this implies $\tilde{RM}(\left(r_{i,\tau}, m_{i,\tau}\right)_{\tau=1}^{t-1}) = \emptyset$. Therefore, the mediation range is (17). Otherwise,

$$\Pr^{\left(\bar{\sigma}^k_{-i}, \bar{\mu}^k(\bar{\sigma})\right)}\left(\begin{array}{c} \left\{\tilde{m}_{i,t-1} \neq r_{i,t-1}^A\right\} \cap \left\{m_{j,t-1} = \tilde{m}_{j,t-1} = r_{j,t-1}^A \ \forall j \neq i\right\} \\ |\mathfrak{h}_i^{t-1}, r_{i,t-1}, m_{i,t-1}, a_{i,t-1}, s_{i,t}, r_{i,t} \end{array}\right) > 0.$$

Hence, (18) holds. So again the mediation range is (17).

12.2.10 Fixed Point Argument

Let $\bar{\mu}^k\left(\bar{\sigma}\right)$ denote the mediator's strategy $\bar{\mu}^k$ as a function of $\bar{\sigma}$ as constructed in Sections 12.2.7 and 12.2.8. Define $FBR_i^k(\bar{\sigma}) \subset \Sigma_i$ to be the set of strategies $\sigma_i \in \Sigma_i$ such that:

1. For each faithful history $\mathfrak{h}_i^t$, $\sigma_i^R\left(\mathfrak{h}_i^t\right) = r_{i,t}$ with $r_{i,t} = (a_{i,t-1}, s_{i,t})$.

2. For each faithful history $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t})$ with $m_{i,t} \neq \star$, $\bar{\sigma}_i^A\left((\mathfrak{h}_i^t, r_{i,}, m_{i,t})\right) = m_{i,t}$.

3. For each faithful history $(\mathfrak{h}_i^t, r_{i,t}, m_{i,t})$ with $m_{i,t} = \star$, and for each unfaithful history, player i plays an arbitrary best response to the belief that $(\mathfrak{h}_0^t, \mathfrak{h}_{-i}^t)$ is distributed according to $\Pr^{\left(\bar{\sigma}^k_{-i}, \bar{\mu}^k(\bar{\sigma})\right)}(\mathfrak{h}_0^t, \mathfrak{h}_{-i}^t | \mathfrak{h}_i^t)$ and, in all future periods $\tau \geq t+1$, $(\sigma_\tau, \mu_\tau) = \left(\bar{\sigma}_\tau, \bar{\mu}_\tau^k\left(\bar{\sigma}\right)\right)$.

(Intuitively, $FBR_i^k(\bar{\sigma})$ is the set of strategies for player i that are faithful after faithful histories, and are optimal against $\left(\bar{\sigma}, \bar{\mu}^k\left(\bar{\sigma}\right)\right)$ after $m_{i,t} = \star$ and at unfaithful histories.)

Let $FBR^k\left(\bar{\sigma}\right) = \prod_{i=1}^N FBR_i^k\left(\bar{\sigma}\right)$. For each k, the mapping $FBR^k : \Sigma \to \Sigma$ is non-empty, convex, and upper-hemicontinuous. So a fixed point exists by Kakutani's theorem. Let σ^k be a fixed point, and let $\mu^k = \bar{\mu}^k\left(\sigma^k\right)$. Finally, let $(\sigma^*, \mu^*) = \lim_k \left(\sigma^k, \mu^k\right)$, taking a subsequence if needed.

12.3 Belief Verification 1: Faithfulness

We prove that a faithful player always believes her opponents are faithful. In what follows, a sequence of auxiliary recommendations $(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}$ is said to be *compatible* with a faithful history $\mathfrak{h}_i^t = \left((s_{i,\tau}, m_{i,\tau}, a_{i,\tau})_{\tau=1}^{t-1}, s_{i,t}\right)$ if, for each $\tau \leq t-1$, (i) if $m_{i,\tau} \neq \star$ then $\tilde{m}_{i,\tau} = m_{i,\tau}$, and (ii) if $m_{i,\tau} = \star$ and there exists $\tau' \geq \tau + 1$ with $m_{i,\tau'} \in A_{i,\tau'}$, then $\tilde{m}_{i,\tau} \neq a_{i,\tau}$.

(Intuitively, if $(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}$ is not compatible with $\mathfrak{h}_i^t$, then this provisional recommendation sequence is precluded by the construction of the transition probabilities ε^k.)

Lemma 4 *For any faithful history* $\mathfrak{h}_i^t = \left((s_{i,\tau}, m_{i,\tau}, a_{i,\tau})_{\tau=1}^{t-1}, s_{i,t}\right)$ *and any compatible sequence of auxiliary recommendations* $(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}$*, we have*

$$\lim_k \Pr^{\left(\sigma^k, \mu^k\right)}\left(\mathcal{F} | \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}\right) = 1. \tag{19}$$

Proof. Since unfaithful histories for players $-i$ arise with probability at most $f_4(k)$, it suffices to construct a history for players $-i$ that arises without probability-$f_4(k)$ events given $\mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}$. We show this is possible by letting $\hat{\theta}_\tau = \rho$ and $\theta_\tau = \pi$ as long as this is feasible given $\mathfrak{h}_i^t$.

Given $\mathfrak{h}_i^t$, fix the largest $t^* \leq t$ such that, for all $\tau \leq t^*$, if $m_{i,\tau} = \star$ and $m_{i,\tau'} \in A_{i,\tau'}$ for some $\tau' \in \{\tau+1, ..., t-1\}$, then $\tilde{m}_{i,\tau} \neq a_{i,\tau}$. Note that, by definition of the mediation range, $m_{i,\tau} = \star$ for all $\tau \geq t^* + 1$.

Fix $(\bar{a}_{-i,\tau-1}, \bar{s}_{-i,\tau})_{\tau=1}^{t}$ (including $(\bar{a}_{0,\tau-1}, \bar{s}_{0,\tau})_{\tau=1}^{t}$) such that

$$p(s_{i,\tau}, \bar{s}_{-i,\tau} | (a_{i,t'-1}, \bar{a}_{-i,t'-1}, s_{i,t'}, \bar{s}_{-i,t'})_{t'=1}^{\tau-1}, a_{i,\tau-1}, \bar{a}_{-i,\tau-1}) > 0 \quad \forall \tau \leq t.$$

Let $r_{j,\tau} = (\bar{a}_{j,\tau-1}, \bar{s}_{j,\tau})$ for all $j \neq i, \tau \leq t$.

Next, let $\tilde{m}_{-i,1} \in A_{-i,1}$ satisfy

$$\Pr^{\left(\sigma^k, \mu^k\right)}(\tilde{m}_{-i,1} | r_1, \tilde{m}_{i,1}) \geq \underline{f}_3(k).$$

Note that there exists at least one such $\tilde{m}_{-i,1}$ by (10). Recursively, given $(\tilde{m}_\tau)_{\tau=1}^{t'-1}$ and $(r_\tau)_{\tau=1}^{t'}$, let $\tilde{m}_{-i,t'} \in A_{-i,t'}$ satisfy

$$\Pr^{\left(\sigma^k, \mu^k\right)}\left(\tilde{m}_{-i,t'} | (\tilde{m}_\tau)_{\tau=1}^{t'-1}, (r_\tau)_{\tau=1}^{t'}\right) \geq \underline{f}_3(k).$$

Finally, for each $j \neq i$ and $t' \leq t^*$, let $m_{j,t'} = \tilde{m}_{j,t'}$ if $\tilde{m}_{j,t'} = \bar{a}_{j,t'}$, and let $m_{j,t'} = \star$ otherwise.

Now, note that if $\hat{\theta}_\tau = \rho$ then any recommendation in the mediation range is sent with probability at least $\underline{f}_3(k)$. In addition, $\varepsilon^k\left(\mathfrak{h}_0^{\tau+1}\right) \geq \underline{f}_3(k)$ for every $\mathfrak{h}_0^{\tau+1}$ such that $\varepsilon^k\left(\mathfrak{h}_0^{\tau+1}\right) > 0$. Therefore, the event that $\mathfrak{h}_i^t$ and $(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}$ arise, $\hat{\theta}_\tau = \rho$ for all $\tau \leq t^*$, and $\left(\mathfrak{h}_j^t\right)_{j \neq i}$ is faithful and satisfies $(a_{-i,\tau-1}, s_{-i,\tau}) = (\bar{a}_{-i,\tau-1}, \bar{s}_{-i,\tau})$ for all $\tau \leq t^*$ occurs with probability at least $\left(f_1(k) f_2(k) \underline{f}_3(k)\right)^{(N+1)T}$. This asymptotically dominates $f_4(k)$, completing the proof. ■

12.4 Belief Verification 2: Convergence to CPPBE Beliefs

The following lemma verifies the key belief convergence properties. Equation (13) is established in the course of its proof.

Lemma 5 *For any faithful history $\mathfrak{h}_i^t$ and any compatible sequence of auxiliary recommendations $(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}$, the following hold:*

1. *For any $m_{i,t} = \tilde{m}_{i,t} \in A_{i,t}$, if player i receives $m_{i,t}$, then she believes that $\theta_t = \rho$ with probability 1, and her belief about $\tilde{m}_{-i,t}$ is the same as in the original equilibrium: for any $\dot{\mathfrak{h}}_0^t$ and $\tilde{m}_{-i,t}$,*

$$\begin{aligned} & \lim_k \Pr^{\left(\sigma^k,\mu^k\right)}\left(\dot{\mathfrak{h}}_0^t, m_{-i,t} = \tilde{m}_{-i,t}, \theta_t = \rho | \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, m_{i,t} = \tilde{m}_{i,t}\right) \\ = \; & \lim_k F\left(\alpha^k\right)\left(\dot{\mathfrak{h}}_0^t, \tilde{m}_{-i,t} | c\left(\mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}\right), \tilde{m}_{i,t}\right). \end{aligned} \tag{20}$$

2. *For any $m_{i,t} = \tilde{m}_{i,t} \in A_{i,t}$, $a_{i,t} = \tilde{m}_{i,t}$, and $s_{i,t+1}$, if player i receives $m_{i,t}$, plays $a_{i,t}$, and observes $s_{i,t+1}$, then her belief about θ_t is degenerate:*

$$\lim_k \Pr^{\left(\sigma^k,\mu^k\right)}\left(\theta_t = \rho | \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, m_{i,t} = \tilde{m}_{i,t}, a_{i,t}, s_{i,t+1}\right) \in \{0,1\}. \tag{21}$$

   Moreover, if the limit in (21) equals 1, then player i's belief is the same as in the original equilibrium: for any $\dot{\mathfrak{h}}_0^t$ and $\tilde{m}_{-i,t}$,

$$\begin{aligned} & \lim_k \Pr^{\left(\sigma^k,\mu^k\right)}\left(\dot{\mathfrak{h}}_0^t, \tilde{m}_{-i,t}, a_{-i,t}, s_{-i,t+1} | \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, m_{i,t} = \tilde{m}_{i,t}, a_{i,t}, s_{i,t+1}\right) \\ = \; & \lim_k F\left(\alpha^k\right)\left(\dot{\mathfrak{h}}_0^t, \tilde{m}_{-i,t}, a_{-i,t}, s_{-i,t+1} | c(\mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}), \tilde{m}_{i,t}, a_{i,t}, s_{i,t+1}\right). \end{aligned} \tag{22}$$

The proof of Lemma 5 is the most technical part of the proof and is deferred to Section 12.6. The lemma express the key idea that a player's belief can switch from a point mass on $\theta_t = \rho$ to a point mass on $\theta_t = \pi$ 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 $m_{i,t+1} \in A_{i,t+1}$ (applying part 1 of the lemma for $t+1$).

