The Revelation Principle in Multistage Games*

Takuo Sugaya and Alexander Wolitzky

Stanford GSB and MIT

July 22, 2019

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 an incentive-compatible direct 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, or equivalently if and only if it is a sequential communication equilibrium outcome as defined by Myerson [Myerson, R.B. (1986), "Multistage Games with Communication," *Econometrica*, 54, 323-358].

*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, Bocconi, Collegio Carlo Alberto, Cowles, Harvard-MIT, LSE, NYU, Princeton, Queen Mary, Stony Brook, UCL, UCLA, and UCSD; and the anonymous referees. Wolitzky acknowledges financial support from the NSF and the Sloan Foundation.

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 direct 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 result was developed throughout the 1970s, 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. The importance of the communication revelation principle 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 (NE): this approach leads to the concept of *communication equilibrium* introduced by Forges (1986). But NE 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 (SE) 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 SE in multistage games. We will show that the communication revelation principle holds for PBE but fails for SE. Nonetheless, our main result establishes that the implementation side of the revelation principle remains valid even for SE.

[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).

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 direct communication: 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 direct communication 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*, which are actions that cannot be motivated by any belief consistent with a player's own information and the presumption that her opponents will avoid codominated actions in the future.[2] Myerson's paper also proves 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 direct communication? Second, is there an equivalence between implementation in CPPBE and implementation in SE, 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 for CPPBE, and we prove that implementation in CPPBE is equivalent to implementation in SE. The first of these results may be viewed as a formalization of ideas implicit in Myerson.

[2] We review Myerson's characterization and the definition of codomination in Section 2.5.

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

The second—which we refer to as the *implementation revelation principle* for SE—is quite subtle and (to us) unexpected. Indeed, we show that the *communication* RP fails for SE.[4] We thus show

$$SCE = CPPBE\,(C^*) = \bigcup_{C\in\mathcal{C}} CPPBE\,(C) = \bigcup_{C\in\mathcal{C}} SE\,(C) = SE\,(C^{**}) \supsetneq SE\,(C^*)\,, \qquad (1)$$

where "solution concept(C)" means the set of equilibrium outcome distributions with communication system C, C^* is the direct communication system, $\mathcal{C}$ is the set of "all" communication systems, and C^{**} denotes a particular "quasi-direct" communication system, which we describe later on.

Our results have a concise and practical message for applied dynamic mechanism design: to calculate the set of outcomes implementable in *sequential* equilibrium by *any* communication system, simply calculate the set of outcomes implementable in *Nash* equilibrium *without codominated actions*, using *direct* communication. These two sets are always the same, even though actually implementing some outcomes as sequential equilibria might require using the richer communication system C^{**}.

Let us preview the intuition for our main result: the implementation RP for SE, or $SCE \subseteq \bigcup_{C\in\mathcal{C}} SE\,(C)$.[5] By Myerson's result, a SCE is a communication equilibrium where all actions that are ever played (on or off path) can be motivated by some belief consistent with a player's own information. Such a belief can be generated in accordance with Kreps-Wilson consistency by specifying that all players tremble with substantial probability (along a sequence of strategy profiles converging to the equilibrium) and then honestly report their trembles to the mediator, and the mediator appropriately conditions his recommendations on the reported actions. An obstacle to this construction is that a player who trembles

[4] 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). More relevant are Dhillon and Mertens (1996; Example 1), who show that the communication RP fails for the solution concept of "perfect correlated equilibrium," which describes outcomes that can be implemented in trembling-hand perfect equilibrium; and Gerardi and Myerson (2007), who show that both the communication and implementation RPs fail for SE when the mediator cannot tremble (see Section 5.2).

[5] The reverse inclusion is a corollary of the (simpler) communication RP for CPPBE: $\bigcup_{C\in\mathcal{C}} SE\,(C) \subseteq \bigcup_{C\in\mathcal{C}} CPPBE\,(C)$ because SE is a refinement CPPBE, and $\bigcup_{C\in\mathcal{C}} CPPBE\,(C) = CPPBE\,(C^*) = SCE$ by the communication RP for CPPBE.

to an action for which she must be punished in equilibrium will not honestly report her deviation. To circumvent this problem, the mediator may (with probability converging to 0) promise in advance that he will disregard a player's report almost-surely. However, to afford the mediator the ability to make such an advance promise, the communication system must be enriched with an extra message. This augmentation results in the quasi-direct communication system C^{**}. Note that the mediator's "promise to ignore reports" is made with equilibrium probability 0, so our construction is "canonical on path."[6]

The failure of the communication RP for SE can be avoided if the game satisfies appropriate full support conditions. We clarify what conditions are required. In particular, the communication RP is always valid in single-agent settings, as well as in settings with adverse selection but no moral hazard (a class of games which encompasses much of the recent literature on dynamic mechanism design).

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 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 equilibrium (PBE) as our solution concept."

Our results provide a foundation for this approach, while also showing that Nash and PBE are outcome-equivalent in pure adverse selection settings like this one.[7]

[6] In our definition of sequential equilibrium, therefore, it is important to allow the mediator to tremble. See Section 5.2 for a discussion of what happens if the mediator cannot tremble.

[7] 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 discuss this point in Section 5.2.

For variants of other well-known models, naïvely applying the RP can produce serious mistakes. For example, Kremer, Mansour, and Perry (2014) consider a setting where an information designer wants to facilitate social learning by encouraging players to explore a risky option; a leading example is persuading commuters to explore different routes.[8] The authors claim that the communication RP is established for their setting by Myerson. In fact, it is valid only because of special features of the model, especially the absence of payoff externalities among the players. In a variant of their model where commuters choose routes repeatedly and face congestion externalities, the communication RP for SE is not guaranteed. However, our theorem implies that the implementation RP remains valid.

More generally, our simple positive results for single-agent settings and settings with adverse selection but no moral hazard imply that the subtleties at the heart of our paper are most relevant for multi-agent, multi-stage games with moral hazard: that is, *multi-agent dynamic information design*. Papers on this topic include Gershkov and Szentes (2009), Aoyagi (2010), Halac, Kartik, Liu (2014), Kremer, Mansour, and Perry (2014), Che and Hörner (2017), Sugaya and Wolitzky (2017), Ely (2017), Doval and Ely (2019), and Makris and Renou (2019). We hope our work will allow this emerging literature to use the RP with more confidence. To this end, we provide a compact summary of our results at the end of the paper.

1.1 Example

Before presenting the model, we give an example that illustrates how letting the mediator make advance promises to disregard players' reports can expand the set of implementable outcomes. This phenomenon is the basic explanation both for why (for SE) the communication RP fails and the implementation RP holds.

There are two players (in addition to the mediator) and three periods.

In period 1, player 1 takes an action $a_1 \in \{A, B, C\}$.

In period 2, player 1 observes a signal $\theta \in \{n, p\}$, with each realization equally likely. Then, the mediator ("player 0") takes an action $a_0 \in \{A, B\}$.

[8] See also Che and Hörner (2017).

In period 3, the mediator and player 2 observe a common signal $s \in \{0,1\}$, where $s = 1$ if and only if $a_0 \neq a_1$. Then, player 2 takes an action $a_2 \in \{N, P\}$ ("Not punish," "Punish").

Player 1's payoff equals $1_{\{a_0 \neq a_1 \wedge a_1 \neq C\}} - 1_{\{a_2 = P\}} - 3 \times 1_{\{a_1 = C\}}$, and player 2's payoff equals $-1_{\{(a_1,\theta) \neq (C,p)\}} 1_{\{a_2 = P\}}$. In particular, player 1 wants to mismatch her action with the mediator's action; action C is strictly dominated for player 1; and player 2 is willing to punish player 1 iff $a_1 = C$ and $\theta = p$.

Consider the outcome distribution $\frac{1}{2}(A, A, N) + \frac{1}{2}(B, B, N)$. It is trivial to construct a canonical NE (i.e., a NE with direct communication where in equilibrium players report their information honestly and obey the mediator's recommendations) that implements this outcome: the mediator sends message/recommendation $m_1 = A$ and $m_1 = B$ with equal probability, plays $a_0 = m_1$, and recommends $m_2 = N$ if $s = 0$ and $m_2 = P$ if $s = 1$; meanwhile, players are honest and obedient. Moreover, this NE is sequential iff player 2 believes with probability 1 that $(a_1, \theta) = (C, p)$ when $s = 1$ and $m_2 = P$. Thus, $\frac{1}{2}(A, A, N) + \frac{1}{2}(B, B, N)$ is implementable in sequential equilibrium iff this belief is consistent.

Myerson shows that any NE outcome where players do not take codominated actions at any (possibly off-path) history is a SCE outcome. Our main result is that exactly the same set of outcomes is implementable in SE, but we may need to allow non-canonical equilibria. In this example, the action $a_2 = P$ is not codominated at the history following following signal $s = 1$, because $a_2 = P$ is an optimal action for player 2 if $(a_1, \theta) = (C, p)$. Our main result thus implies that $\frac{1}{2}(A, A, N) + \frac{1}{2}(B, B, N)$ is implementable in some SE. We now explain intuitively why $\frac{1}{2}(A, A, N) + \frac{1}{2}(B, B, N)$ is not implementable in any canonical SE, but is implementable in a non-canonical SE.[9]

Impossibility for canonical SE: In a canonical SE, player 1 obeys all recommendations from the mediator. Since $a_1 = C$ is strictly dominated, this implies that the mediator can only ever recommend $m_1 \in \{A, B\}$. If player 1 mistakenly plays the strictly dominated action C after such a recommendation, she will subsequently (for each possible realization of θ) make whatever report $\left(\hat{a}_1, \hat{\theta}\right)$ minimizes the probability that $m_2 = P$. Since $s = 1$ whenever $a_1 = C$, Bayes' rule then implies that $\Pr\left((a_1, \theta) = (C, p) \,|\, s = 1, m_2 = P\right) = \frac{1}{2}$. Hence, player 2 will not follow the recommendation $m_2 = P$ when $s = 1$, so the desired

[9] For the details, see the proof of Proposition 5.

outcome is not implementable in a canonical SE.

Possibility for non-canonical SE: Why does enriching the communication system overturn this negative result? Suppose the mediator can tremble by giving a "free pass" to player 1 in period 1. If player 1 gets a free pass in period 1, the mediator will always recommend $m_2 = N$, barring another mediator tremble. Now, when player 2 is recommended $m_2 = P$, he can believe that the mediator trembled by giving player 1 a free pass in period 1, player 1 trembled to $a_1 = C$, player 1 honestly reported $\left(\hat{a}_1, \hat{\theta}\right) = (C, p)$, and the mediator trembled again by recommending $m_2 = P$. This new possibility can rationalize player 2's belief that $(a_1, \theta) = (C, p)$.

More precisely, consider the following sequence of strategy profiles, indexed by $k \in \mathbb{N}$:

Mediator's strategy: In period 1, the mediator recommends A and B with equal probability, while trembling to a third message, "$\star$" (the "free pass"), with probability $\frac{1}{k}$. In period 2, if $m_1 \in \{A, B\}$, the mediator plays $m_0 = m_1$; if $m_1 = \star$, he plays A and B with probability $\frac{1}{2}$ each. In period 3, if $m_1 \in \{A, B\}$, the mediator recommends $m_2 = N$ if $s = 0$ and $m_2 = P$ if $s = 1$; if $m_1 = \star$, with probability $1 - \frac{1}{k}$ he recommends $m_2 = N$ (regardless of $\left(\hat{a}_1, \hat{\theta}\right)$ and s), and with probability $\frac{1}{k}$ he recommends $m_2 = P$ if $\left(\hat{a}_1, \hat{\theta}\right) = (C, p)$ and $m_2 = N$ otherwise.