12.5 Sequential Rationality

We verify that σ_i^* is a best response to $\left(\sigma_{-i}^*, \mu^*\right)$. If player i has been unfaithful, σ_i^* is a best response by upper hemi-continuity of the best response correspondence (as $(\sigma^*, \mu^*) = \lim_k \left(\sigma^k, \mu^k\right)$, where σ_i^k is a best response to $\left(\sigma_{-i}^k, \mu^k\right)$ at every unfaithful history). So assume player i's history $\mathfrak{h}_i^t$ is faithful.

We consider two cases, depending on whether the player's most recent message from the mediator is an action recommendation or message $\star$.

If $m_{i,t} \in A_i$ (so $m_{i,t} = \tilde{m}_{i,t}$), we show that following σ_i^* is optimal for each possible value of $(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}$ (and hence is optimal conditional on player i's information). By (20), player i

believes $(\dot{\mathfrak{h}}_0^t, \tilde{m}_{-i,t})$ is distributed according to

$$\lim_k F\left(\alpha^k\right)\left(\dot{\mathfrak{h}}_0^t, \tilde{m}_{-i,t} | c\left(\mathfrak{h}_i^t, \left(\tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}\right), \tilde{m}_{i,t}\right),$$

and she believes $\theta_t = \rho$ with probability 1. As she also believes her opponents have been faithful with probability 1 (by Lemma 4), this implies future play will be given by α^k. Hence, player i's expected payoff from each action is the same as in the original CPPBE (σ, μ, β), so playing $a_{i,t} = m_{i,t}$ is optimal.

In addition, after observing $s_{i,t+1}$, (21) holds. If the limiting belief is 1, then (22) implies truthful reporting is optimal (by the same argument as for obedience in period t). If it is 0, then with probability 1 the mediator sends message $\star$ in every future period regardless of the players' reports (as $\lim_k \varepsilon^k(\mathfrak{h}_0^{t+1}) = 0 \ \forall \mathfrak{h}_0^{t+1}$). In this case, any report is optimal.

If instead $m_{i,t} = \star$, then σ_i^* prescribes an optimal action at history $(\mathfrak{h}_i^t, m_{i,t})$ by upper hemi-continuity. In addition, $m_{i,t} = \star$ implies $\theta_t = \pi$, and hence with probability 1 the mediator sends message $\star$ in every future period regardless of the players' reports. So any report is optimal.

12.6 Proof of Lemma 5

We use the following elementary fact:

Lemma 6 *For a finite set X and functions $f^k : X \to \mathbb{R}_{>0}, g^k : X \to \mathbb{R}_{>0}, h^k : X \to \mathbb{R}_{>0}$ indexed by $k \in \mathbb{N}$, if $\lim_k \frac{f^k(x)}{g^k(x)} = 1 \ \forall x \in X$ and $\lim_k \frac{f^k(x)}{\sum_{x' \in X} f^k(x') h^k(x')} \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ exists, then $\lim_k \frac{f^k(x)}{\sum_{x' \in X} f^k(x') h^k(x')} = \lim_k \frac{g^k(x)}{\sum_{x' \in X} g^k(x') h^k(x')}$.*

Proof. As $\left(\frac{f^k(x)}{\sum_{x' \in X} f^k(x') h^k(x')}\right)_k$ is a sequence of positive numbers converging to a limit, $\lim_k \sum_{x' \in X} \frac{f^k(x') h^k(x')}{f^k(x)}$ exists as well. Hence, we have

$$\begin{aligned}
\lim_k \sum_{x' \in X} \frac{f^k(x') h^k(x')}{f^k(x)} &= \lim_k \sum_{x' \in X} \frac{f^k(x')}{g^k(x')} \frac{g^k(x)}{f^k(x)} \frac{g^k(x')}{g^k(x)} h^k(x') \\
&= \sum_{x' \in X} \lim_k \frac{f^k(x')}{g^k(x')} \lim_k \frac{g^k(x)}{f^k(x)} \lim_k \frac{g^k(x') h^k(x')}{g^k(x)} \\
&= \lim_k \sum_{x' \in X} \frac{g^k(x') h^k(x')}{g^k(x)}.
\end{aligned}$$

If $\lim_k \frac{f^k(x)}{\sum_{x' \in X} f^k(x') h^k(x')}$ is finite, then (with convention $0 = 1/\infty$)

$$\begin{aligned}
\lim_k \frac{f^k(x)}{\sum_{x' \in X} f^k(x') h^k(x')} &= \left[\lim_k \sum_{x' \in X} \frac{f^k(x') h^k(x')}{f^k(x)}\right]^{-1} \\
&= \left[\lim_k \sum_{x' \in X} \frac{g^k(x') h^k(x')}{g^k(x)}\right]^{-1} \\
&= \lim_k \frac{g^k(x)}{\sum_{x' \in X} g^k(x') h^k(x')}.
\end{aligned}$$

If $\lim_k \frac{f^k(x)}{\sum_{x' \in X} f^k(x') h^k(x')} = \infty$, then $\lim_k \sum_{x' \in X} \frac{f^k(x') h^k(x')}{f^k(x)} = 0$, so $\lim_k \sum_{x' \in X} \frac{g^k(x') h^k(x')}{g^k(x)} = 0$. Since $\frac{g^k(x') h^k(x')}{g^k(x)} > 0$ for all k, this implies $\lim_k \frac{g^k(x)}{\sum_{x' \in X} g^k(x') h^k(x')} = \infty$. ■

Turning to the proof of Lemma 5, we reduce notation by writing $\Pr^k$ for $\Pr^{(\sigma^k, \mu^k)}$ and F^k for $F(\alpha^k)$. We first derive Equation (13):

Lemma 7 *For any faithful history $\mathfrak{h}_i^t$, any compatible sequence of auxiliary recommendations $(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}$, and any canonical history $\dot{\mathfrak{h}}_0^t$, we have*

$$\lim_k \frac{\Pr^k\left(\dot{\mathfrak{h}}_0^t | \mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho\right)}{F^k\left(\dot{\mathfrak{h}}_0^t | c(\mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1})\right)} = 1. \tag{23}$$

Proof. We claim it suffices to show that, for each $\mathfrak{h}_0^t, m_{i,t}, \tilde{m}_t, a_t$,

$$\lim_k \frac{\Pr^k\left(m_{i,t}, \tilde{m}_t, a_t, \hat{\theta}_{t+1} = \rho | \mathcal{F}, \mathfrak{h}_0^t, \hat{\theta}_t = \rho\right)}{F^k\left(\tilde{m}_t, a_t | c(\mathfrak{h}_0^t)\right)} = 1. \tag{24}$$

Intuitively, (24) implies (23) because, given $\mathcal{F}$, $\mathfrak{h}_0^t$ contains all the information in $\left(\mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}\right)$. To prove that (24) implies (23), suppose (23) holds for $t = \bar{t}$ and (24) holds for every t. Since $\hat{\theta}_{\bar{t}+1} = \rho$ implies $\hat{\theta}_{\bar{t}} = \rho$, we have

$$\lim_k \frac{\Pr^k\left(\dot{\mathfrak{h}}_0^{\bar{t}+1} | \mathcal{F}, \mathfrak{h}_i^{\bar{t}+1}, (\tilde{m}_{i,\tau})_{\tau=1}^{\bar{t}}, \hat{\theta}_{\bar{t}+1} = \rho\right)}{F^k\left(\dot{\mathfrak{h}}_0^{\bar{t}+1} | c(\mathfrak{h}_i^{\bar{t}+1}, (\tilde{m}_{i,\tau})_{\tau=1}^{\bar{t}})\right)}$$

$$= \lim_k \frac{\sum_{\mathfrak{h}_0^{\bar{t}}:c(\mathfrak{h}_0^{\bar{t}})=\dot{\mathfrak{h}}_0^{\bar{t}}} \left(\begin{array}{c} \Pr^k\left(\mathfrak{h}_0^{\bar{t}}|\mathcal{F},\mathfrak{h}_i^{\bar{t}},(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_{\bar{t}}=\rho\right) \\ \times \left(\begin{array}{c} \Pr^k\left(m_{i,\bar{t}},\tilde{m}_{\bar{t}},a_{\bar{t}},\hat{\theta}_{\bar{t}+1}=\rho|\mathcal{F},\mathfrak{h}_0^{\bar{t}},\mathfrak{h}_i^{\bar{t}},\hat{\theta}_{\bar{t}}=\rho\right) \\ \times \Pr^k\left(s_{\bar{t}+1}|\mathcal{F},\mathfrak{h}_0^{\bar{t}},\mathfrak{h}_i^{\bar{t}},m_{i,\bar{t}},\tilde{m}_{\bar{t}},a_{\bar{t}},\hat{\theta}_{\bar{t}+1}=\rho\right) \end{array} \right) \end{array} \right)}{\left(\begin{array}{c} \sum_{\mathfrak{h}_0^{\bar{t}}} \left(\begin{array}{c} \Pr^k\left(\mathfrak{h}_0^{\bar{t}}|\mathcal{F},\mathfrak{h}_i^{\bar{t}},(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_{\bar{t}}=\rho\right) \\ \times \sum_{\tilde{m}'_{-i,\bar{t}},a'_{-i,\bar{t}}} \left(\begin{array}{c} \Pr^k\left(m_{i,\bar{t}},\tilde{m}_{i,\bar{t}},\tilde{m}'_{-i,\bar{t}},a_{i,\bar{t}},a'_{-i,\bar{t}},\hat{\theta}_{\bar{t}+1}=\rho|\mathcal{F},\mathfrak{h}_0^{\bar{t}},\mathfrak{h}_i^{\bar{t}},\hat{\theta}_{\bar{t}}=\rho\right) \\ \times \Pr^k\left(s_{i,\bar{t}+1}|\mathcal{F},\mathfrak{h}_0^{\bar{t}},\mathfrak{h}_i^{\bar{t}},m_{i,\bar{t}},\tilde{m}_{i,\bar{t}},\tilde{m}'_{-i,\bar{t}},a_{i,\bar{t}},a'_{-i,\bar{t}},\hat{\theta}_{\bar{t}+1}=\rho\right) \end{array} \right) \end{array} \right) \\ \times F^k\left(\dot{\mathfrak{h}}_0^{\bar{t}+1}|c(\mathfrak{h}_i^{\bar{t}+1},(\tilde{m}_{i,\tau})_{\tau=1}^{\bar{t}})\right) \end{array} \right)}.$$