Players' strategies: If $m_1 \in \{A, B\}$, player 1 takes $a_1 = m_1$ and trembles to each other action with probability $\frac{1}{k^4}$; if $m_1 = \star$, she plays A and B with probability $\frac{1}{2}$ each, while trembling to C with probability $\frac{1}{k}$. Player 1 always reports her action and signal honestly. Player 2 always takes $a_2 = m_2$.

Note that honesty is always optimal for player 1: if $m_1 \in \{A, B\}$, then any deviation from $a_1 = m_1$ leads to $a_2 = P$ with probability 1 regardless of player 1's report; while if $m_1 = \star$, then $a_2 = N$ with probability 1 regardless of her report.

Now suppose player 2 observes $s = 1$ and $m_2 = P$. There are two possible explanations: either (i) player 1 trembled after $m_1 \in \{A, B\}$, or (ii) the mediator trembled to $m_1 = \star$, player 1 trembled to $a_1 = C$, player 1 honestly reported $\left(\hat{a}_1, \hat{\theta}\right) = (C, p)$, and the mediator trembled again to $m_2 = P$. Case (i) occurs with probability of order $\frac{1}{k^4}$, while case (ii) occurs with probability of order $\frac{1}{k^3}$. Hence, player 2 believes with probability 1 that $(a_1, \theta) = (C, p)$. This belief rationalizes $a_2 = P$, as is required to implement the desired outcome.

The remainder of the paper is organized as follows. Section 2 defines the model and solution concepts and presents the main theorem, which collects results from Sections 3 and 4. Section 3 presents the (positive and negative) communication RP results. Section 4 presents the implementation RP for SE. Section 5.1 summarizes our results, and Section 5.2 discusses possible extension. Proofs are deferred to the appendix.

2 Multistage Games with Communication

2.1 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) "trembles" by the mediator may be treated differently than trembles by the other players.[10] In each period t, each player i (including the mediator) has a set of possible signals $S_{i,t}$, a set of possible actions $A_{i,t}$, 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). Note also the artificial assumption that the mediator "communicates with himself," which simplifies notation. Throughout the paper, for any set $Y_{i,t}$ indexed by i and t, we let $Y_i^t = \prod_{\tau=1}^{t-1} Y_{i,\tau}$, $Y_t = \prod_{i=0}^{N} Y_{i,t}$, $Y^t = \prod_{\tau=1}^{t-1} Y_\tau$, and $Y = Y^{T+1}$. For example, Y^t is the vector of Y_τ's at the beginning of period t.

Let $H^t = S^t \times R^t \times M^t \times A^t$ 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 $Z = H^{T+1}$ denote the set of terminal nodes of the game. Let $X^t = S^t \times A^t$ denote the set of possible histories of signals and actions ("payoff-relevant histories") at the beginning of

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

period t. Given a complete history $h^t = (s^t, r^t, m^t, a^t) \in H^t$, let $\mathring{h}^t = (s^t, a^t)$ denote the projection of h^t onto X^t. Let $X = X^{T+1} = S \times A$ denote the set of 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 $s_t \in S_t$ is drawn with probability $p(s_t|x^t)$, where $x^t \in X^t$ is the current payoff-relevant history. Player i observes $s_{i,t}$, the i^{th} component of s_t.

2. Each player i 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.

4. Each player i takes an action $a_{i,t} \in A_{i,t}$.

We refer to the tuple $\Gamma := (N, T, S, A, u, p)$ as the *base game* and refer to the pair $C := (R, M)$ as the *communication system*. Assume without loss of generality that $S_{i,t} = \bigcup_{x^t \in X^t} \operatorname{supp} p_i(\cdot|x^t)$ for all i, t, where p_i denotes the marginal distribution of p. The *implementation problem* asks, for a given base game Γ and a given equilibrium concept, which distributions $\rho \in \Delta(X)$ arise in equilibrium for some communication system C?

We now introduce histories, strategies, and beliefs. For each i, let $H_i^t = S_i^t \times R_i^t \times M_i^t \times A_i^t$ denote the set of player i'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 $h_i^t = (s_i^t, r_i^t, m_i^t, a_i^t)$ denote the projection of h^t onto H_i^t. Similarly, $\mathring{h}_i^t = (s_i^t, a_i^t)$ denotes the payoff-relevant component of h_i^t . We also let $H_i^{R,t} = H_i^t \times S_{i,t}$ and $H_i^{A,t} = H_i^{R,t} \times R_{i,t} \times M_{i,t}$ denote *reporting* and *acting* histories for player i, respectively.

A behavioral strategy for player i 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 $\sigma_{i,t}^R : H_i^{R,t} \to \Delta(R_{i,t})$ and $\sigma_{i,t}^A : H_i^{A,t} \to \Delta(A_{i,t})$. This standard definition requires that a player uses the same mixing probability at all nodes in the same information set. Let Σ_i be the set of player i's strategies, and let $\Sigma = \prod_{i=0}^N \Sigma_i$.

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 $\beta_{i,t}^R : H_i^{R,t} \to$

$\Delta\left(H^{R,t}\right)$ and $\beta^A_{i,t} : H^{A,t}_i \to \Delta\left(H^{A,t}\right)$.[11] We write $\sigma^R_{i,t}\left(r_{i,t}|h^{R,t}_i\right)$ for $\sigma^R_{i,t}\left(h^{R,t}_i\right)(r_{i,t})$, and similarly for $\sigma^A_{i,t}$, $\beta^R_{i,t}$, and $\beta^A_{i,t}$. When the meaning is unambiguous, we omit the superscript R or A and the subscript t from σ_i and β_i, so that, for example, σ_i can take as its argument either $h^{R,t}_i$ or $h^{A,t}_i$.

A *mediation plan* is a function $f = (f_t)^T_{t=1}$, where $f_t : R^{t+1} \to M_t$ maps a profile of reports up to and including period t to a profile of period-t messages.[12] A *mixed mediation strategy* is a distribution $\mu \in \Delta(F)$, where F denotes the set of pure mediation plans. A *behavioral mediation strategy* is a function $\phi = (\phi_t)^T_{t=1}$, where $\phi_t : R^{t+1} \times M^t \to M_t$ maps past reports and messages to current messages. Since the mediator can receive signals and take actions in our model, he must choose both a mediation plan f and a report/action strategy σ_0. However, we can equivalently view the mediator as choosing only f, while a separate "dummy player" chooses σ_0. The distinctive feature of the mediator is thus the choice of f, while the strategy σ_0 plays no special role in the analysis and is included only for the sake of generality. As we will see, whether it is most convenient to view the mediator as choosing a pure, mixed, or behavioral mediation strategy depends on the solution concept under consideration. All three perspectives will be used in this paper. In contrast, we always view players (including player 0) as choosing behavioral strategies.

Denote the probability distribution on Z induced by behavioral strategy profile σ and mediation plan f by $\Pr^{\sigma,f}$, and denote the corresponding distribution for a mixed or behavioral mediation strategy by $\Pr^{\sigma,\mu}$ or $\Pr^{\sigma,\phi}$, respectively. Denote the corresponding probability distribution on X (the "outcome distribution") by $\rho^{\sigma,f}$, $\rho^{\sigma,\mu}$, or $\rho^{\sigma,\phi}$. As usual, probabilities are computed assuming that all randomizations (by the players and the mediator) are stochastically independent. We refer to a pair (σ, f), (σ,μ), or (σ,ϕ) as simply a *profile*.

We extend players' payoff functions from terminal histories to profiles in the usual way, writing $\bar{u}_i(\sigma, f)$ for player i's expected payoff at the beginning of the game under profile (σ, f), and writing $\bar{u}_i(\sigma, f|h^t)$ for player i's expected payoff conditional on reaching the complete history h^t. Note that $\bar{u}_i(\sigma, f|h^t)$ does not depend on player i's beliefs, as h^t is a single node in the game tree. The quantities $\bar{u}_i(\sigma,\mu)$, $\bar{u}_i(\sigma,\phi)$, and $\bar{u}_i(\sigma,\phi|h^t)$ are defined

[11] Since the mediator is indifferent over outcomes, there are no optimality conditions on the mediator's strategy, and hence no need to introduce beliefs for the mediator.

[12] Myerson (1986) calls such a function a *feedback rule*.

analogously. In contrast, we avoid the "bad" notation $\bar{u}_i(\sigma, \mu | h^t)$, as this cannot be defined when $\Pr^{\sigma,\mu}(h^t) = 0$.

A *Nash equilibrium (NE)* is a profile (σ, μ) such that $\bar{u}_i(\sigma, \mu) \geq \bar{u}_i(\sigma_i', \sigma_{-i}, \mu)$ for all $i \neq 0, \sigma_i' \in \Sigma_i$. (Or put ϕ in place of μ; the definitions are equivalent by Kuhn's theorem.) In the context of games with communication, a NE is also called a *communication equilibrium* (Forges, 1986).

2.2 Conditional Probability Perfect Bayesian Equilibrium

We consider perfect Bayesian equilibria in which beliefs are derived from a common conditional probability system (CPS) on $F \times Z$. Recall that a CPS on a finite set Ω is a function $\mu(\cdot|\cdot) : 2^\Omega \times 2^\Omega \backslash \emptyset \to [0,1]$ such that (i) for all non-empty $C \subseteq \Omega$, $\mu(\cdot|C)$ is a probability distribution on C, and (ii) for all $A \subseteq B \subseteq C \subseteq \Omega$ with $B \neq \emptyset$, we have $\mu(A|B)\mu(B|C) = \mu(A|C)$. Given a CPS $\bar{\mu}$ on $F \times Z$, $f \in F$, and $Y, Y' \subset Z$, we write $\bar{\mu}(f) = \sum_{z \in Z} \bar{\mu}(f, z)$ and $\bar{\mu}(Y|f, Y') = \sum_{y \in Y} \bar{\mu}(y|f, Y')$. A *conditional probability perfect Bayesian equilibrium (CPPBE)* is a profile (σ, μ) together with a CPS $\bar{\mu}$ on $F \times Z$ such that

- *[CPS Consistency]* For all $f, i, t, h_i^{R,t}, h_i^{A,t}, h^{R,t}, h^{A,t}, r_t, m_t, a_t, s_{t+1}$, we have

$$\begin{aligned} &\bar{\mu}(f) = \mu(f), && \bar{\mu}\left(r_t|f, h^{R,t}\right) = \textstyle\prod_{i=0}^N \sigma_i^R\left(r_{i,t}|h_i^{R,t}\right), \\ &\bar{\mu}\left(a_t|f, h^{A,t}\right) = \textstyle\prod_{i=0}^N \sigma_i^A\left(a_{i,t}|h_i^{A,t}\right), && \bar{\mu}\left(s_{t+1}|f, h^{A,t}, a_t\right) = p\left(s_{t+1}|\mathring{h}^{A,t}, a_t\right), \\ &\bar{\mu}\left(m_t|f, h^{R,t} = \left(s^{t+1}, r^t, m^t, a^t\right), r_t\right) = 1_{\{m_t = f(r^t, r_t)\}}. \end{aligned} \tag{2}$$

  (In these equations, the first argument of $f(\cdot|\cdot)$ must be read as a subset of Z. For example, $\bar{\mu}\left(r_t|f, h^{R,t}\right) = \sum_{z \in Z'} \bar{\mu}\left(z|f, h^{R,t}\right)$, where Z' is the set of all terminal nodes z with period-t report profile r_t.)