Since $\mathfrak{h}_i^{\bar{t}}$ is determined by $\mathfrak{h}_0^{\bar{t}}$ (given $\mathcal{F}$), we can omit $\mathfrak{h}_i^{\bar{t}}$ from the conditioning event when we condition on $\mathcal{F},\mathfrak{h}_0^{\bar{t}}$. In addition, given $\mathcal{F}$, the distribution of $s_{\bar{t}+1}$ is determined by the payoff-relevant history:

$$\Pr^k\left(s_{\bar{t}+1}|\mathcal{F},\mathfrak{h}_0^{\bar{t}},m_{i,\bar{t}},\tilde{m}_{\bar{t}},a_{\bar{t}},\hat{\theta}_{\bar{t}+1}=\rho\right) = p\left(s_{\bar{t}+1}|\mathring{c}(c(\mathfrak{h}_0^{\bar{t}})),a_{\bar{t}}\right).$$

Hence, the last expression equals

$$\lim_k \frac{\sum_{\mathfrak{h}_0^{\bar{t}}:c(\mathfrak{h}_0^{\bar{t}})=\dot{\mathfrak{h}}_0^{\bar{t}}} \left(\begin{array}{c} \Pr^k\left(\mathfrak{h}_0^{\bar{t}}|\mathcal{F},\mathfrak{h}_i^{\bar{t}},(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_{\bar{t}}=\rho\right) \\ \times \left(\begin{array}{c} \Pr^k\left(m_{i,\bar{t}},\tilde{m}_{\bar{t}},a_{\bar{t}},\hat{\theta}_{\bar{t}+1}=\rho|\mathcal{F},\mathfrak{h}_0^{\bar{t}},\hat{\theta}_{\bar{t}}=\rho\right) \\ \times p\left(s_{\bar{t}+1}|\mathring{c}(c(\mathfrak{h}_0^{\bar{t}})),a_{\bar{t}}\right) \end{array} \right) \end{array} \right)}{\left(\begin{array}{c} \sum_{\mathfrak{h}_0^{\bar{t}}} \left(\begin{array}{c} \Pr^k\left(\mathfrak{h}_0^{\bar{t}}|\mathcal{F},\mathfrak{h}_i^{\bar{t}},(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_{\bar{t}}=\rho\right) \\ \times \sum_{\tilde{m}'_{-i,\bar{t}},a'_{-i,\bar{t}}} \left(\begin{array}{c} \Pr^k\left(m_{i,\bar{t}},\tilde{m}_{i,\bar{t}},\tilde{m}'_{-i,\bar{t}},a_{i,\bar{t}},a'_{-i,\bar{t}},\hat{\theta}_{\bar{t}+1}=\rho|\mathcal{F},\mathfrak{h}_0^{\bar{t}},\hat{\theta}_{\bar{t}}=\rho\right) \\ \times p\left(s_{i,\bar{t}+1}|\mathring{c}(c(\mathfrak{h}_0^{\bar{t}})),a_{i,\bar{t}},a'_{-i,\bar{t}}\right) \end{array} \right) \end{array} \right) \\ \times F^k\left(\dot{\mathfrak{h}}_0^{\bar{t}+1}|c(\mathfrak{h}_i^{\bar{t}+1},(\tilde{m}_{i,\tau})_{\tau=1}^{\bar{t}})\right) \end{array} \right)}.$$

Using Lemma 6 to substitute (24) into the second term in the numerator and denominator, this equals

$$\lim_k \frac{\sum_{\mathfrak{h}_0^{\bar{t}}:c(\mathfrak{h}_0^{\bar{t}})=\dot{\mathfrak{h}}_0^{\bar{t}}} \left(\Pr^k\left(\mathfrak{h}_0^{\bar{t}}|\mathcal{F},\mathfrak{h}_i^{\bar{t}},(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_{\bar{t}}=\rho\right) F^k\left(\tilde{m}_{\bar{t}},a_{\bar{t}}|c(\mathfrak{h}_0^{\bar{t}})\right) p\left(s_{\bar{t}+1}|\mathring{c}(c(\mathfrak{h}_0^{\bar{t}})),a_{\bar{t}}\right)\right)}{\left(\begin{array}{c} \sum_{\mathfrak{h}_0^{\bar{t}}} \left(\begin{array}{c} \Pr^k\left(\mathfrak{h}_0^{\bar{t}}|\mathcal{F},\mathfrak{h}_i^{\bar{t}},(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_{\bar{t}}=\rho\right) \\ \times \sum_{\tilde{m}'_{-i,\bar{t}},a'_{-i,\bar{t}}} \left(F^k\left(m_{i,\bar{t}},\tilde{m}_{i,\bar{t}},\tilde{m}'_{-i,\bar{t}},a_{i,\bar{t}},a'_{-i,\bar{t}}|c(\mathfrak{h}_0^{\bar{t}})\right) p\left(s_{i,\bar{t}+1}|\mathring{c}(c(\mathfrak{h}_0^{\bar{t}})),a_{i,\bar{t}},a'_{-i,\bar{t}}\right)\right) \end{array} \right) \\ \times F^k\left(\dot{\mathfrak{h}}_0^{\bar{t}+1}|c(\mathfrak{h}_i^{\bar{t}+1},(\tilde{m}_{i,\tau})_{\tau=1}^{\bar{t}})\right). \end{array} \right)}.$$

Since everything except for the first term in the denominator depends only on $c(\mathfrak{h}_0^{\bar{t}})$, by the inductive hypothesis this equals

$$
\begin{aligned}
&\lim_k \frac{F^k\left(\dot{\mathfrak{h}}_0^{\bar{t}}|c(\mathfrak{h}_i^{\bar{t}},(\tilde{m}_{i,\tau})_{\tau=1}^{t-1})\right)\sum_{\mathfrak{h}_0^{\bar{t}}:c(\mathfrak{h}_0^{\bar{t}})=\dot{\mathfrak{h}}_0^{\bar{t}}} F^k\left(\tilde{m}_{\bar{t}},a_{\bar{t}}|c(\mathfrak{h}_0^{\bar{t}})\right)p\left(s_{\bar{t}+1}|\mathring{c}(c(\mathfrak{h}_0^{\bar{t}})),a_{\bar{t}}\right)}{\sum_{\widetilde{\dot{\mathfrak{h}}}_0^{\bar{t}}}\left(\begin{array}{c} F^k\left(\widetilde{\dot{\mathfrak{h}}}_0^{\bar{t}}|c(\mathfrak{h}_i^{\bar{t}},(\tilde{m}_{i,\tau})_{\tau=1}^{t-1})\right) \\ \times\sum_{\tilde{m}'_{-i,\bar{t}},a'_{-i,\bar{t}}} F^k\left(\tilde{m}_{i,\bar{t}},\tilde{m}'_{-i,\bar{t}},a_{i,\bar{t}},a'_{-i,\bar{t}}|\widetilde{\dot{\mathfrak{h}}}_0^{\bar{t}}\right)p\left(s_{i,\bar{t}+1}|\mathring{c}(\widetilde{\dot{\mathfrak{h}}}_0^{\bar{t}}),a_{i,\bar{t}},a'_{-i,\bar{t}}\right)\end{array}\right)} \\
&\times F^k\left(\dot{\mathfrak{h}}_0^{\bar{t}+1}|c(\mathfrak{h}_i^{\bar{t}+1},(\tilde{m}_{i,\tau})_{\tau=1}^{\bar{t}})\right)^{-1}. \qquad (25)
\end{aligned}
$$

Finally, the first term in (25) equals $F^k\left(\dot{\mathfrak{h}}_0^{\bar{t}+1}|c(\mathfrak{h}_i^{\bar{t}+1},(\tilde{m}_{i,\tau})_{\tau=1}^{\bar{t}})\right)$ by Bayes' rule, so (25) equals 1. Hence, (23) holds for $\bar{t}+1$, and it therefore holds for every t by induction.[23]

We now prove (24). For any $\tilde{m}_t, a_t \in A_t$, let $I^{**}(\tilde{m}_t, a_t) = \{i \in I : \tilde{m}_{i,t} = a_{i,t}\}$. We establish (24) separately for the case where $\tilde{m}_t = a_t$ (i.e., $I^{**}(\tilde{m}_t, a_t) = I$) and the case where $\tilde{m}_t \neq a_t$ (i.e., $I^{**}(\tilde{m}_t, a_t) = I$).