- *[Sequential rationality of reports]* For all $i \neq 0, t, \sigma_i' \in \Sigma_i$, and $h_i^{R,t} \in H_i^{R,t}$, we have

$$\sum_{f \in F, h^{R,t} \in H^{R,t}} \bar{\mu}\left(f, h^{R,t}|h_i^{R,t}\right) \bar{u}_i\left(\sigma, f|h^{R,t}\right) \geq \sum_{f \in F, h^{R,t} \in H^{R,t}} \bar{\mu}\left(f, h^{R,t}|h_i^{R,t}\right) \bar{u}_i\left(\sigma_i', \sigma_{-i}, f|h^{R,t}\right). \tag{3}$$

- *[Sequential rationality of actions]* For all $i \neq 0, t, \sigma_i' \in \Sigma_i$, and $h_i^{A,t} \in H_i^{A,t}$, we have

$$\sum_{f \in F, h^{A,t} \in H^{A,t}} \bar{\mu}\left(f, h^{A,t} | h_i^{A,t}\right) \bar{u}_i\left(\sigma, f | h^{A,t}\right) \geq \sum_{f \in F, h^{A,t} \in H^{A,t}} \bar{\mu}\left(f, h^{A,t} | h_i^{A,t}\right) \bar{u}_i\left(\sigma_i', \sigma_{-i}, f | h^{A,t}\right). \tag{4}$$

Note that, in an unmediated game, (i) $\bar{\mu}$ reduces to a CPS on Z, (ii) $\bar{\mu}\left(f, h^{A,t} | h_i^{A,t}\right)$ reduces to a belief $\beta\left(h^{A,t} | h_i^{A,t}\right)$, (iii) (3) disappears, and (iv) (4) reduces to the standard definition of sequential rationality. In the context of unmediated games, the CPPBE 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 top of CPPBE leads to an equivalence with sequential equilibrium in general games. In contrast, our main result is that CPPBE and sequential equilibrium are always outcome-equivalent in games with communication. 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.[13]

For games with communication, the definition of CPPBE follows Myerson in specifying that the mediator plays a mixed (rather than behavioral) strategy and in defining $\bar{\mu}$ as a CPS on $F \times Z$, rather than Z alone. This approach models the mediator's trembles in normal form rather than agent-normal form and thus leads to a more permissive solution concept.[14] If the mediator instead trembled in agent-normal form, the communication RP for CPPBE (Proposition 1 below) would not hold.[15]

[13] Mailath (2019) defines a notion of "almost perfect Bayesian equilibrium," which appears to coincide with CPPBE in unmediated multistage games, though this remains to be proved. 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.

[14] The observation that normal form trembles lead to more permissive solution concepts than agent-normal form trembles is due to Selten (1978).

[15] In particular, the proof of Proposition 5 would go through for canonical CPPBE if the mediator trembled in agent-normal form.

2.3 Sequential Equilibrium

Our definition of sequential equilibrium in a game with communication is simply Kreps-Wilson (1982) sequential equilibrium in the $N+1$ player game where the mediator is treated just like any other player. That is, a *sequential equilibrium (SE)* is an assessment (σ, ϕ, β) consisting of behavioral strategies σ for the players, a behavioral strategy ϕ for the mediator, and beliefs β for the players, such that

- *[Sequential rationality of reports]* For all $i \neq 0, t, \sigma_i'$, and $h_i^{R,t}$, we have

$$\sum_{h^{R,t} \in H^{R,t}} \beta_i \left(h^{R,t} | h_i^{R,t} \right) \bar{u}_i \left(\sigma, \phi | h^{R,t} \right) \geq \sum_{h^{R,t} \in H^{R,t}} \beta_i \left(h^{R,t} | h_i^{R,t} \right) \bar{u}_i \left(\sigma_i', \sigma_{-i}, \phi | h^{R,t} \right).$$

- *[Sequential rationality of actions]* For all $i \neq 0, t, \sigma_i'$, and $h_i^{A,t}$, we have

$$\sum_{h^{A,t} \in H^{A,t}} \beta_i \left(h^{A,t} | h_i^{A,t} \right) \bar{u}_i \left(\sigma, \phi | h^{A,t} \right) \geq \sum_{h^{A,t} \in H^{A,t}} \beta_i \left(h^{A,t} | h_i^{A,t} \right) \bar{u}_i \left(\sigma_i', \sigma_{-i}, \phi | h^{A,t} \right).$$

- *[Kreps-Wilson Consistency]* There exists a sequence of full-support behavioral strategy profiles $\left(\sigma^k, \phi^k\right)_{k=1}^{\infty}$ such that $\lim_{k \to \infty} \left(\sigma^k, \phi^k\right) = (\sigma, \phi)$;

$$\beta_i \left(h^{R,t} | h_i^{R,t} \right) = \lim_{k \to \infty} \frac{\Pr^{\sigma^k, \phi^k} \left(h^{R,t} \right)}{\Pr^{\sigma^k, \phi^k} \left(h_i^{R,t} \right)}$$

  for all $i \neq 0$, all $h_i^{R,t} \in H_i^{R,t}$, and all $h^{R,t} \in H^{R,t}$ with i-component $h_i^{R,t}$; and

$$\beta_i \left(h^{A,t} | h_i^{A,t} \right) = \lim_{k \to \infty} \frac{\Pr^{\sigma^k, \phi^k} \left(h^{A,t} \right)}{\Pr^{\sigma^k, \phi^k} \left(h_i^{A,t} \right)}$$

  for all $i \neq 0$, all $h_i^{A,t} \in H_i^{A,t}$, and all $h^{A,t} \in H^{A,t}$ with i-component $h_i^{A,t}$.

In this definition, the mediator takes a behavioral strategy and trembles in agent-normal form, just like each of the players. An alternative definition, where the mediator cannot tremble at all, would yield a more restrictive version of sequential equilibrium, for which our main results would not hold. We discuss this possibility in Section 5.2.

Theorem 1 of Myerson (1986) shows that every sequence of full-support probability distributions induces a CPS (see also Rényi, 1955). Hence, for any base game and any communication system C, we have $NE(C) \supseteq CPPBE(C) \supseteq SE(C)$, where for any solution concept $\mathcal{S} \in \{NE, CPPBE, SE\}$, $\mathcal{S}(C)$ denotes the set of outcome distributions $\rho \in \Delta(X)$ that arise in equilibrium under solution concept $\mathcal{S}$ with communication system C.

2.4 The Main Result

Given a base game (N, T, S, A, u, p), the *direct communication system* $C^* = (R^*, M^*)$ is given by $R^*_{i,t} = A_{i,t-1} \times S_{i,t}$ and $M^*_{i,t} = A_{i,t}$, for all i and t: that is, players' reports are actions and signals and the mediator's messages are "recommended" actions. Define the *quasi-direct communication system* $C^{**} = (R^*, M^{**})$ by $R^*_{i,t} = A_{i,t-1} \times S_{i,t}$ and $M^{**}_{i,t} = A_{i,t} \cup \{\star\}$, for all i and t, where $\star$ denotes an arbitrary extra message. That is, under quasi-direct communication, in every period a single extra message from the mediator to each player is permitted. Denote the set of all finite communication systems by $\mathcal{C}$.[16] We prove

Theorem 1 *For any base game,*

$$\bigcup_{C \in \mathcal{C}} CPPBE(C) = CPPBE(C^*) = \bigcup_{C \in \mathcal{C}} SE(C) = SE(C^{**}). \tag{5}$$

In addition, for some base game,

$$\bigcup_{C \in \mathcal{C}} SE(C) \supsetneq SE(C^*). \tag{6}$$

The first equality in (5) is the *communication revelation principle for CPPBE.* This formalizes the "revelation principle" implicit in Myerson. We actually prove a stronger version of this equality, namely that with communication system C^* it suffices to consider "canonical" equilibria, where players are honest and obedient. This is proved as Proposition 4 below.

[16] This notion may be formalized by identifying each report or message with a natural number, so that, letting $\mathcal{N}$ denote the set of all finite subsets of the natural numbers, we have $\mathcal{C} = \prod_{t=1}^{T} \prod_{i=0}^{N} (\mathcal{N} \times \mathcal{N})$.

The second equality in (5) is the *implementation revelation principle for SE*. This is the key part of the theorem. It characterizes the set of outcomes implementable in SE, and in particular shows that it is possible to combine the tractability of Myerson's characterization of canonical CPPBE (reviewed in the next subsection) with Kreps-Wilson consistency.

The third equality in (5) may be called the *quasi-communication revelation principle for SE*. It shows that any outcome implementable in SE with any communication system is also implementable in SE with quasi-direct communication. The role of the extra message $\star$ is analogous to the "free pass" message in the opening example and is explained further in Section 4. The second and third equalities are both proved as Proposition 7.

Finally, (6) is the *failure of the communication revelation principle for SE*. This was previewed in the opening example and is proved below as Proposition 5.

2.5 Mediation Range, Sequential Communication Equilibrium, and Myerson's Characterization

Much of the importance of Theorem 1 comes from the fact that the set $CPPBE\,(C^*)$ admits a remarkably simple characterization: it is the set of *sequential communication equilibrium* outcomes, which Myerson shows equals the set of communication equilibrium outcomes with *canonical strategies* in which no player is ever recommended a *codominated action*. To present our proof of Theorem 1 and to understand its significance, it is necessary to review these concepts here.

A sequential communication equilibrium is essentially a canonical CPPBE. To formalize this, note that at first glance there appears to be a conflict between the possibility that the mediator can tremble and the requirement that, in a canonical equilibrium, players obey the mediator's recommendations. For example, a player will never follow a recommendation to play a strictly dominated action, so obedience and mediator trembles are inconsistent if the mediator can tremble to *any* message. We must therefore consider restrictions on the mediator's possible messages.

Following Myerson, a *mediation range* $Q = (Q_i)_{i \neq 0}$ specifies a set of possible messages $Q_i\left(r_i^t, m_i^t, r_{i,t}\right) \subseteq M_{i,t}$ that can be received by each player $i \neq 0$ when the history of communications between player i and the mediator is given by $(r_i^t, m_i^t, r_{i,t})$. Denote the set of mediation plans consistent with mediation range Q by

$$F|_Q = \left\{ f \in F : f_i\left(r_i^{t+1}\right) \in Q_i\left(r_i^{t-1}, \left(f_i\left(r_i^{\tau+1}\right)\right)_{\tau=1}^{t-1}, r_{i,t}\right) \quad \forall i, t, r_i^{t+1} \right\}.$$

Next, given a canonical game (Γ, C^*) and a mediation range Q, the *fully canonical strategy profile* σ^* is defined by letting players report honestly and obey the mediator's recommendation at all histories consistent with Q: that is, we have

- *[Honesty]* $\sigma_i^R\left(h_i^{R,t}\right) = (a_{i,t-1}, s_{i,t})$ for all $h_i^{R,t} \in H_i^{R,t}$ such that $m_{i,\tau} \in Q_i\left(r_i^\tau, m_i^\tau, r_{i,\tau}\right)$ for all $\tau < t$, and
- *[Obedience]* $\sigma_i^A\left(h_i^{A,t}\right) = m_{i,t}$ for all $h_i^{A,t} \in H_i^{A,t}$ such that $m_{i,\tau} \in Q_i\left(r_i^\tau, m_i^\tau, r_{i,\tau}\right)$ for all $\tau \leq t$.

Later on, this will be contrasted with a more general notion of *canonical strategies*, where honesty and obedience are required only for players who have not previously lied to the mediator. Also, since the artificial assumption that the mediator "communicates with himself" is purely for notational convenience, there is no loss in assuming throughout the paper that $C_0 = C_0^*$ and $\sigma_0 = \sigma_0^*$.

Denote the set of terminal, canonical-game histories consistent with Q and honest behavior by the players by

$$Z|_Q = \left\{ z \in Z : r_{i,t} = (a_{i,t-1}, s_{i,t}) \text{ and } m_{i,t} \in Q_i\left(r_i^t, m_i^t, r_{i,t}\right) \quad \forall i, t \right\}.$$

A *sequential communication equilibrium (SCE)* is then a mixed mediation strategy $\mu \in \Delta(F)$ in a canonical game (Γ, C^*) together with a mediation range Q and a CPS $\bar{\mu}$ on $F|_Q \times Z|_Q$ such that

- *[CPS consistency]* For all $f \in F|_Q$, i, t, $h^{R,t} = (s^{t+1}, r^t, m^t, a^t) \in Z^{R,t}|_Q$, $h^{A,t} =$

$(s^{t+1}, r^{t+1}, m^{t+1}, a^t) \in Z^{R,t}|_Q$, m_t, a_t, and s_{t+1}, we have

$$\bar{\mu}(f) = \mu(f), \qquad \bar{\mu}\left(r_t | f, h^{R,t}\right) = 1_{\{r_t = (a_{t-1}, s_t)\}},$$
$$\bar{\mu}\left(a_t | f, h^{A,t}\right) = 1_{\{a_t = m_t\}}, \qquad \bar{\mu}\left(s_{t+1} | f, h^{A,t}, a_t\right) = p\left(s_{t+1} | \mathring{h}^{A,t}, a_t\right),$$
$$\bar{\mu}\left(m_t | f, h^{R,t}, r_t\right) = 1_{\{m_t = f(r^t, r_t)\}}.$$

- *[Sequential rationality of honesty]* For all $i \neq 0, t, \sigma'_i \in \Sigma_i$, and $h_i^{R,t} = \left(s_i^{t+1}, r_i^t, m_i^t, a_i^t\right) \in H_i^{R,t}$ such that $r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$ and $m_{i,\tau} \in Q_i(r_i^\tau, m_i^\tau, r_{i,\tau})$ for all $\tau < t$, we have

$$\sum_{f \in F|_Q, h^{R,t} \in H^{R,t}} \bar{\mu}\left(f, h^{R,t} | h_i^{R,t}\right) \bar{u}_i\left(\sigma^*, f | h^{R,t}\right) \geq \sum_{f \in F|_Q, h^{R,t} \in H^{R,t}} \bar{\mu}\left(f, h^{R,t} | h_i^{R,t}\right) \bar{u}_i\left(\sigma'_i, \sigma^*_{-i}, f | h^{R,t}\right).$$

- *[Sequential rationality of obedience]* For all $i \neq 0, t, \sigma'_i \in \Sigma_i$, and $h_i^{A,t} = \left(s_i^{t+1}, r_i^{t+1}, m_i^{t+1}, a_i^t\right) \in H_i^{A,t}$ such that $r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$ and $m_{i,\tau} \in Q_i(r_i^\tau, m_i^\tau, r_{i,\tau})$ for all $\tau \leq t$, we have

$$\sum_{f \in F|_Q, h^{A,t} \in H^{A,t}} \bar{\mu}\left(f, h^{A,t} | h_i^{A,t}\right) \bar{u}_i\left(\sigma^*, f | h^{A,t}\right) \geq \sum_{f \in F|_Q, h^{A,t} \in H^{A,t}} \bar{\mu}\left(f, h^{A,t} | h_i^{A,t}\right) \bar{u}_i\left(\sigma'_i, \sigma^*_{-i}, f | h^{A,t}\right). \tag{7}$$

The definition of a SCE is similar to the definition of a CPPBE with the direct communication system C^*. (Of course, the general definition of a CPPBE allows $C \neq C^*$.) There are however two differences:

1. SCE requires not only direct communication, but also canonical equilibrium (i.e., players are required to be honest and obedient).

2. SCE imposes sequential rationality for each player only at histories consistent with the pre-specified mediation range Q, and only at histories at which the player has not previously lied to the mediator.

These differences turn out to be immaterial.

Proposition 1 *For any base game Γ and any outcome distribution $\rho \in \Delta(X)$, there exists a SCE $\mu \in \Delta(F)$ with $\rho = \rho^{\sigma^*, \mu}$ if and only if $\rho \in CPPBE(C^*)$.*

Myerson characterizes the set of SCE in terms of codominated actions. Intuitively, an action $a_{i,t}$ is codominated for player i at history $(x_i^t, s_{i,t})$ if it is not optimal for any belief over $\left(x_{-i}^t, s_{-i,t}\right)$, given that other players avoid codominated actions in future periods. Formally, consider a canonical game (Γ, C^*). Given a correspondence B that specifies a set of actions $B_i(x_i^t, s_{i,t}) \subset A_{i,t}$ for each i, t, and $(x_i^t, s_{i,t}) \in X_i^t \times S_{i,t}$, let $E^t(B) = \left\{f \in F : f_i\left(r^{\tau+1}\right) \notin B_i\left(r_i^{\tau+1}\right) \quad \forall i, \tau > t, r^{\tau+1} \in X^\tau \times S_\tau\right\}$ be the set of pure mediation strategies that avoid actions in B after period t. Such a correspondence B is a *codomination correspondence* if, for every period t and every probability distribution $\pi \in \Delta\left(F \times X^t \times S_t\right)$ with $\pi\left(E^t(B) \times X^t \times S_t\right) = 1$ and $\pi\left(f, x^t, s_t\right) > 0$ for some $\left(f, x^t, s_t\right)$ with $f_i\left(x^t, s_t\right) \in B_i\left(x_i^t, s_{i,t}\right)$ for some i, we have

$$\sum_{\substack{\left(f, x^t, s_t\right) \in F \times X^t \times S_t \\ : f_i\left(x^t, s_t\right) = a_{i,t}}} \pi\left(f, x^t, s_t\right) \bar{u}\left(\sigma^*, f | x^t, s_t\right) < \sum_{\substack{\left(f, x^t, s_t\right) \in F \times X^t \times S_t \\ : f_i\left(x^t, s_t\right) = a_{i,t}}} \pi\left(f, x^t, s_t\right) \bar{u}\left(\sigma_i', \sigma_{-i}^*, f | x^t, s_t\right)$$

for some i, $(x_i^t, s_{i,t})$, $a_{i,t} \in B_i\left(x_i^t, s_{i,t}\right)$, and $\sigma_i' \in \Sigma_i$. Finally, let D denote the union of all codomination correspondences (which is itself a codomination correspondence), and say that an action $a_{i,t} \in A_{i,t}$ is *codominated* at history $(x_i^t, s_{i,t})$ if $a_i \in D_i(x_i^t, s_{i,t})$. Myerson's main result is

Proposition 2 (Myerson (1986; Theorem 2, Lemma 1)) *For any base game Γ and any outcome distribution $\rho \in \Delta(X)$, there exists a SCE $\mu \in \Delta(F)$ with $\rho = \rho^{\sigma^*, \mu}$ if and only if there exists a NE (σ^*, μ) in (Γ, C^*) such that $\rho = \rho^{\sigma^*, \mu}$ and $Q_i\left(r_i^t, m_i^t, r_{i,t}\right) \cap D_i(r_i^{t+1}) = \emptyset$ for all i, t, r_i^{t+1}, and m_i^t.*

3 The Communication RP

This section analyzes the communication RP. In addition to proving $\bigcup_{C \in \mathcal{C}} CPPBE(C) = CPPBE(C^*)$ and $\bigcup_{C \in \mathcal{C}} SE(C) \supsetneq SE(C^*)$ (the first and last parts of Theorem 1), we also clarify for what solution concepts and classes of games the classical communication RP holds (which, in addition to requiring $C = C^*$, requires that players are honest and obedient).

The section is organized as follows. Section 3.1 states the communication RP. Section

3.2 presents the communication RP for CPPBE. Section 3.3 notes that, in general, the communication RP fails for SE. Section 3.4 gives conditions under which the communication RP holds for SE.

3.1 Statement of the Communication RP

Given a canonical game (Γ, C^*) a strategy profile $\sigma \in \Sigma$ together with a mediation range Q is *canonical* if the following conditions hold:

1. *[Previously honest players are honest]* $\sigma_i^R\left(h_i^{R,t}\right) = (a_{i,t-1}, s_{i,t})$ for all $h_i^{R,t} \in H_i^{R,t}$ such that $r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$ and $m_{i,\tau} \in Q_i\left(r_i^\tau, m_i^\tau, r_{i,\tau}\right)$ for all $\tau < t$.

2. *[Previously honest players are obedient]* $\sigma_i^A\left(h_i^{A,t}\right) = m_{i,t}$ for all $h_i^{A,t} \in H_i^{A,t}$ such that $r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$ and $m_{i,\tau} \in Q_i\left(r_i^\tau, m_i^\tau, r_{i,\tau}\right)$ for all $\tau \leq t$.

For any game $G = (\Gamma, C)$ and mediation range Q, let $G|_Q$ denote the game where the mediator is restricted to sending messages in Q: that is, if at history (h^t, s_t, r_t) the history of communications between some player i and the mediator is $(r_i^t, m_i^t, r_{i,t})$, then all messages $m_{i,t} \notin Q_i\left(r_i^t, m_i^t, r_{i,t}\right)$ are deleted from the game tree. Note that, for any Q, any strategy profile $\sigma \in \Sigma$ induces a strategy profile in the restricted game $G|_Q$.

Given a canonical game (Γ, C^*) a canonical strategy profile $\sigma \in \Sigma$ together with a mediation range Q such that σ is an equilibrium in $G|_Q$ is a *canonical equilibrium*.

For any game $G = (\Gamma, C)$, let $G^* = (\Gamma, C^*)$ denote the canonical game with the same base game as G. A classical statement of the RP is as follows:[17]

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 G^*.

Forges (1986) established the communication RP for NE.

Proposition 3 (Forges (1986; Proposition 1)) *The communication RP holds for NE.*

[17] Townsend (1988) extends the RP by requiring a player to be honest 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). Theorem 1 shows that enriching the communication system in the way does not expand the set of implementable outcomes.