Case 1: $\tilde{m}_t = a_t$. In this case, given $\mathcal{F}$, if $\theta_t = \pi$ then $\varepsilon^k\left(\mathfrak{h}_0^{t+1}\right) = 0$, and therefore $\hat{\theta}_{t+1} = \pi$. Hence, $\hat{\theta}_{t+1} = \rho$ implies $\theta_t = \rho$, and thus $a_t = m_t = \tilde{m}_t$ (given $\mathcal{F}$). Therefore, by Bayes' rule,

$$
\begin{aligned}
&\lim_k \frac{\Pr^k\left(m_{i,t},\tilde{m}_t,a_t,\hat{\theta}_{t+1}=\rho|\mathcal{F},\mathfrak{h}_0^t,\hat{\theta}_t=\rho\right)}{F^k\left(\tilde{m}_t,a_t|c(\mathfrak{h}_0^t)\right)} \\
= &\lim_k \frac{\Pr^k\left(\tilde{m}_t|\mathcal{F},\mathfrak{h}_0^t,\hat{\theta}_t=\theta_t=\rho\right)\Pr^k\left(\hat{\theta}_{t+1}=\rho|\mathcal{F},\mathfrak{h}_0^t,\tilde{m}_t,\hat{\theta}_t=\theta_t=\rho\right)}{F^k\left(\tilde{m}_t|c(\mathfrak{h}_0^t)\right)F^k\left(a_t|c(\mathfrak{h}_0^t),\tilde{m}_t\right)}.
\end{aligned}
$$

Now, for each $\mathfrak{h}_0^t, \tilde{m}_t$,

$$
\begin{aligned}
\lim_k \Pr^k\left(\hat{\theta}_{t+1}=\rho|\mathcal{F},\mathfrak{h}_0^t,\tilde{m}_t,\hat{\theta}_t=\theta_t=\rho\right) &= 1 \text{ as } f_2(k)\to 0, \\
\lim_k F^k\left(a_t|c(\mathfrak{h}_0^t),\tilde{m}_t\right) &= 1 \text{ as } a_t=\tilde{m}_t.
\end{aligned}
$$

Hence, by Lemma 6, the above expression equals

$$
\lim_k \frac{\Pr^k\left(\tilde{m}_t|\mathcal{F},\mathfrak{h}_0^t,\hat{\theta}_t=\theta_t=\rho,I^{**}(\tilde{m}_t,a_t)=I\right)}{F^k\left(\tilde{m}_t|c(\mathfrak{h}_0^t)\right)}.
$$

Since $\Pr^k\left(\tilde{m}_t|\mathcal{F},\mathfrak{h}_0^t,\hat{\theta}_t=\theta_t=\rho\right) = F^k\left(\tilde{m}_t|c(\mathfrak{h}_0^t)\right)$ by (10), (24) holds.

Case 2: $\tilde{m}_t \neq a_t$. In this case, given $\mathcal{F}$, $m_{j,t} = \star$ for some player j, and hence $\theta_t = \pi$. Moreover, given $\mathcal{F}$ and $\hat{\theta}_{t+1} = \rho$, for each player j, $j \in I^{**}(\tilde{m}_t, a_t)$ if and only if $a_{j,t} = m_{j,t} =$

[23] Note that we have also shown that the left hand side of (23) converges, as every convergent subsequence converges to 1 and the existence of a divergent subsequence is contradicted by (25).

$\tilde{m}_{j,t}$ (equivalently, $j \in I_t^*$), as otherwise $\varepsilon^k\left(\mathfrak{h}_0^{t+1}\right) = 0$ and $\hat{\theta}_{t+1}$ would equal π. Hence,

$$\begin{aligned}
& \lim_k \frac{\Pr^k\left(m_{i,t}, \tilde{m}_t, a_t, \hat{\theta}_{t+1} = \rho | \mathcal{F}, \mathfrak{h}_0^t, \hat{\theta}_t = \rho\right)}{F^k\left(\tilde{m}_t, a_t | c(\mathfrak{h}_0^t)\right)} \\
= & \lim_k \frac{\left(\begin{array}{c} \Pr^k\left(\tilde{m}_t | \mathcal{F}, \mathfrak{h}_0^t, \hat{\theta}_t = \rho\right) \\ \times \Pr^k\left(I_t^* = I^{**}\left(\tilde{m}_t, a_t\right), a_t, \hat{\theta}_{t+1} = \rho | \mathcal{F}, \mathfrak{h}_0^t, \tilde{m}_t, \hat{\theta}_t = \rho\right) \end{array}\right)}{F^k\left(\tilde{m}_t | c(\mathfrak{h}_0^t)\right) F^k\left(a_t | c(\mathfrak{h}_0^t), \tilde{m}_t\right)}.
\end{aligned}$$

For each $\mathfrak{h}_0^t, \tilde{m}_t$ and each k,

$$\begin{aligned}
\Pr^k\left(\tilde{m}_t | \mathcal{F}, \mathfrak{h}_0^t, \hat{\theta}_t = \rho\right) &= F^k\left(\tilde{m}_t | c(\mathfrak{h}_0^t)\right) \quad \text{by (10)}, \\
\Pr^k\left(I_t^* = I^{**}\left(\tilde{m}_t, a_t\right), a_t, \hat{\theta}_{t+1} = \rho | \mathcal{F}, \mathfrak{h}_0^t, \tilde{m}_t, \hat{\theta}_t = \rho\right) &= F^k\left(a_t | c(\mathfrak{h}_0^t), \tilde{m}_t\right) \quad \text{by (15)}.
\end{aligned}$$

Hence, Lemma 6 implies the above expression equals 1, so (24) holds. ■

We can now prove Lemma 5.

Proof of (20): Given Lemma 4, it suffices to show

$$\lim_k \frac{\Pr^k\left(\dot{\mathfrak{h}}_0^t, m_{-i,t} = \tilde{m}_{-i,t}, \theta_t = \rho | \mathcal{F}, \mathfrak{h}_i^t, \left(\tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}, m_{i,t} = \tilde{m}_{i,t}\right)}{F^k\left(\dot{\mathfrak{h}}_0^t, \tilde{m}_{-i,t} | c(\mathfrak{h}_i^t, \left(\tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}), \tilde{m}_{i,t}\right)} = 1. \tag{26}$$

Since $m_{i,t} = \tilde{m}_{i,t}$ implies $\hat{\theta}_t = \rho$, we can condition on $\hat{\theta}_t = \rho$ in the numerator. In addition, by Bayes' rule, the numerator equals

$$\frac{\Pr^k\left(\dot{\mathfrak{h}}_0^t, m_{-i,t} = \tilde{m}_{-i,t}, m_{i,t} = \tilde{m}_{i,t}, \theta_t = \rho | \mathcal{F}, \mathfrak{h}_i^t, \left(\tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}, \hat{\theta}_t = \rho\right)}{\Pr^k\left(m_{i,t} = \tilde{m}_{i,t} | \mathcal{F}, \mathfrak{h}_i^t, \left(\tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}, \hat{\theta}_t = \rho\right)}.$$

Note that

$$\begin{aligned}
& \Pr^k\left(\dot{\mathfrak{h}}_0^t, m_{-i,t} = \tilde{m}_{-i,t}, m_{i,t} = \tilde{m}_{i,t}, \theta_t = \rho | \mathcal{F}, \mathfrak{h}_i^t, \left(\tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}, \hat{\theta}_t = \rho\right) \\
= & \Pr^k\left(\dot{\mathfrak{h}}_0^t | \mathcal{F}, \mathfrak{h}_i^t, \left(\tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}, \hat{\theta}_t = \rho\right) \Pr^k\left(\theta_t = \rho | \mathcal{F}, \dot{\mathfrak{h}}_0^t, \mathfrak{h}_i^t, \left(\tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}, \hat{\theta}_t = \rho\right) \\
& \times \Pr^k\left(\tilde{m}_t | \mathcal{F}, \dot{\mathfrak{h}}_0^t, \mathfrak{h}_i^t, \left(\tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}, \hat{\theta}_t = \theta_t = \rho\right) \Pr^k\left(m_t = \tilde{m}_t | \mathcal{F}, \dot{\mathfrak{h}}_0^t, \mathfrak{h}_i^t, \left(\tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}, \hat{\theta}_t = \theta_t = \rho, \tilde{m}_t\right) \\
= & \Pr^k\left(\dot{\mathfrak{h}}_0^t | \mathcal{F}, \mathfrak{h}_i^t, \left(\tilde{m}_{i,\tau}\right)_{\tau=1}^{t-1}, \hat{\theta}_t = \rho\right) \Pr^k\left(\theta_t = \rho | \hat{\theta}_t = \rho\right) F^k\left(\tilde{m}_t | \dot{\mathfrak{h}}_0^t\right),
\end{aligned}$$

since the distribution of θ_t is determined by $\hat{\theta}_t$ (by (12)), the distribution of $\tilde{m}_t$ is determined by $\dot{\mathfrak{h}}_0^t$ (by (10)), and $m_t = \tilde{m}_t$ with probability 1 if $\theta_t = \rho$. Hence, the left-hand side of (26) equals

$$\lim_k \frac{\Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t=\rho\right)\Pr^k\left(\theta_t=\rho|\hat{\theta}_t=\rho\right)F^k\left(\tilde{m}_t|\dot{\mathfrak{h}}_0^t\right)}{\left(\begin{array}{c}\sum_{\tilde{\dot{\mathfrak{h}}}_0^t}\Pr^k\left(\tilde{\dot{\mathfrak{h}}}_0^t|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t=\rho\right)\Pr^k\left(\theta_t=\rho|\hat{\theta}_t=\rho\right)F^k\left(\tilde{m}_{i,t}|\tilde{\dot{\mathfrak{h}}}_0^t\right)\\ \times F^k\left(\dot{\mathfrak{h}}_0^t, \tilde{m}_{-i,t}|c(\mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}), \tilde{m}_{i,t}\right)\end{array}\right)}.$$

By Lemma 6 and (23), this expression equals 1, as desired:

$$\lim_k \frac{F^k\left(\dot{\mathfrak{h}}_0^t|c(\mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1})\right)F^k\left(\tilde{m}_t|\dot{\mathfrak{h}}_0^t\right)}{\left(\begin{array}{c}\sum_{\tilde{\dot{\mathfrak{h}}}_0^t}F^k\left(\tilde{\dot{\mathfrak{h}}}_0^t|c(\mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1})\right)F^k\left(\tilde{m}_{i,t}|\tilde{\dot{\mathfrak{h}}}_0^t\right)\\ \times F^k\left(\dot{\mathfrak{h}}_0^t, \tilde{m}_{-i,t}|c(\mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}), \tilde{m}_{i,t}\right)\end{array}\right)} = 1.$$

Proof of (21) and (22): Again, it suffices to establish these equalities conditional on $\mathcal{F}$ (by Lemma 4). Intuitively, before observing $s_{i,t+1}$, player i's beliefs coincides with those under F^k, by (20). If there exists a positive probability recommendation $\tilde{m}_{-i,t}$ such that $s_{i,t+1}$ arises with positive probability given action profile $(a_{i,t}, a_{-i,t}=\tilde{m}_{-i,t})$ ("Case 1" below), then player i believes $\theta_t=\rho$ with probability 1, and her beliefs continue to coincide with F^k. Otherwise ("Case 2"), since action trembles occur with much smaller probability when $\theta_t=\rho$ than when $\theta_t=\pi$, player i believes $\theta_t=\pi$ with probability 1.

Formally, we use the following equation:

Lemma 8 *For any $\mathfrak{h}_i^t$, compatible $(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}$, $\tilde{m}_{i,t}$, and $s_{i,t+1}$,*

$$\begin{aligned}
&\lim_k \Pr^k\left(\theta_t=\rho|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, m_{i,t}=\tilde{m}_{i,t}, s_{i,t+1}\right) \qquad (27)\\
=\; &\lim_k \frac{\sum_{\dot{\mathfrak{h}}_0^t}\left(\begin{array}{c}\Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t=\rho, \tilde{m}_{i,t}\right)\\ \times\sum_{\tilde{m}_{-i,t}}\left(F^k\left(\tilde{m}_{-i,t}|\tilde{m}_{i,t}, \dot{\mathfrak{h}}_0^t\right)(1-f_2(k))\,p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t), a_{i,t}, a_{-i,t}=\tilde{m}_{-i,t}\right)\right)\end{array}\right)}{\begin{array}{c}\sum_{\dot{\mathfrak{h}}_0^t}\left(\begin{array}{c}\Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t=\rho, \tilde{m}_{i,t}\right)\\ \times\sum_{\tilde{m}_{-i,t}}\left(F^k\left(\tilde{m}_{-i,t}|\tilde{m}_{i,t}, \dot{\mathfrak{h}}_0^t\right)(1-f_2(k))\,p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t), a_{i,t}, a_{-i,t}=\tilde{m}_{-i,t}\right)\right)\end{array}\right)\\ +\sum_{\dot{\mathfrak{h}}_0^t}\left(\begin{array}{c}\Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t=\rho, \tilde{m}_{i,t}\right)\\ \times\sum_{\tilde{m}_{-i,t}}\left(\begin{array}{c}F^k\left(\tilde{m}_{-i,t}|\tilde{m}_{i,t}, \dot{\mathfrak{h}}_0^t\right)(f_2(k))^2\\ \sum_{a_{-i,t}}\Pr^k\left(a_{-i,t}|\mathcal{F}, \dot{\mathfrak{h}}_0^t, \mathfrak{h}_i^t, \theta_t=\pi, \tilde{m}_t\right)p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t), a_t\right)\end{array}\right)\end{array}\right)\end{array}}.
\end{aligned}$$

In this equation, the numerator represents the probability that $\theta_t=\rho$, where $a_{-i,t}=\tilde{m}_{-i,t}$ given $\mathcal{F}$ and $\theta_t=\rho$. The denominator adds the probability that $\theta_t\neq\rho$, accounting for the $f_2(k)$-probability event that $m_{i,t}=\star$. The equation follows mechanically from Bayes' rule and simplifying terms. The details are deferred to the end of the proof.

Given $\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}$, let us partition the set of canonical histories $\dot{H}_0^t \ni \dot{\mathfrak{h}}_0^t$ into sets $\dot{H}_0^t[1], ..., \dot{H}_0^t[L] \subseteq \dot{H}_0^t$ according to the following criteria:

$$\begin{aligned}
\lim_k \frac{\Pr^k\left(\mathfrak{h}|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right)}{\Pr^k\left(\mathfrak{h}'|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right)} &\in (0,\infty) \text{ for each } \mathfrak{h}, \mathfrak{h}' \in \dot{H}_0^t[l] \text{ with } l = 1, ..., L; \\
\lim_k \frac{\Pr^k\left(\mathfrak{h}|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right)}{\Pr^k\left(\mathfrak{h}'|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right)} &= 0 \text{ for each } \mathfrak{h} \in \dot{H}_0^t[l], \mathfrak{h}' \in \dot{H}_0^t[l'] \text{ with } l' < l.
\end{aligned}$$

That is, each $\dot{H}_0^t[l]$ is a set of canonical histories with mutually bounded likelihood ratios, and each $\dot{H}_0^t[l']$ with $l' < l$ is a set of qualitatively more likely histories.

Given $a_{i,t}, s_{i,t+1}$, let l^* be the smallest number l such that there exists $\dot{\mathfrak{h}}_0^t \in \dot{H}_0^t[l]$ with $p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t), a_{i,t}, a_{-i,t}\right) > 0$ for some $a_{-i,t}$. Since $l < l^*$ and $\dot{\mathfrak{h}}_0^t \in \dot{H}_0^t[l]$ imply

$$p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t), a_{i,t}, a_{-i,t} = \tilde{m}_{-i,t}\right) = 0 \quad \forall a_{-i,t} \in A_{-i,t},$$

in computing (27), we may restrict attention to $\dot{\mathfrak{h}}_0^t \in \dot{H}_0^t[l]$ with $l \geq l^*$.

We now show that, in computing (27), we may further restrict attention to $\dot{\mathfrak{h}}_0^t \in \dot{H}_0^t[l^*]$:

$$\begin{aligned}
&\lim_k \Pr^k\left(\theta_t = \rho|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, m_{i,t} = \tilde{m}_{i,t}, s_{i,t+1}\right) \qquad (28) \\
= &\lim_k \frac{\sum_{\dot{\mathfrak{h}}_0^t \in \dot{H}_0^t[l^*]}\left(\begin{array}{c}\Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right) \\ \times \sum_{\tilde{m}_{-i,t}}\left(F^k\left(\tilde{m}_{-i,t}|\dot{\mathfrak{h}}_0^t, \tilde{m}_{i,t}\right)(1 - f_2(k))\, p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t), a_{i,t}, a_{-i,t} = \tilde{m}_{-i,t}\right)\right)\end{array}\right)}{\begin{array}{c}\sum_{\dot{\mathfrak{h}}_0^t \in \dot{H}_0^t[l^*]}\left(\begin{array}{c}\Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right) \\ \times \sum_{\tilde{m}_{-i,t}}\left(F^k\left(\tilde{m}_{-i,t}|\tilde{m}_{i,t}, \dot{\mathfrak{h}}_0^t\right)(1 - f_2(k))\, p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t), a_{i,t}, a_{-i,t} = \tilde{m}_{-i,t}\right)\right)\end{array}\right) \\ + \sum_{\dot{\mathfrak{h}}_0^t \in \dot{H}_0^t[l^*]}\left(\begin{array}{c}\Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right) \\ \times \sum_{\tilde{m}_{-i,t}}\left(\begin{array}{c}F^k\left(\tilde{m}_{-i,t}|\tilde{m}_{i,t}, \dot{\mathfrak{h}}_0^t\right)(f_2(k))^2 \\ \sum_{a_{-i,t}} \Pr^k\left(a_{-i,t}|\mathcal{F}, \dot{\mathfrak{h}}_0^t, \mathfrak{h}_i^t, \theta_t = \pi, \tilde{m}_t\right) p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t), a_t\right)\end{array}\right)\end{array}\right)\end{array}}.
\end{aligned}$$

To see why, for each $\mathfrak{h} \in \dot{H}_0^t[l^*]$ and $\mathfrak{h}' \in \dot{H}_0^t[l]$ with $l > l^*$, we have

$$\frac{\Pr^k\left(\mathfrak{h}'|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right)}{\Pr^k\left(\mathfrak{h}|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right)} = \frac{\frac{\Pr^k\left(\mathfrak{h}'|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho\right)\Pr^k\left(\tilde{m}_{i,t}|\mathcal{F}, \mathfrak{h}', \mathfrak{h}_i^t, \hat{\theta}_t = \rho\right)}{\sum_{\tilde{\mathfrak{h}}}\Pr^k\left(\tilde{\mathfrak{h}}|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho\right)\Pr^k\left(\tilde{m}_{i,t}|\mathcal{F}, \tilde{\mathfrak{h}}, \mathfrak{h}_i^t, \hat{\theta}_t = \rho\right)}}{\frac{\Pr^k\left(\mathfrak{h}|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho\right)\Pr^k\left(\tilde{m}_{i,t}|\mathcal{F}, \mathfrak{h}, \mathfrak{h}_i^t, \hat{\theta}_t = \rho\right)}{\sum_{\tilde{\mathfrak{h}}}\Pr^k\left(\tilde{\mathfrak{h}}|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho\right)\Pr^k\left(\tilde{m}_{i,t}|\mathcal{F}, \tilde{\mathfrak{h}}, \mathfrak{h}_i^t, \hat{\theta}_t = \rho\right)}},$$

where the summations are taken over $\tilde{\mathfrak{h}} \in \dot{H}_0^t$. For each $\tilde{m}_{i,t}$, $\tilde{\mathfrak{h}}$, and $\mathfrak{h}_i^t$, by (12) and (10),

$$\Pr^k\left(\tilde{m}_{i,t}|\mathcal{F}, \tilde{\mathfrak{h}}, \mathfrak{h}_i^t, \hat{\theta}_t = \rho\right) = F^k\left(\tilde{m}_{i,t}|\tilde{\mathfrak{h}}\right).$$

Hence, for sufficiently large k, for each $\dot{\mathfrak{h}}_0^t \in \dot{H}_0^t[l^*]$ and $\mathfrak{h}' \in \dot{H}_0^t[l]$ with $l > l^*$,

$$\begin{aligned}
\frac{\frac{\Pr^k\left(\mathfrak{h}'|\mathcal{F},\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_t=\rho\right)F^k(\tilde{m}_{i,t}|\mathfrak{h}')}{\sum_{\tilde{\mathfrak{h}}}\Pr^k\left(\tilde{\mathfrak{h}}|\mathcal{F},\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_t=\rho\right)F^k\left(\tilde{m}_{i,t}|\tilde{\mathfrak{h}}\right)}}{\frac{\Pr^k\left(\mathfrak{h}|\mathcal{F},\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_t=\rho\right)F^k(\tilde{m}_{i,t}|\mathfrak{h})}{\sum_{\tilde{\mathfrak{h}}}\Pr^k\left(\tilde{\mathfrak{h}}|\mathcal{F},\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_t=\rho\right)\Pr^k\left(\tilde{m}_{i,t}|\tilde{\mathfrak{h}}\right)}} &\leq 2\frac{\frac{F^k\left(\mathfrak{h}'|c(\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}),\tilde{m}_{i,t}\right)F^k(\tilde{m}_{i,t}|\mathfrak{h}')}{\sum_{\tilde{\mathfrak{h}}}F^k\left(\tilde{\mathfrak{h}}|c(\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}),\tilde{m}_{i,t}\right)F^k\left(\tilde{m}_{i,t}|\tilde{\mathfrak{h}}\right)}}{\frac{F^k\left(\mathfrak{h}|c(\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}),\tilde{m}_{i,t}\right)F^k(\tilde{m}_{i,t}|\mathfrak{h})}{\sum_{\tilde{\mathfrak{h}}}F^k\left(\tilde{\mathfrak{h}}|c(\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}),\tilde{m}_{i,t}\right)F^k\left(\tilde{m}_{i,t}|\tilde{\mathfrak{h}}\right)}} \quad \text{by Lemma 6 and (23)} \\
&= 2\frac{F^k\left(\mathfrak{h}'|c(\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}),\tilde{m}_{i,t}\right)}{F^k\left(\mathfrak{h}|c(\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}),\tilde{m}_{i,t}\right)} \quad \text{by Bayes' rule} \\
&\leq 2\bar{f}_3(k) \quad \text{by (8).}
\end{aligned}$$

Note that, for each $a_{-i,t}$,

$$\begin{aligned}
&\Pr^k\left(a_{-i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\rho,\tilde{m}_t\right) \geq \qquad (29) \\
&\underbrace{f_2(k)}_{\Pr^k\left(\theta_t=\pi|\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\rho,\tilde{m}_t\right)} \times \underbrace{(1-f_2(k))^{N-1}}_{\leq\Pr^k\left(m_j=\star\ \forall j\neq i|\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\rho,\theta_t=\pi,\tilde{m}_t\right)} \times \underbrace{(f_1(k))^N}_{\leq\Pr^k\left(a_{-i,t}|\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\rho,\theta_t=\pi,\tilde{m}_t,m_j=\star\ \forall j\neq i\right)},
\end{aligned}$$

which asymptotically dominates $2\bar{f}_3(k)$. Hence, for any $\mathfrak{h} \in \dot{H}_0^t[l^*]$ and $\mathfrak{h}' \in \dot{H}_0^t[l]$ with $l > l^*$, the likelihood ratio of $\mathfrak{h}'$ to $\mathfrak{h}$ converges to 0 in (27). This implies (28).