As we build on this result, we reprise the proof in our notation in the appendix. The intuition is that, in any non-canonical game (Γ, C), we may view each player as reporting her signals and actions to a "personal mediator" under her control, who then communicates with a "central mediator" via communication system C, and then recommends actions to the player. Each player may as well be honest and obedient vis a vis her personal mediator, since she controls her personal mediator's strategy. Now, view the collection of the N personal mediators together with the central mediator as a single mediator in the canonical game (Γ, C^*), where player i's personal mediator now automatically executes her equilibrium communication strategy from the non-canonical game. Then it remains optimal for each player to be honest and obedient, as each player has access to fewer deviations when she cannot directly control her personal mediator.

3.2 The Communication RP for CPPBE

Our first substantive result formalizes the communication RP for CPPBE implicit in Myerson.

Proposition 4 *The communication RP holds for CPPBE. In particular,* $\bigcup_{C \in \mathcal{C}} CPPBE\,(C) = CPPBE\,(C^*)$.

This is much more subtle than the corresponding result for NE. The issue is that it is not obvious how to translate a CPS $\bar{\mu}$ on $F \times Z$ in a non-canonical game to a corresponding CPS in the canonical game. We therefore give an indirect proof, building on Myerson's results. By Proposition 1, an outcome distribution $\rho \in \Delta(X)$ arises in a SCE if and only if it arises in a canonical CPPBE. By Proposition 2, an outcome distribution arises in a SCE if and only if it arises in a NE in which codominated actions are never played with positive probability. Since every (possibly non-canonical) CPPBE is a NE, to prove Proposition 4 it suffices to show that, in every CPPBE, codominated actions are never played with positive probability. This is proved as Lemma 4 in the appendix.

In proving Propositions 1 and 4, we introduce the notions of a "quasi-strategy," which is simply a partially defined strategy, and a "quasi-equilibrium," which is a profile of quasi-strategies where incentive constraints are satisfied wherever strategies are defined. We say that a quasi-equilibrium is "valid" if no unilateral deviation by a player can ever lead to a history where another player's quasi-strategy is undefined. We show that it makes no difference whether we consider fully specified CPPBE or (valid) quasi-CPPBE. This is useful in proving Propositions 1 and 4, as it saves us from having to specify what a player does after she lies to the mediator or after she receives a message outside the mediation range, and it also lets us assume that a previously honest player always believes her opponents have also been honest. A similar approach is useful in the SE analysis in Section 9.

3.3 Failure of the Communication RP for SE

In contrast to the situation for NE or CPPBE, we have

Proposition 5 *The communication RP does not hold for SE. Furthermore,* $\bigcup_{C \in \mathcal{C}} SE(C) \supsetneq SE(C^*)$.

The failure of the communication RP for SE was previewed in the introduction. The stronger result that $\bigcup_{C \in \mathcal{C}} SE(C) \supsetneq SE(C^*)$ (without restricting to honest and obedient strategies) is proved by extending the opening example to include two extra players in such a way as to ensure that action C must be recommended at some history. This implies that a recommendation to play C cannot be used to substitute for the extra "free pass" message, so the set of possible messages must be expanded. In addition, the analysis of these examples is robust to perturbations of the payoff functions and the target outcome distribution, so the failure of the communication RP for SE is generic.

3.4 Sufficient Conditions for the Communication RP for SE

In light of Proposition 5, it is natural to ask when the communication RP for SE 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 SE outcome distribution. This result may be "folk knowledge," 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 applicable to many models of dynamic moral hazard (e.g., Garrett and Pavan, 2012) and dynamic information design (e.g., Ely, 2017).

Third, the communication RP holds in games of *pure adverse selection*: that is, if $|A_{i,t}| = 1$ for all $i \neq 0$ and $t \in \{1, ..., T\}$. 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).

Recall that, given a distribution $\rho \in \Delta(X)$, ρ_i denotes the projection of ρ onto X_i. Let $\|\cdot\|$ denote the sup norm on $\Delta(X)$: for distributions $\rho, \rho' \in \Delta(X)$, $\|\rho - \rho'\| = \max_{x \in X} |\rho(x) - \rho'(x)|$.

Proposition 6 *The following hold:*

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

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

3. *If G is a game of pure adverse selection then any NE outcome distribution is a canonical SE outcome distribution.*

Parts 1 and 2 are fairly straightforward. Part 3 follows from noting that the construction in Proposition 7 is canonical in pure adverse selection games.

4 The Implementation RP

Our main result is that every SCE outcome is implementable in SE for some communication system (in particular, with quasi-direct communication). The failure of the communication RP for SE thus poses no obstacle to the characterization of SE-implementable outcomes.

Proposition 7 *For any base game, an outcome distribution arises in SE for some communication system if and only if it arises in a canonical CPPBE (equivalently, a SCE).*

In addition, every such outcome distribution arises in SE with quasi-direct communication. Hence,

$$CPPBE\left(C^*\right) = \bigcup_{C \in \mathcal{C}} SE\left(C\right) = SE\left(C^{**}\right).$$

Proposition 7 completes the proof of Theorem 1.

As every SE is a CPPBE and the communication RP holds for CPPBE (Proposition 4), we have

$$CPPBE\left(C^*\right) \supseteq \bigcup_{C \in \mathcal{C}} SE\left(C\right).$$

The substance of Proposition 7 is thus that every canonical CPPBE outcome distribution arises as a SE with quasi-direct communication:

$$CPPBE\left(C^*\right) \subseteq SE\left(C^{**}\right).$$

Recall that the quasi-direct communication system $C^{**} = (R^*, M^{**})$ 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 $i \neq 0$ and t, where $\star$ denotes an arbitrary extra message. In our construction, message $\star$ is not used on path. Moreover, players are honest and follow all recommendations other than $\star$, as long as they have done so in the past. The construction is thus "almost" canonical.[18] In addition, the extra message $\star$ 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 $\star$.

Message $\star$ corresponds to the "free pass" in the opening example. 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 plays $a_{i,t} = m_{i,t}$ and trembles with much smaller probability.) In addition, after receiving $\star$, a player's future reports to the mediator are inconsequential (barring future mediator trembles), so honesty is optimal. Based on these honest reports, the mediator's future trembles can be specified so that, conditional on a player receiving a future recommendation to take any non-codominated action, the player's beliefs are those required to motivate that action. For instance, in the

[18] A second way in which our construction is not canonical is that a previously honest but disobedient player may not be honest. This difference from Myerson's approach arises because the SE solution concept limits the consistent beliefs available to a disobedient player (in particular, she cannot believe that her own deviation was correlated with deviations by other players), which makes ensuring honesty difficult.

example, when player 2 receives recommendation $m_2 = P$, he believes that the mediator trembled first to $m_1 = \star$ and then to $m_2 = P$ after $(\hat{a}_1, \hat{\theta}) = (C, p)$, which generates the required belief to motivate $a_2 = P$. Note that it is the possibility that one's *opponents* received message $\star$, trembled, and then reported truthfully that motivates a given player to follow her recommendation.

Propositions 4 and 7 jointly imply that the implementation RP holds for any notion of PBE which is stronger than CPPBE but weaker than SE. 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."

We end this section by sketching the proof of Proposition 7.

It is useful to first briefly review Myerson's proof that the outcome of every NE where codominated actions are never played is a SCE, as we build on this proof. Myerson first constructs an arbitrary SCE with the property that all non-codominated actions are recommended at each history with positive probability along a sequence of move distributions converging to the equilibrium. He then constructs another equilibrium where the mediator mixes this "full-support" SCE with the target NE. By specifying that trembles are much more likely in the former equilibrium, after any history in the mixed equilibrium that lies off-path in the target NE, players believe that the full-support CPPBE is being played, and therefore follow all non-codominated recommendations. Taking the mixing probability to 0 yields a SCE with the same outcome as the target NE, in which all non-codominated recommendations are incentive compatible.

Our construction starts with an arbitrary "full-support" NE: more precisely, a NE in the unmediated, perturbed game where each player must take each action at each history with independent probability at least ε_k. Here, $\varepsilon_k \to 0$ along a sequence of strategy profiles indexed by k converging to the equilibrium, but this convergence is slow compared to other possible trembles: that is, action trembles in the full-support NE are relatively likely. In the SE we construct, the mediator uses the off-path message $\star$ to signal to a player that the full-support NE is being played. Since the full-support NE is an equilibrium in the unmediated game, a player who receives message $\star$ believes that her future reports are almost-surely inconsequential, and thus reports honestly.[19] In particular, when the mediator implements the full-support NE, he recommends $m_{i,t} \in A_{i,t}$ according to the equilibrium strategy of player i with probability $1 - \sqrt{\varepsilon_k}$ and recommends $m_{i,t} = \star$ with probability $\sqrt{\varepsilon_k}$, independently across players. Player i obeys each recommendation $m_{i,t} \in A_{i,t}$ (with negligible tremble probability) and, and after $\star$, takes $a_{i,t}$ according to her equilibrium strategy but trembles with probability $\sqrt{\varepsilon_k}$. Since the mediator tremble to $m_{i,t} = \star$ is independent across players, from the other players' perspectives, it is as if player i plays her NE strategy while trembling with probability $\sqrt{\varepsilon_k} \times \sqrt{\varepsilon_k} = \varepsilon_k$.

In order to provide on-path incentives, the mediator must also be able to recommend specific, non-codominated punishment actions. To make these recommendations incentive compatible, we mix in trembles to mediation plans that recommend all motivatable actions (as in Myerson's construction). A key step in our construction is to show that, since trembles in the full-support NE are relatively likely and players who believe this equilibrium is being played report truthfully, the mediator tremble probabilities can be chosen to generate the beliefs required to motivate each non-codominated action.

An important difficulty is posed by histories that involve multiple surprising signals or recommendations: for example, a player may receive a 0-probability recommendation to play some action a in period t and update her beliefs about the mediation plan accordingly, but may then observe another surprising (i.e., conditional 0-probability) recommendation to play some action a' in a later period t'. We need to ensure that every non-codominated recommendation in period t' is incentive compatible, no matter what recommendations were made in earlier periods. This is challenging, because there is no guarantee that the mediation plan that motivates action a in period t is consistent with the mediation plan that motivates action a' in period t'.

To deal with this, we introduce an additional layer of trembles, whereby the mediator may tremble to recommend any motivatable action even while he still "intends" to implement the full-support NE. These trembles are less likely than both action trembles within the full-support NE and the mediator trembles to mediation plans that rationalize non-codominated

[19]This step is absent in Myerson's proof, as the SCE/CPPBE solution concept allows mediator trembles to be directly correlated with player trembles about which the mediator has no information.

actions. Therefore, when a player receives a 0-probability recommendation to play action a in period t, she believes with probability 1 that the mediator has trembled to the mediation plan that motivates action a; but when she later receives another surprising recommendation to play action a' in period t', she switches to believing that, in fact, her period-t recommendation was due to a recommendation tremble "within" the full-support NE (and thus that, in retrospect, she might have been better-off disobeying the period-t recommendation), while the current, period-t' recommendation to play a' indicates a tremble to the mediation plan that motivates a'.[20]

To complete the construction, this "motivating equilibrium" (the mixture of the full-support NE, the mediation plans that motivate each non-codominated action, and the additional layer of trembles) is mixed with the original target NE, with almost all weight on the latter. Players therefore believe that the mediator follows the target NE until they observe a 0-probability signal or recommendation. Subsequently, players assign probability 1 to the motivating equilibrium, and therefore obey all non-codominated recommendations. Since the original NE did not involve codominated actions, all on- and off-path recommendations are incentive compatible.

5 Conclusion

5.1 Summary

Our main result is that the *implementation RP* holds for SE: to calculate the set of outcomes implementable in SE by any communication system, it suffices to calculate the set of canonical CPPBE outcomes, or equivalently the set of outcomes of canonical NE in which players avoid codominated actions.

We also show that the stronger *communication RP* holds for CPPBE, but not for SE. In particular, while the set of SE-implementable outcomes equals the set of canonical CPPBE outcomes, it may be necessary to allow one extra message to implement some of these

[20] This additional layer of trembles is also not needed in Myerson's proof, because the SCE/CPPBE solution concept allows mediator trembles to off-path recommendations to be directly correlated with the earlier player trembles needed to rationalize such recommendations.

outcomes as SE.

There are however some important settings where the communication RP does hold for SE. These include games where no player can perfectly detect another's deviation, games with a single agent, and games of pure adverse (a class that includes much of the dynamic mechanism design literature).

5.2 Discussion

Sequential Equilibrium without Mediator Trembles In defining sequential equilibrium in games with communication, one must take a position on 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. In the current paper, the mediator can tremble. If the mediator cannot tremble, one obtains a more restrictive version of sequential equilibrium, which in the previous version of this paper we called "machine sequential equilibrium" (MSE), to indicate that the mediator follows his equilibrium strategy mechanically and without error. Gerardi and Myerson (2007; Example 3) showed that, in general, both the communication and implementation RPs can fail for MSE: that is, $\bigcup_{C\in\mathcal{C}} MSE(C) \supsetneq MSE(C^*)$ and $CPPBE(C^*) \supsetneq \bigcup_{C\in\mathcal{C}} MSE(C)$. However, the previous version of this paper showed that Claims 1 and 2 of Proposition 6 hold for MSE, as does a "virtual-implementation" version of Claim 3. Thus, whether the mediator can tremble or not is "almost irrelevant" in games of pure adverse selection.

Infinite Games The dynamic mechanism design literature often assumes a continuum of types or actions to facilitate the use of the envelope theorem, while we restrict to finite games to have a well-defined notion of SE. We conjecture that the communication RP for CPPBE can be extended to infinite games under suitable measurability conditions. In the current paper, we prove the communication RP for CPPBE in two steps: (i) we show that players avoid codominated actions in any CPPBE, and (ii) we appeal to Myerson's Theorem 2 to show that the outcome of any NE where players avoid codominated actions is implementable in a canonical CPPBE. The extension of (i) to infinite games is straightforward; the difficulty is extending (ii), as Myerson also worked with finite games, and in particular relied on characterizing CPS's as limits of full-support move distributions, which is not possible in infinite games. However, we believe Myerson's results can be generalized to infinite games by instead relying on an alternative characterization of CPS's as lexicographic probability systems (Halpern, 2010). This is an interesting question for future research.

Non-Multistage Games Some recent models of dynamic information design go beyond multistage games to consider general extensive-form games that lack a common notion of period (e.g., Doval and Ely, 2019). Modeling communication equilibrium in general extensive-form games is a long-standing open issue, and different approaches are possible (e.g., Forges, 1986; von Stengel and Forges, 2008). What version of the RP might apply in such games is thus another open question.

References

[1] Aoyagi, M. (2010), "Information Feedback in a Dynamic Tournament," *Games and Economic Behavior*, 70, 242-260.

[2] Athey, S. and I. Segal (2013), "An Efficient Dynamic Mechanism," *Econometrica*, 81, 2463-2485.

[3] Battaglini, M. (2005), "Long-Term Contracting with Markovian Consumers," *American Economic Review*, 95, 637-658.

[4] Battigalli, P. (1996), "Strategic Independence and Perfect Bayesian Equilibria," *Journal of Economic Theory*, 70, 201-234.

[5] Bergemann, D. and J. Välimäki (2010), "The Dynamic Pivot Mechanism," *Econometrica*, 78, 771-789.

[6] Bester, H. and R. Strausz (2000), "Imperfect Commitment and the Revelation Principle: The Multi-Agent Case," *Economics Letters*, 69, 165-171.

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

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

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

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

[11] Doval, L. and J. Ely (2019), "Sequential Information Design," *working paper.*

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

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

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

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

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

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

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

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

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

[21] Gershkov, A. and B. Szentes (2009), "Optimal Voting Schemes with Costly Information Acquisition," *Journal of Economic Theory*, 144, 36-68.

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

[23] Halac, M., N. Kartik, and Q. Liu (2016), "Optimal Contracts for Experimentation," *Review of Economic Studies*, 83, 1040-1091.

[24] Halpern, J.Y. (2010), "Lexicographic Probability, Conditional Probability, and Nonstandard Probability," *Games and Economic Behavior*, 68, 155-179.

[25] 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.

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

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

[28] Mailath, G.J. (2019), *Modeling Strategic Behavior*, World Scientific Press.

[29] Makris, M. and L. Renou (2019), "Information Design in Multi-Stage Games," *working paper.*

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

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

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

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

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

[35] Rényi, A. (1955), "On a New Axiomatic Theory of Probability," *Acta Mathematica Hungarica*, 6, 285-335.

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

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

[38] Von Stengel, B. and F. Forges (2008), "Extensive-Form Correlated Equilibrium: Definition and Computational Complexity," *Mathematics of Operations Research*, 33, 1002-1022.

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

Appendix: Omitted Proofs

6 Proof of Proposition 3

We give a more formal version of the argument of Forges (1986; Proposition 1).

Fix a game G and a NE (σ, ϕ). We construct a canonical NE $\left(\tilde{\sigma}, \tilde{\phi}\right)$ in G^* with $\rho^{\tilde{\sigma},\tilde{\phi}} = \rho^{\sigma,\phi}$.

Let $\tilde{\sigma} = \sigma^*$: players are honest and obedient at every history. The mediator's strategy is constructed as follows:

Denote player i's period t report by $\tilde{r}_{i,t} = (\tilde{a}_{i,t-1}, \tilde{s}_{i,t}) \in A_{i,t-1} \times S_{i,t}$, with $A_{i,0} = \emptyset$. 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 $\sigma_i^R (\tilde{s}_{i,1})$ (player i's equilibrium strategy in G, given period-1 signal $\tilde{s}_{i,1}$), independently across players. Given the resulting vector of fictitious reports $r_1 = (r_{i,1})_i$, the mediator draws a vector of "fictitious messages" $m_1 \in M_1$ (the set of possible messages in G) according to $\phi(m_1|r_1)$. Next, given $(\tilde{s}_{i,1}, r_{i,1}, m_{i,1})$, the mediator draws an action recommendation $\tilde{m}_{i,1} \in A_{i,1}$ according to $\sigma_i^A (\tilde{m}_{i,1}|\tilde{s}_{i,1}, r_{i,1}, m_{i,1})$, independently across players. Finally, the mediator sends message $\tilde{m}_{i,1}$ to player i.

Recursively, for $t = 2, \ldots, T$, given player i's reports $\tilde{r}_{i,\tau} = (\tilde{a}_{i,\tau-1}, \tilde{s}_{i,\tau})$ for each $\tau \leq t$ and the fictitious reports and messages $(r_{i,\tau}, m_{i,\tau})$ for each $\tau < t$, the mediator draws $r_{i,t} \in R_{i,t}$ according to $\sigma_i^R(\tilde{s}_i^t, r_i^t, m_i^t, \tilde{a}_i^t, \tilde{s}_{i,t})$, independently across players.[21] Given the resulting vector $r_t = (r_{i,t})_i$, the mediator draws $m_t \in M_t$ according to $\phi(m_t|r^t, m^t, r_t)$. Next, given $(\tilde{s}_{i,t}, r_{i,t}, m_{i,t})$, the mediator draws $\tilde{m}_{i,t} \in A_{i,t}$ according to $\sigma_i^A(\tilde{m}_{i,t}|\tilde{s}_i^t, r_i^t, m_i^t, \tilde{a}_i^t, \tilde{s}_{i,t}, r_{i,t}, m_{i,t})$, independently across players.[22] Finally, the mediator sends message $\tilde{m}_{i,t}$ to player i.

That this profile satisfies $\rho^{\left(\tilde{\sigma},\tilde{\phi}\right)} = \rho^{(\sigma,\phi)}$ follows by induction from the beginning of the game: given that players are honest and obedient, $(\tilde{r}_i^t, r_i^t, m_i^t)$ equals player i's period-t history in the non-canonical game G, so, conditional on each vector $(\tilde{r}_i^t, r_i^t, m_i^t)_i$, the variables $r_{i,t}$, $m_{i,t}$, and $a_{i,t}$ are all chosen with the same probabilities as in G.

We claim that the profile is also a NE. To see this, fix a player i. We show that player i cannot attain a payoff above $u_i\left(\tilde{\sigma}, \tilde{\phi}\right)$ from using any strategy, even if the notion of a "strategy" for player i is extended to allow player i to observe $(s_{i,t}, \tilde{r}_{i,t}, r_{i,t}, m_{i,t}, \tilde{m}_{i,t}, a_{i,t})$ in each period t (rather than only $(s_{i,t}, \tilde{r}_{i,t}, \tilde{m}_{i,t}, a_{i,t})$), and furthermore to directly choose $r_{i,t}$ as a function of $(s_i^t, \tilde{r}_i^t, r_i^t, m_i^t, \tilde{m}_i^t, a_i^t, s_{i,t}, \tilde{r}_{i,t})$ (rather than having the mediator choose $r_{i,t}$ as a pre-specified function of $(\tilde{r}_i^t, r_i^t, m_i^t, \tilde{r}_{i,t})$). (In this thought experiment, the rest of the mediator's strategy—i.e., the determination of $r_{-i,t}$, m_t, and $\tilde{m}_t$ as a function of the other variables available to the mediator—is held constant.) Let $\bar{\Sigma}_i$ be the set of such *extended strategies* for player i.

We first note that, for any extended strategy $\bar{\sigma}_i$, there is another extended strategy $\bar{\sigma}_i' \in \bar{\Sigma}_i$ that does not condition on $\tilde{m}_i^t$ and satisfies $\rho^{\bar{\sigma}'_i,\tilde{\sigma}_{-i},\tilde{\phi}} = \rho^{\bar{\sigma}_i,\tilde{\sigma}_{-i}\tilde{\phi}}$. This follows because neither the mediator nor any other player $j \neq i$ conditions on $\tilde{m}_i^t$; formally, such a strategy can be constructed by letting player i draw an alternative message $\hat{m}_{i,t}$ with

[21] If $(\hat{s}_i^t, \hat{a}_i^t, \hat{s}_{i,t})$ is not a possible history for player i (that is, if there is no $\left(s_{-i}^t, a_{-i}^t, s_{-i,t}\right)$ and (σ', ϕ') such that $\Pr^{(\sigma',\phi')}(\left(\hat{s}_i^t, s_{-i}^t\right), \left(\hat{a}_i^t, a_{-i}^t\right), (\hat{s}_{i,t}, s_{-i,t})) > 0$), the mediator can draw $r_{i,t} \in R_{i,t}$ arbitrarily (e.g., uniformly at random).