We now prove (21) and (22). Consider the following two cases:

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

$$\lim_k F^k\left(\tilde{m}_{-i,t}|\tilde{m}_{i,t},\dot{\mathfrak{h}}_0^t\right) > 0 \text{ and } p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t),a_{i,t},a_{-i,t}=\tilde{m}_{-i,t}\right) > 0. \qquad (30)$$

Substituting this into (28) and using $f_2(k) \to 0$, player i believes $\theta_t = \rho$:

$$\lim_k \Pr^k\left(\theta_t=\rho|\mathcal{F},\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},m_{i,t}=\tilde{m}_{i,t},a_{i,t},s_{i,t+1}\right) = 1. \qquad (31)$$

Moreover, we establish (22) as follows: By Lemma 4 (conditioning on $\mathcal{F}$) and (31) (requiring $\theta_t = \rho$),

$$\begin{aligned}
&\lim_k \frac{\Pr^k\left(\dot{\mathfrak{h}}_0^t,\tilde{m}_{-i,t},a_{-i,t},s_{-i,t+1}|\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},m_{i,t}=\tilde{m}_{i,t},a_{i,t},s_{i,t+1}\right)}{F^k\left(\dot{\mathfrak{h}}_0^t,\tilde{m}_{-i,t},a_{-i,t},s_{-i,t+1}|c(\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}),\tilde{m}_{i,t},a_{i,t},s_{i,t+1}\right)} \\
=\ &\lim_k \frac{\Pr^k\left(\dot{\mathfrak{h}}_0^t,\theta_t=\rho,\tilde{m}_{-i,t},a_{-i,t},s_{-i,t+1}|\mathcal{F},\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},m_{i,t}=\tilde{m}_{i,t},a_{i,t},s_{i,t+1}\right)}{F^k\left(\dot{\mathfrak{h}}_0^t,\tilde{m}_{-i,t},a_{-i,t},s_{-i,t+1}|c(\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}),\tilde{m}_{i,t},a_{i,t},s_{i,t+1}\right)}. \qquad (32)
\end{aligned}$$

As in (27), the numerator equals

$$\lim_k \frac{\left(\begin{array}{c}\Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F},\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_t=\rho,\tilde{m}_{i,t}\right)\\ \times\sum_{\tilde{m}_{-i,t}}\left(F^k\left(\tilde{m}_{-i,t}|\dot{\mathfrak{h}}_0^t,\tilde{m}_{i,t}\right)(1-f_2(k))\,p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t),a_{i,t},a_{-i,t}=\tilde{m}_{-i,t}\right)\right)\end{array}\right)}{\begin{array}{c}\sum_{\tilde{\dot{\mathfrak{h}}}_0^t\in\dot{H}_0^t}\left(\begin{array}{c}\Pr^k\left(\tilde{\dot{\mathfrak{h}}}_0^t|\mathcal{F},\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_t=\rho,\tilde{m}_{i,t}\right)\\ \times\sum_{\tilde{m}_{-i,t}}\left(F^k\left(\tilde{m}_{i,t},\tilde{m}_{-i,t}|\tilde{\dot{\mathfrak{h}}}_0^t\right)(1-f_2(k))\,p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t),a_{i,t},a_{-i,t}=\tilde{m}_{-i,t}\right)\right)\end{array}\right)\\ +\sum_{\tilde{\dot{\mathfrak{h}}}_0^t\in\dot{H}_0^t}\left(\begin{array}{c}\Pr^k\left(\tilde{\dot{\mathfrak{h}}}_0^t|\mathcal{F},\tilde{\dot{\mathfrak{h}}}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_t=\rho,\tilde{m}_{i,t}\right)\\ \times\sum_{\tilde{m}_{-i,t}}\left(\begin{array}{c}F^k\left(\tilde{m}_{-i,t}|\tilde{\dot{\mathfrak{h}}}_0^t,\tilde{m}_{i,t}\right)(f_2(k))^2\\ \sum_{a_{-i,t}}\Pr^k\left(a_{-i,t}|\mathcal{F},\tilde{\dot{\mathfrak{h}}}_0^t,\mathfrak{h}_i^t,\theta_t=\pi,\tilde{m}_t\right)p\left(s_{i,t+1}|\mathring{c}(\tilde{\dot{\mathfrak{h}}}_0^t),a_t\right)\end{array}\right)\end{array}\right)\end{array}}.$$

By (30) and $f_2(k)\to 0$, the second term of the denominator drops out in the limit. Finally, by (23), $f_2(k)\to 0$, and Lemma 6, the numerator of (32) equals

$$\lim_k \frac{\left(\begin{array}{c}F^k\left(\dot{\mathfrak{h}}_0^t|c(\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}),\tilde{m}_{i,t}\right)\\ \times\sum_{\tilde{m}_{-i,t}}\left(F^k\left(\tilde{m}_{-i,t}|\dot{\mathfrak{h}}_0^t,\tilde{m}_{i,t}\right)p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t),a_{i,t},a_{-i,t}=\tilde{m}_{-i,t}\right)\right)\end{array}\right)}{\sum_{\tilde{\dot{\mathfrak{h}}}_0^t\in\dot{H}_0^t}\left(\begin{array}{c}F^k\left(\tilde{\dot{\mathfrak{h}}}_0^t|c(\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}),\tilde{m}_{i,t}\right)\\ \times\sum_{\tilde{m}_{-i,t}}F^k\left(\tilde{m}_{i,t},\tilde{m}_{-i,t}|\tilde{\dot{\mathfrak{h}}}_0^t\right)p\left(s_{i,t+1}|\mathring{c}(\tilde{\dot{\mathfrak{h}}}_0^t),a_{i,t},a_{-i,t}=\tilde{m}_{-i,t}\right)\end{array}\right)}.$$

As the fraction inside the limit equals

$$F^k\left(\dot{\mathfrak{h}}_0^t,\tilde{m}_{-i,t},a_{-i,t},s_{-i,t+1}|c(\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}),\tilde{m}_{i,t},a_{i,t},s_{i,t+1}\right),$$

(32) implies (22).

Case 2: There do not exist $\dot{\mathfrak{h}}_0^t\in\dot{H}_0^t[l^*]$ and $\tilde{m}_{-i,t}$ such that (30) holds. In this case, we claim that player i believes $\theta_t=\pi$:

$$\lim_k \Pr^k\left(\theta_t=\rho|\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},m_{i,t}=\tilde{m}_{i,t},a_{i,t},s_{i,t+1}\right)=0.$$

To see why, for each $\dot{\mathfrak{h}}_0^t \in \dot{H}_0^t[l^*]$, $s_{i,t+1}$ arises with probability at most $\bar{f}_3(k)$ given $\theta_t = \rho$: the numerator of (28) (and the first term in the denominator) sums the terms

$$\begin{aligned}
& \Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right) \\
& \times \sum_{\tilde{m}_{-i,t}} \left(F^k\left(\tilde{m}_{-i,t}|\tilde{m}_{i,t}, \dot{\mathfrak{h}}_0^t\right)(1 - f_2(k))\, p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t), a_{i,t}, a_{-i,t} = \tilde{m}_{-i,t}\right)\right) \\
\leq\ & \Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right)\bar{f}_3(k),
\end{aligned}$$

since $\lim_k F^k\left(\tilde{m}_{-i,t}|\tilde{m}_{i,t}, \dot{\mathfrak{h}}_0^t\right) = 0$ implies $F^k\left(\tilde{m}_{-i,t}|\tilde{m}_{i,t}, \dot{\mathfrak{h}}_0^t\right) \leq \bar{f}_3(k)$ by (8).

On the other hand, $s_{i,t+1}$ arises with higher probability given $\theta_t = \pi$: the second term in the denominator of (28) sums the terms

$$\begin{aligned}
& \Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right) \\
& \times \textstyle\sum_{\tilde{m}_{-i,t}} \begin{pmatrix} F^k\left(\tilde{m}_{-i,t}|\tilde{m}_{i,t}, \dot{\mathfrak{h}}_0^t\right)(f_2(k))^2 \\ \sum_{a_{-i,t}} \Pr^k\left(a_{-i,t}|\mathcal{F}, \dot{\mathfrak{h}}_0^t, \mathfrak{h}_i^t, \theta_t = \pi, \tilde{m}_t\right) p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t), a_t\right) \end{pmatrix} \\
\geq\ & \begin{array}{c} \Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right)(f_2(k))^2(1 - f_2(k))^{N-1}(f_1(k))^N \\ \times \min_{s_{i,t+1}, a_t : p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t), a_t\right) > 0} p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t), a_t\right) \end{array},
\end{aligned}$$

by (29), which asymptotically dominates $\Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right)\bar{f}_3(k)$. Hence, (28) equals 0.