[22] Again, if $(\hat{s}_i^t, \hat{a}_i^t, \hat{s}_{i,t})$ is not a possible history, the mediator can draw $\tilde{m}_{i,t}$ randomly.

probability $\sigma_i^A(\hat{m}_{i,t}|\tilde{s}_i^t, r_i^t, m_i^t, \tilde{a}_i^t, \tilde{s}_{i,t}, r_{i,t}, m_{i,t})$ at each history $(\tilde{s}_i^t, r_i^t, m_i^t, \tilde{a}_i^t, \tilde{s}_{i,t}, r_{i,t}, m_{i,t})$, and subsequently follow $\bar{\sigma}_i$ with $\hat{m}_{i,t}$ in place of $\tilde{m}_{i,t}$. Next, for any extended strategy $\bar{\sigma}_i'$ that does not condition on $\tilde{m}_i^t$, there exists another extended strategy $\bar{\sigma}_i''$ that also does not condition on $\tilde{r}_i^t$ and satisfies $\rho^{\left(\bar{\sigma}_i'', \tilde{\sigma}_{-i}, \tilde{\phi}\right)} = \rho^{\left(\bar{\sigma}_i', \tilde{\sigma}_{-i}, \tilde{\phi}\right)}$. This follows because, given that r_i^t is now under the control of player i rather than the mediator, $\tilde{r}_i^t$ enters the mediator's strategy only through the determination of $\tilde{m}_i^t$, which strategy $\bar{\sigma}_i'$ does not condition on; formally, such a strategy can be constructed by letting player i draw an alternative report $\hat{r}_{i,t}$ with probability $\bar{\sigma}_i'^R(\hat{r}_{i,t}|s_i^t, \hat{r}_i^t, r_i^t, m_i^t, \hat{m}_i^t, a_i^t, s_{i,t})$ at each history $(s_i^t, \hat{r}_i^t, r_i^t, m_i^t, \hat{m}_i^t, a_i^t, s_{i,t})$. Here, $\hat{r}_i^t = \emptyset$ for $t = 1$ and for $t \geq 2$, recursively replace $\tilde{r}^t$ with $\hat{r}^t$.

Note that an extended strategy for player i that does not condition on $\tilde{m}_i^t$ or $\tilde{r}_i^t$ can be viewed as, for each t, a mapping from $(S_i^t, R_i^t, M_i^t, A_i^t, S_{i,t})$ to $R_{i,t}$, and a mapping from $(S_i^t, R_i^t, M_i^t, A_i^t, S_{i,t}, R_{i,t}, M_{i,t})$ to $A_{i,t}$—or, equivalently, as a strategy in the non-canonical game G. Moreover, for any such strategy $\bar{\sigma}_i$, we have $\rho^{\bar{\sigma}_i, \tilde{\sigma}_{-i}, \tilde{\phi}} = \rho^{\bar{\sigma}_i, \sigma_{-i}, \phi}$, by the same reasoning as for why $\rho^{\tilde{\sigma}, \tilde{\phi}} = \rho^{\sigma, \phi}$. Since σ is a NE of G, it follows that player i cannot attain a payoff above $u_i\left(\tilde{\sigma}, \tilde{\phi}\right)$ in the canonical game by using any extended strategy that does not condition on $\tilde{m}_i^t$ or $\tilde{r}_i^t$. Hence, player i cannot attain such a payoff by using any actually available strategy in the canonical game. That is, $\tilde{\sigma}$ is a NE in the canonical game.

7 Results for CPPBE

This section contains our analysis of CPPBE, culminating in the proofs of Propositions 1 and 4.

7.1 Quasi-Strategies and Quasi-CPPBE

Fix a game $G = (\Gamma, C)$. For each player i, a *quasi-strategy* (χ_i, J_i) for player i consists of

1. A set of histories $J_i = \prod_{t=1}^{T+1} J_i^{R,t} \cup J_i^{R,t+} \cup J_i^{A,t} \cup J_i^{A,t+}$ with $J_i^{R,t} \subset H_i^{R,t}, J_i^{R,t+} \subset H_i^{R,t} \times R_{i,t}, J_i^{A,t} \subset H_i^{A,t}, J_i^{A,t+} \subset H_i^{A,t} \times A_{i,t}$ for each t, such that (i) for every $h_i^{R,t} \in J_i^{R,t}$ there exists $h_i^{T+1} \in J_i^{T+1}$ that coincides with $h_i^{R,t}$ up to period t, and (ii) for every $h_i^{T+1} \in J_i^{T+1}$ and every $h_i^{R,t} \in H_i^{R,t}$ that coincides with h_i^{T+1} up to period t, we have $h_i^{R,t} \in J_i^{R,t}$, (iii) the same conditions hold for $J_i^{R,t+}, J_i^{A,t}, J_i^{A,t+}$.

2. A function $\chi_i = \left(\chi_i^{R,t}, \chi_i^{A,t}\right)_{t=1}^{T}$, where $\chi_i^{R,t} : J_i^t \to \Delta(R_{i,t})$ and $\chi_i^{A,t} : J_i^t \to \Delta(A_{i,t})$ for each t.

Let $J = \prod_{i=0}^{N} J_i$. Note that $h^t \in J^t$ if and only if $h_i^t \in J_i^t$ for all i.
A *quasi-strategy* for the mediator $(\psi, P, F|_P)$ consists of

1. A set of reports $P = \prod_{t=1}^{T+1} P^t$ with $P^t \subset R^t$ for each t, such that for every $r^t \in P^t$ there exists $r^{T+1} \in P^{T+1}$ that coincides with r^t up to period t.

2. A set $F|_P$, where each $f \in F|_P$ consists of, for each $t = 1, \ldots, T$, a function $f_t : P^t \to M_t$.

3. A probability distribution $\psi \in \Delta(F|_P)$.

We say a quasi-strategy profile $(\chi, \psi, J, P, F|_P)$ is *valid* if

1. For each $f \in F|_P$, $h^t \in J^t$, $i \neq 0$, $\sigma_i \in \Sigma_i$, and $\left(s_{-i}^{T+1}, r^{T+1}, m_{-i}^{T+1}, a_{-i}^{T+1}\right)$ with either $\left(s_{-i}^{T+1}, r_{-i}^{T+1}, m_{-i}^{T+1}, a_{-i}^{T+1}\right) \notin J_{-i}^{T+1}$ or $r^{T+1} \notin P$, we have $\Pr^{(\sigma_i,\chi_{-i},f)}\left(\left(s_{-i}^{T+1}, r^{T+1}, m_{-i}^{T+1}, a_{-i}^{T+1}\right)|h^t\right) = 0$. That is, no unilateral player-deviation leads to a history where either the mediator's or another player's quasi-strategy is undefined. This implies in particular that $\Pr^{(\sigma_i,\chi_{-i},f)}(\cdot|h^t)$ is a well-defined probability distribution on Z, and $\bar{u}_i\left(\sigma_i, \chi_{-i}, f|h^{A,t}\right)$ is a well-defined expected payoff.

2. For each $i \neq 0$ and $h_i^t \in J_i^t$, there exist $h_{-i}^t \in J_{-i}^t$ and $f \in F|_P$ such that $p(s_\tau|x^\tau) > 0$ and $m_{i,\tau} = f\left(r^{\tau+1}\right)$ for all $\tau \leq t$, where s_τ, x^τ, and r^τ are the corresponding components of history $\left(h_i^t, h_{-i}^t\right)$. That is, every history in J_i^t is "explicable" by some opponents' history $h_{-i}^t \in J_{-i}^t$.

Finally, a *quasi-CPPBE* $\left(\chi, \psi, J, P, F|_P, \bar{\psi}\right)$ is a valid quasi-strategy profile $(\chi, \psi, J, P, F|_P)$ together with a CPS $\bar{\psi}$ on $F|_P \times J$ such that no player has a profitable deviation from a history in J_i:

1. *[CPS consistency]* For all $f, t, h_i^{R,t} \in J_i^{R,t}, h_i^{A,t} \in J_i^{A,t}, h^{R,t} \in J^{R,t}, h^{A,t} \in J^{A,t}, r_t, m_t, a_t, s_{t+1}$, we have

$$\bar{\psi}(f) = \psi(f), \qquad \bar{\psi}\left(r_t|f, h^{R,t}\right) = \prod_{i=0}^N \chi_i^R\left(r_{i,t}|h_i^{R,t}\right),$$
$$\bar{\psi}\left(a_t|f, h^{A,t}\right) = \prod_{i=0}^N \chi_i^A\left(a_{i,t}|h_i^{A,t}\right), \quad \bar{\psi}\left(s_{t+1}|f, h^{A,t}, a_t\right) = p\left(s_{t+1}|\mathring{h}^{A,t}, a_t\right)$$
$$\bar{\psi}\left(m_t|f, h^{R,t}, r_t\right) = 1_{\{m_t = f(r^t, r_t)\}},$$

2. *[Sequential rationality of reports]* For all $i \neq 0, t, \sigma_i' \in \Sigma_i$, and $h_i^{R,t} \in J_i^{R,t}$, we have

$$\sum_{f \in F|_P, h^{R,t} \in J} \bar{\psi}\left(f, h^{R,t}|h_i^{R,t}\right) \bar{u}_i\left(\chi, f|h^{R,t}\right) \geq \sum_{f \in F|_P, h^{R,t} \in J} \bar{\psi}\left(f, h^{R,t}|h_i^{R,t}\right) \bar{u}_i\left(\sigma_i', \chi_{-i}, f|h^{R,t}\right). \tag{8}$$

3. *[Sequential rationality of actions]* For all $i \neq 0, t, \sigma_i' \in \Sigma_i$, and $h_i^{A,t} \in J_i^{A,t}$, we have

$$\sum_{f \in F|_P, h^{A,t} \in J} \bar{\psi}\left(f, h^{A,t}|h_i^{A,t}\right) \bar{u}_i\left(\chi, f|h^{A,t}\right) \geq \sum_{f \in F|_P, h^{A,t} \in J} \bar{\psi}\left(f, h^{A,t}|h_i^{A,t}\right) \bar{u}_i\left(\sigma_i', \chi_{-i}, f|h^{A,t}\right). \tag{9}$$

Let $\rho^{(\chi,\psi)} \in \Delta(X)$ denote the outcome distribution induced by valid quasi-strategy profile (χ, ψ). The following lemma says that it is without loss to consider quasi-CPPBE rather than fully specified CPPBE.

Lemma 1 *For any game $G = (\Gamma, C)$ and any outcome distribution $\rho \in \Delta(X)$, we have $\rho \in CPPBE(C)$ if and only if $\rho = \rho^{(\chi,\psi)}$ for some quasi-CPPBE profile $\left(\chi, \psi, J, P, F|_P, \bar{\psi}\right)$.*

Proof. Fix a game G. One direction of the lemma is immediate: If $(\sigma, \mu, \bar{\mu})$ is a CPPBE in G and we let $\chi = \sigma$, $\psi = \mu$, $J = Z$, $P = R$, $F|_P = F$, and $\bar{\psi} = \bar{\mu}$, then $(\chi, \psi, J, P, F|_P, \bar{\psi})$ is a quasi-CPPBE with $\rho^{(\chi,\psi)} = \rho^{(\sigma,\mu)}$.