Proof of Lemma 8. Given $\mathcal{F}$ and $m_{i,t} = \tilde{m}_{i,t}$, we have $m_{i,t} = \tilde{m}_{i,t} = a_{i,t}$. Since $m_{i,t} = \tilde{m}_{i,t}$ implies $\hat{\theta}_t = \rho$,

$$\begin{aligned}
& \lim_k \Pr^k\left(\theta_t = \rho|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, m_{i,t} = \tilde{m}_{i,t}, s_{i,t+1}\right) \\
=\ & \lim_k \Pr^k\left(\theta_t = \rho|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, m_{i,t} = \tilde{m}_{i,t}, s_{i,t+1}\right) \\
=\ & \lim_k \frac{\Pr^k\left(\theta_t = \rho, m_{i,t} = \tilde{m}_{i,t} = a_{i,t}, s_{i,t+1}|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right)}{\begin{pmatrix} \Pr^k\left(\theta_t = \rho, m_{i,t} = \tilde{m}_{i,t} = a_{i,t}, s_{i,t+1}|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right) \\ + \Pr^k\left(\theta_t = \pi, m_{i,t} = \tilde{m}_{i,t} = a_{i,t}, s_{i,t+1}|\mathcal{F}, \mathfrak{h}_i^t, (\tilde{m}_{i,\tau})_{\tau=1}^{t-1}, \hat{\theta}_t = \rho, \tilde{m}_{i,t}\right) \end{pmatrix}}.
\end{aligned}$$

By Bayes' rule, the numerator (and the first term in the denominator) equals

$$\sum_{\dot{\mathfrak{h}}_0^t}\left(\begin{array}{c}\Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F},\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_t=\rho,\tilde{m}_{i,t}\right)\\ \times\sum_{\tilde{m}_{-i,t}}\left(\begin{array}{c}\Pr^k\left(\tilde{m}_{-i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\rho,\tilde{m}_{i,t}\right)\\ \times\Pr^k\left(\theta_t=\rho,m_{i,t}=\tilde{m}_{i,t}=a_{i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\rho,\tilde{m}_t\right)\\ \times\sum_{a_{-i,t}}\left(\begin{array}{c}\Pr^k\left(a_{-i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\theta_t=\rho,\tilde{m}_t,m_{i,t}=\tilde{m}_{i,t}=a_{i,t}\right)\\ \times\Pr^k\left(s_{i,t+1}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\theta_t=\rho,\tilde{m}_t,m_{i,t}=\tilde{m}_{i,t}=a_{i,t},a_t\right)\end{array}\right)\end{array}\right)\end{array}\right).$$

(We omit conditioning on $(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}$ given $\dot{\mathfrak{h}}_0^t$, since $\dot{\mathfrak{h}}_0^t$ determines $(\tilde{m}_{i,\tau})_{\tau=1}^{t-1}$.) For each $\dot{\mathfrak{h}}_0^t$, $\mathfrak{h}_i^t$, $\tilde{m}_{-i,t}$, $a_{-i,t}$, and $s_{i,t+1}$, we have

$$\begin{aligned}\Pr^k\left(\tilde{m}_{-i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\rho,\tilde{m}_{i,t}\right) &= F^k\left(\tilde{m}_{-i,t}|\dot{\mathfrak{h}}_0^t,\tilde{m}_{i,t}\right) \text{ by (10)},\\ \Pr^k\left(\theta_t=\rho,m_{i,t}=\tilde{m}_{i,t}=a_{i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\rho,\tilde{m}_t\right) &= 1-f_2(k) \text{ by (12) and (10)},\\ \Pr^k\left(a_{-i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\theta_t=\rho,\tilde{m}_t,m_{i,t}=\tilde{m}_{i,t}=a_{i,t}\right) &= 1 \text{ if } a_{-i,t}=\tilde{m}_{-i,t},\\ \Pr^k\left(a_{-i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\theta_t=\rho,\tilde{m}_t,m_{i,t}=\tilde{m}_{i,t}=a_{i,t}\right) &= 0 \text{ if } a_{-i,t}\neq\tilde{m}_{-i,t}\\ (\text{as } \tilde{m}_{-i,t} &\in A_{-i,t} \text{ given } \theta_t=\rho),\\ \Pr^k\left(s_{i,t+1}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\theta_t=\rho,\tilde{m}_t,m_{i,t}=\tilde{m}_{i,t}=a_{i,t},a_t\right) &= p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t),a_t\right).\end{aligned}$$

Hence,

$$\begin{aligned}&\Pr^k\left(\theta_t=\rho,m_{i,t}=\tilde{m}_{i,t},a_{i,t},s_{i,t+1}|\mathcal{F},\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\tilde{m}_{i,t},\hat{\theta}_t=\rho\right) \qquad (33)\\ &= \sum_{\dot{\mathfrak{h}}_0^t}\left(\begin{array}{c}\Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F},\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_t=\rho,\tilde{m}_{i,t}\right)\\ \times\sum_{\tilde{m}_{-i,t}}F^k\left(\tilde{m}_{-i,t}|\dot{\mathfrak{h}}_0^t,\tilde{m}_{i,t}\right)(1-f_2(k))\,p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t),a_{i,t},a_{-i,t}=\tilde{m}_{-i,t}\right)\end{array}\right).\end{aligned}$$

On the other hand, the second term in the denominator equals

$$\sum_{\dot{\mathfrak{h}}_0^t}\left(\begin{array}{c}\Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F},\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_t=\rho,\tilde{m}_{i,t}\right)\\ \times\sum_{\tilde{m}_{-i,t}}\left(\begin{array}{c}\Pr^k\left(\tilde{m}_{-i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\rho,\tilde{m}_{i,t}\right)\\ \times\Pr^k\left(\theta_t=\pi,m_{i,t}=a_{i,t}=\tilde{m}_{i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\rho,\tilde{m}_t\right)\\ \times\sum_{a_{-i,t}}\left(\begin{array}{c}\Pr^k\left(a_{-i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\theta_t=\pi,\tilde{m}_t,m_{i,t}=\tilde{m}_{i,t}=a_{i,t}\right)\\ \times\Pr^k\left(s_{i,t+1}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\theta_t=\pi,\tilde{m}_t,m_{i,t}=\tilde{m}_{i,t}=a_{i,t},a_t\right)\end{array}\right)\end{array}\right)\end{array}\right).$$

For each $\dot{\mathfrak{h}}_0^t$, $\mathfrak{h}_i^t$, $\tilde{m}_{-i,t}$, $a_{-i,t}$, and $s_{i,t+1}$, we have

$$
\begin{aligned}
\Pr^k\left(\tilde{m}_{-i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\rho,\tilde{m}_{i,t}\right) &= F^k\left(\tilde{m}_{-i,t}|\dot{\mathfrak{h}}_0^t,\tilde{m}_{i,t}\right), \\
\Pr^k\left(\theta_t=\pi,m_{i,t}=\tilde{m}_{i,t}=a_{i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\hat{\theta}_t=\rho,\tilde{m}_t\right) &= (f_2(k))^2 \quad \text{by (12) and (11)}, \\
\Pr^k\left(a_{-i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\theta_t=\pi,\tilde{m}_t,m_{i,t}=\tilde{m}_{i,t}=a_{i,t}\right) &= \Pr^k\left(a_{-i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\theta_t=\pi,\tilde{m}_t\right), \\
\Pr^k\left(s_{i,t+1}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\theta_t=\pi,\tilde{m}_t,m_{i,t}=\tilde{m}_{i,t}=a_{i,t},a_t\right) &= p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t),a_t\right).
\end{aligned}
$$

Hence,

$$
\begin{aligned}
&\Pr^k\left(\theta_t=\pi,m_{i,t}=\tilde{m}_{i,t},a_{i,t},s_{i,t+1}|\mathcal{F},\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\tilde{m}_{i,t},\hat{\theta}_t=\rho\right) \qquad (34) \\
&= \sum_{\dot{\mathfrak{h}}_0^t}\left(\begin{array}{c}
\Pr^k\left(\dot{\mathfrak{h}}_0^t|\mathcal{F},\mathfrak{h}_i^t,(\tilde{m}_{i,\tau})_{\tau=1}^{t-1},\hat{\theta}_t=\rho,\tilde{m}_{i,t}\right) \\
\times\sum_{\tilde{m}_{-i,t}}\left(\begin{array}{c}
F^k\left(\tilde{m}_{-i,t}|\dot{\mathfrak{h}}_0^t,\tilde{m}_{i,t}\right)(f_2(k))^2 \\
\times\sum_{a_{-i,t}}\Pr^k\left(a_{-i,t}|\mathcal{F},\dot{\mathfrak{h}}_0^t,\mathfrak{h}_i^t,\theta_t=\pi,\tilde{m}_t\right)p\left(s_{i,t+1}|\mathring{c}(\dot{\mathfrak{h}}_0^t),a_t\right)
\end{array}\right)
\end{array}\right).
\end{aligned}
$$

Combining (33) and (34) yields (27). ■

13 Proof of Proposition 6

The proof is by example. The idea is to extend Example 2 by introducing an extra player and period in such a way as to ensure that actions (C_1, C_2) must be recommended with positive probability on path. This implies that recommendations to play these actions cannot be used to substitute for the extra "free pass" message, so the message set must be expanded.

There are four players and three periods. Player 1's payoff is given by

$$\frac{9}{10}\times\mathbf{1}_{\{a_1=B_1\}}-\mathbf{1}_{\{a_2=C_2\}}\mathbf{1}_{\{a_3=C_3\}}\mathbf{1}_{\{a_4=B_4\}}.$$

The other players' payoffs are the sum of the per-period payoffs described below.

Period 1: No signal is observed. Player 1 takes an action $a_1 \in \{A_1, B_2\}$. The other players' period payoffs are 0.

Period 2: The mediator observes player 1's action a_1. Players 2 and 3 take actions $(a_2, a_3) \in \{A_2, B_2, C_2\} \times \{A_3, B_3, C_3\}$. They receive period payoffs as follows:

- If $a_1 = A_1$ then player 2 and 3's payoffs are given by

	A_3	B_3	C_3
A_2	$0, 0$	$1, 1$	$0, -3$
B_2	$1, 1$	$0, 0$	$0, -3$
C_2	$-3, 0$	$-3, 0$	$-3, -3$

- If $a_1 = B_1$ then player 2 and 3's payoffs are given by

	A_3	B_3	C_3
A_2	$\frac{1}{10}, \frac{1}{10}$	$0, 0$	$0, -3$
B_2	$0, 0$	$0, 0$	$0, -3$
C_2	$-3, 0$	$-3, 0$	$\frac{1}{10}, \frac{1}{10}$

Period 3: Player 4 and the mediator observe a signal $s_4 \in \{0, 1\}$, where $s_4 = 0$ if and only if $(a_2, a_3) \in \{(A_2, A_3), (B_2, B_3)\}$. Player 4 takes an action $a_4 \in \{A_4, B_4\}$ and receives payoff

	$a_1 = B_1$ or $(a_2, a_3) \neq (C_2, C_3)$	$a_1 = A_1$ and $(a_2, a_3) = (C_2, C_3)$
A_4	1	0
B_4	0	1

Finally, each player $i = 2, 3$ receives period payoff $-2 \times 1_{\{a_1 = A_1\}} 1_{\{a_4 = B_4\}} - 1_{\{a_1 = B_1\}} 1_{\{a_i = C_i\}} 1_{\{a_4 = A_4\}}$.

Consider the target outcome distribution where (i) $\frac{1}{2}A_1 + \frac{1}{2}B_1$ is played in period 1, (ii) when A_1 is played in period 1, $\frac{1}{2}(A_2, A_3) + \frac{1}{2}(B_2, B_3)$ is played in period 2, (iii) when B_1 is played in period 1, (A_2, A_3) is played in period 2, and (iv) A_4 is played in period 4.

Claim 1 *The target outcome distribution is a SCE but it is not implementable in PSE with message set $M_i = A_i$ for all $i \neq 0$.*

(Note that we have omitted the time subscript on M_i and A_i since each player moves only once.)

To show the target distribution is implementable as a canonical CPPBE, we first describe the mediator's strategy:

In period 1, the mediator recommends A_1 and B_1 with equal probability. Then,

1. If $m_1 = a_1 = A_1$, the mediator recommends (A_1, A_2) and (B_1, B_2) with equal probability in period 2. In period 3, the mediator recommends A_4 if $s_4 = 0$ and recommends B_4 if $s_4 = 1$.

2. If $m_1 = A_1$ but $a_1 = B_1$, the mediator recommends (C_2, C_3) in period 2. In period 3, the mediator recommends A_4 if $s_4 = 0$ and recommends B_4 if $s_4 = 1$.

3. If $m_1 = B_1$, the mediator recommends (A_2, A_3) in period 2 and A_4 in period 3.

The critical belief to specify is player 4's after observing $s_4 = 1$ and receiving recommendation $m_4 = B_4$. In this case, we specify that player 4 assigns probability 1 to the event $m_1 = a_1 = A_1 \cap (a_2, a_3) = (C_2, C_3)$.

We verify incentive compatibility. For player 1, this follows as (i) if $m_1 = A_1$, a deviation to B_1 is punished, and (ii) if $m_1 = B_1$, $a_1 = B_1$ is myopically optimal and others' play is independent of a_1.

For player $i \in \{2, 3\}$, given an on-path recommendation $m_i \in \{A_i, B_i\}$, incentive compatibility follows as either (i) $a_1 = A_1$ and a deviation to $a_i = \{A, B\} \setminus \{m_i\}$ is punished, or

(ii) $a_1 = B_1$ and $m_i = A_i$ is optimal. Given recommendation $m_i = C_i$, player i believes that $m_1 = A_1$ but $a_1 = B_1$, so player 4 will play B_4 and hence C_i is optimal.

For player 4, following recommendation $m_4 = A_4$ is optimal since A_4 is recommended only at information sets where $a_1 = B_1$ or $(a_2, a_3) \neq (C_2, C_3)$. Finally, following recommendation $m_4 = B_4$ is optimal since B_4 is recommended only when $s_4 = 1$, in which case player 4 assigns probability 1 to $a_1 = A_1 \cap (a_2, a_3) = (C_2, C_3)$.

We now show the target distribution is not implementable in PSE with message set $M_i = A_i$ for all i. Without loss of generality, at each information set for the mediator, let $a_i \in M_i$ denote the message $m_i \in M_i$ that induces action $a_i \in A_i$ with highest probability (given the mediator's beliefs). Note that, in order to implement the target distribution, actions C_2, C_3, and B_4 must be played with positive probability following the event $m_1 = A_1 \cap a_1 = B_1$. (This is necessary to deter deviations by player 1.) As the other actions are played with positive probability under the target distribution, it follows that, for each player i and every message $m_i \in M_i$, action $a_i = m_i$ must be played with positive probability when message m_i is received. Moreover, the standard static revelation principle implies that this probability can be taken to be 1 for player 4 (the last player to move in the game), and the fact that players 2 and 3's actions are perfectly correlated when $a_1 = A_1$ implies that for $i = 2, 3$ this probability can also be taken to be 1 for $m_i \in \{A_i, B_i\}$. Finally, to deter deviations by player 1, each player $i = 2, 3$ must play $a_i = C_i$ with probability at least $9/10$ after receiving message $m_i = C_i$.

Now, when player 2 receives message $m_2 = C_2$ and plays $a_2 = C_2$, she must assign probability at least $10/11$ to the event $a_1 = B_1 \cap a_3 = C_3 \cap m_4 = B_4$. Since $s_4 = 1$ whenever $(a_2, a_3) = (C_2, C_3)$, this implies

$$\Pr\left(a_1 = B_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_2 = a_2 = C_2\right) \geq \frac{10}{11}.$$

Hence,

$$\begin{aligned}
& \Pr\left(a_1 = B_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_2 = C_2\right) \\
= \; & \Pr\left(a_2 = C_2 | m_2 = C_2\right) \Pr\left(a_1 = B_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_2 = a_2 = C_2\right) \\
\geq \; & \frac{9}{10}\frac{10}{11} = \frac{9}{11},
\end{aligned}$$

and

$$\Pr\left(a_1 = A_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_2 = C_2\right) \leq \frac{2}{11}.$$

Similarly,

$$\begin{aligned}
\Pr\left(a_1 = B_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_3 = C_3\right) &\geq \frac{9}{11}, \\
\Pr\left(a_1 = A_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_3 = C_3\right) &\leq \frac{2}{11}.
\end{aligned}$$

On the other hand, when player 4 observes $s_4 = 1$ and receives message $m_4 = B_4$, he must assign probability at least 1/2 to the event $a_1 = A_1 \cap (a_2, a_3) = (C_2, C_3)$: in particular,

$$
\begin{aligned}
\frac{1}{2} &\leq \frac{\begin{pmatrix} \Pr(m_2 = C_2)\Pr(a_1 = A_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_2 = C_2) \\ + \Pr(m_3 = C_3)\Pr(a_1 = A_1, m_2 \neq C_2, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_3 = C_3) \\ + \sum_{m_2 \neq C_2, m_3 \neq C_3} \Pr(m_2, m_3)\Pr(a_1 = A_1, (a_2, a_3) = (C_2, C_3)_3, m_4 = B_4, s_4 = 1 | m_2, m_3) \end{pmatrix}}{\begin{pmatrix} \Pr(m_2 = C_2)\sum_{a_1, a_2, a_3} \Pr(a_1, a_2, a_3, m_4 = B_4, s_4 = 1 | m_2 = C_2) \\ + \sum_{m_2 \neq C_2} \Pr(m_2, m_3)\sum_{a_1, a_2, a_3} \Pr(a_1, a_2, a_3, m_4 = B_4, s_4 = 1 | m_2, m_3) \end{pmatrix}} \\
&\leq \frac{\begin{pmatrix} \Pr(m_2 = C_2)\Pr(a_1 = A_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_2 = C_2) \\ + \Pr(m_3 = C_3)\Pr(a_1 = A_1, m_2 \neq C_2, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_3 = C_3) \\ + \sum_{m_2 \neq C_2, m_3 \neq C_3} \Pr(m_2, m_3)\Pr(a_1 = A_1, (a_2, a_3) = (C_2, C_3)_3, m_4 = B_4, s_4 = 1 | m_2, m_3) \end{pmatrix}}{\begin{pmatrix} \Pr(m_2 = C_2)\sum_{a_1, a_2, a_3} \Pr(a_1, a_2, a_3, m_4 = B_4, s_4 = 1 | m_2 = C_2) \\ + \sum_{m_2 \neq C_2, m_3 \neq C_3} \Pr(m_2, m_3)\sum_{a_1, a_2, a_3} \Pr(a_1, a_2, a_3, m_4 = B_4, s_4 = 1 | m_2, m_3) \end{pmatrix}}.
\end{aligned}
$$

Note that

$$
\begin{aligned}
&\frac{\begin{matrix} \Pr(m_2 = C_2)\Pr(a_1 = A_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_2 = C_2) \\ + \Pr(m_3 = C_3)\Pr(m_2 \neq C_2, a_1 = A_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_3 = C_3) \end{matrix}}{\Pr(m_2 = C_2)\sum_{a_1, a_2, a_3} \Pr(a_1, a_2, a_3, m_4 = B_4, s_4 = 1 | m_2 = C_2)} \\
\leq\ &\frac{\Pr(m_2 = C_2)\Pr(a_1 = A_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_2 = C_2)}{\Pr(m_2 = C_2)\Pr(a_1 = B_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_2 = C_2)} \\
&+ \frac{\Pr(m_3 = C_3)\Pr(a_1 = A_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_3 = C_3)}{\Pr(m_3 = C_3)\Pr(a_1 = B_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_3 = C_3)} \\
\leq\ &\frac{2/11}{9/11} + \frac{2/11}{9/11} = \frac{4}{9}.
\end{aligned}
$$

Now, for all positive numbers a, b, c, d, $(a + c) / (b + d) < (a/b) + (c/d)$. Hence, the preceding inequalities imply

$$
\begin{aligned}
\frac{1}{2} - \frac{4}{9} &\leq \frac{\sum_{m_2 \neq C_2, m_3 \neq C_3} \Pr(m_2, m_3)\Pr(a_1 = A_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_2, m_3)}{\sum_{m_2 \neq C_2, m_3 \neq C_3} \Pr(m_2, m_3)\sum_{a_1, a_2, a_3} \Pr(a_1, a_2, a_3, m_4 = B_4, s_4 = 1 | m_2, m_3)} \\
&\leq \frac{\sum_{m_2 \neq C_2, m_3 \neq C_3} \Pr(m_2, m_3)\Pr(a_1 = A_1, (a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_2, m_3)}{\sum_{m_2 \neq C_2, m_3 \neq C_3} \Pr(m_2, m_3)\sum_{a_2, a_3} \Pr(a_1 = A_1, a_2, m_4 = B_4, s_4 = 1 | m_2, m_3)} \\
&= \frac{\sum_{m_2 \neq C_2, m_3 \neq C_3} \Pr(m_2, m_3 | a_1 = A_1)\Pr((a_2, a_3) = (C_2, C_3), m_4 = B_4, s_4 = 1 | m_2, m_3, a_1 = A_1)}{\sum_{m_2 \neq C_2, m_3 \neq C_3} \Pr(m_2, m_3 | a_1 = A_1)\sum_{a_2, a_3} \Pr(a_2, a_3, m_4 = B_4, s_4 = 1 | m_2, m_3, a_1 = A_1)}.
\end{aligned}
$$

But this equals 0 by same argument as in the proof of Proposition 4. Contradiction.