For the converse, fix a quasi-CPPBE $(\chi, \psi, J, P, F|_P, \bar{\psi})$. We say that a *move distribution on* $F_P \times J$ is a triple $(\alpha^F, \alpha^R, \alpha^A)$, where $\alpha^F \in \Delta(F|_P)$, $\alpha^R = (\alpha^{R,t})_{t=1}^T$ with $\alpha^{R,t} : F|_P \times J^{R,t} \to \Delta(R_t)$, and $\alpha^A = (\alpha^{A,t})_{t=1}^T$ with $\alpha^{A,t} : F|_P \times J^{A,t} \to \Delta(A_t)$. A move distribution on $F_P \times J$ has *full support* if (i) for each $f \in F|_P$, $\alpha^F(f) > 0$, (ii) for each $f \in F|_P$ and $h^{R,t} \in J^{R,t}$, $\alpha^{R,t}(r_t|f, h^{R,t}) > 0$ if and only if $(h^{R,t}, r_t) \in J^{R,t}$, and (iii) for each $f \in F|_P$ and $h^{A,t} \in J^{A,t}$, $\alpha^{A,t}(a_t|f, h^{A,t}) > 0$ if and only if $(h^{A,t}, a_t) \in J^{A,t}$.

By Myerson's Theorem 1, every CPS is the limit of conditional probabilities derived from a sequence of full support move distributions. Thus, there exists a sequence of move distributions $(\alpha^{F,k}, \alpha^{R,k}, \alpha^{A,k})_{k\in\mathbb{N}}$ with full support on $F|_P \times J$ such that (i) $\alpha^{F,k}(f) \to \psi(f)$ for all $f \in F|_P$, (ii) $\alpha^{R,k}(r_t|f, h^{R,t}) \to \prod_{i=0}^N \chi_i^R\left(r_{i,t}|h_i^{R,t}\right)$ for all $r_t \in R_t, h^{R,t} \in J^{R,t}$, and (iii) $\alpha^{A,k}(a_t|f, h^{A,t}) \to \prod_{i=0}^N \chi_i^A(a_{i,t}|h_i^{A,t})$ for all $a_t \in A_t, h^{A,t} \in J^{A,t}$. For each k, let

$$\varepsilon_k = \min_{f\in F|_P, t, (h^{R,t}, r_t, m_t, a_t)\in J^{t+1}} \min\{\alpha^{F,k,t}(f), \alpha^{R,k,t}(r_t|f, h^{R,t}), \alpha^{A,k,t}(a_t|f, h^{R,t}, r_t, m_t)\}. \tag{10}$$

Given $f \in F|_P$, let $F(f) = \{f' \in F : f'(h^t) = f(h^t) \quad \forall t, h^t \in J^t\}$. For $k \in \mathbb{N}$, consider the following auxiliary game (Γ^k, C):

1. The mediator uses a mixed mediation strategy $\mu^k \in \Delta(F)$ defined as follows: (i) with probability $1 - \frac{\varepsilon_k}{k}$, draw $f \in F|_P$ according to $\alpha^{F,k} \in \Delta(F|_P)$, and then draw $f' \in F$ uniformly at random from $F(f)$; (ii) with probability $\frac{\varepsilon_k}{k}$, draw $f' \in F$ uniformly at random from F.

2. Each player i chooses probability distributions $\sigma_i^k(\cdot|h_i^{R,t}) \in \Delta(R_{i,t})$ and $\sigma_i^k(\cdot|h_i^{A,t}) \in \Delta(A_{i,t})$ for each $t, h_i^{R,t} \in H_i^{R,t}\backslash J_i^{R,t}, h_i^{A,t} \in H_i^{A,t}\backslash J_i^{A,t}$. At histories $h_i^t \in J_i^t$, player i has no choice to make, and we formally set $\sigma_i^k(\cdot|h_i^{R,t}) = \chi_i\left(\cdot|h_i^{R,t}\right)$ and $\sigma_i^k(\cdot|h_i^{A,t}) = \chi_i\left(\cdot|h_i^{A,t}\right)$.

3. Given f (drawn from μ^k) and strategy profile σ^k, the distribution of H^{T+1} is determined recursively as follows:

   Given $f \in F$ and $h^{R,t} \in H^{R,t}$, each $r_t \in R_t$ is drawn with probability

$$\begin{cases} \left(1 - \frac{\varepsilon_k}{k}\left|\{\tilde{r}_t \in R_t : (h^{R,t}, \tilde{r}_t) \notin J^{R,t+}\}\right|\right) \alpha^{R,k}(r_t|f, h^{R,t}) & \text{if } h^{R,t} \in J^{R,t} \wedge (h^{R,t}, r_t) \in J^{R,t+}, \\ \frac{\varepsilon_k}{k} & \text{if } h^{R,t} \in J^{R,t+} \wedge (h^{R,t}, r_t) \notin J^{R,t+}, \\ \prod_{i=0}^N \left(\left(1 - \frac{\varepsilon_k}{k}|R_{i,t}|\right)\sigma_i^{R,k}(r_{i,t}|h_i^{R,t}) + \frac{\varepsilon_k}{k}\right) & \text{if } h^{R,t} \notin J^{R,t}. \end{cases}$$

Given $f \in F$ and $h^{A,t} \in H^{A,t}$, each $a_t \in A_t$ is drawn with probability

$$\begin{cases} \left(1 - \frac{\varepsilon_k}{k} \left|\{(\tilde{a}_t, \tilde{s}_{t+1}) \in A_t \times S_{t+1} : \left(h^{R,t}, \tilde{a}_t, \tilde{s}_{t+1}\right) \notin J^{A,t+}\}\right|\right) \times \alpha^{A,k}(a_t|f, h^{A,t}) & \text{if } h^{A,t} \in J^{A,t} \wedge \left(h^{R,t}, a_t\right) \in J^{A,t+}, \\ \frac{\varepsilon_k}{k} & \text{if } h^{A,t} \in J^{A,t} \wedge \left(h^{R,t}, a_t\right) \notin J^{A,t+}, \\ \prod_{i=0}^{N} \left(\left(1 - \frac{\varepsilon_k}{k} \left|A_{i,t}\right|\right) \sigma_i^{R,k}(r_{i,t}|h_i^{R,t}) + \frac{\varepsilon_k}{k}\right) & \text{if } h^{A,t} \notin J^{A,t}. \end{cases}$$

Then, given $a_t \in A_t$, each $s_{t+1} \in S_{t+1}$ is drawn with probability $p\left(s_{t+1}|\mathring{h}^{A,t}, a_t\right)$.

4. Given realized outcome x, player i's payoff is $u_i(x)$.

Note that the strategy set in $\left(\Gamma^k, C\right)$ is a product of simplices. In addition, each player i's utility is continuous in σ^k and affine (and hence quasi-concave) in σ_i^k. Hence, the Debreu-Fan-Glicksberg theorem guarantees existence of a NE in $\left(\Gamma^k, C\right)$. Moreover, since $\left(\mu^k, \sigma^k\right)$ has full support on $F \times Z$ for any strategy profile σ^k in $\left(\Gamma^k, C\right)$, Bayes' rule defines a CPS $\bar{\mu}^k$ on $F \times Z$.

So let $(\mu^k, \sigma^k, \bar{\mu}^k)_k$ denote a sequence of NE σ^k and corresponding CPS's $\bar{\mu}^k$ in $\left(\Gamma^k, C\right)$. Taking a subsequence if necessary to guarantee convergence, let $(\mu, \sigma, \bar{\mu}) = \lim_k(\mu^k, \sigma^k, \bar{\mu}^k)$. We claim that $(\mu, \sigma, \bar{\mu})$ is a CPPBE in (Γ, C). Since $\bar{\mu}$ is a CPS as the limit of conditional probabilities, it remains to verify sequential rationality. We consider reporting histories $h_i^{R,t}$; the argument for acting histories $h_i^{A,t}$ is symmetric.

There are two cases, depending on whether or not $h_i^{R,t} \in J_i^{R,t}$. If $h_i^{R,t} \notin J_i^{R,t}$, then $h_i^{T+1} \notin J_i^{T+1}$ for all h_i^{T+1} that follow $h_i^{R,t}$, so by inspection player i's expected payoff conditional on $h_i^{R,t}$ is continuous in μ^k, σ^k, ε_k, and k. Since $\sigma_i^k\left(\cdot|h_i^{R,t}\right)$ is sequentially rational in $\left(\Gamma^k, C\right)$ (as $\left(\mu^k, \sigma^k\right)$ is a NE in $\left(\Gamma^k, C\right)$ with full support), it follows that $\sigma_i\left(\cdot|h_i^{R,t}\right)$ is sequentially rational in (Γ, C).

Now consider the case where $h_i^{R,t} \in J_i^{R,t}$. Note that, for each $h_i^{T+1} \in J_i^{T+1}$ and $\left(f, h_{-i}^{T+1}\right) \notin F|_P \times J_{-i}^{T+1}$, there exists $\left(f', h_{-i}'^{T+1}\right) \in F|_P \times J_{-i}^{T+1}$ such that

$$\lim_k \frac{\bar{\mu}^k(f, h_i^{T+1}, h_{-i}^{T+1})}{\bar{\mu}^k(f', h_i^{T+1}, h_{-i}'^{T+1})} = 0.$$

This follows because in $\left(\Gamma^k, C\right)$ each "tremble" leading to a history outside J occurs with probability at most ε_k/k, while every history $h_i^{T+1} \in J_i^{T+1}$ occurs with positive probability given move distribution $\left(\alpha^{F,k}, \alpha^{R,k}, \alpha^{A,k}\right)$ (this is an implication of the second condition in the definition of a valid quasi-strategy profile), and with this distribution each move occurs with probability at least ε_k.

Therefore, for each $f \in F|_P$ and $h_i^{R,t} \in J_i^{R,t}$, we have $\bar{\mu}(f, h_{-i}^{R,t}|h_i^{R,t}) = \bar{\psi}(f, h_{-i}^{R,t}|h_i^{R,t})$, and the conditional probability that $\left(f, h_i^{R,t}\right) \in F|_P \times J_i^{R,t}$ equals 1. Hence, the fact that (8) holds with CPS $\bar{\psi}$ implies that $\sigma_i(\cdot|h_i^{R,t}) = \chi_i\left(\cdot|h_i^{R,t}\right)$ is sequentially rational in (Γ, C). ■

7.2 Quasi-CPPBE and Mediation Ranges

The following lemma says that the ability to restrict the mediator's messages to lie in some mediation range does not expand the set of implementable outcomes. Let $\mathcal{Q}(G)$ denote the set of all possible mediation ranges in game G.

Lemma 2 *For any game G, we have $CPPBE(G) = \bigcup_{Q \in \mathcal{Q}(G)} CPPBE(G|_Q)$.*

Proof. If we let Q be the trivial mediation range that never excludes any messages (i.e., $Q_i(r_i^t, m_i^t, r_{i,t}) = M_{i,t}$ for all $i, t, r_i^t, m_i^t, r_{i,t}$), then $G = G|_Q$. Hence, $CPPBE(G) \subseteq \bigcup_{Q \in \mathcal{Q}(G)} CPPBE(G|_Q)$.

It remains to show that, for any mediation range Q, we have $CPPBE(G|_Q) \subseteq CPPBE(G)$. Fix a mediation range Q and a CPPBE $(\sigma, \mu, \bar{\mu})$ in $G|_Q$. For each i and t, let

$$\begin{aligned} J_i^{R,t} &= \left\{ h_i^{R,t} \in H_i^{R,t} : m_{i,\tau} \in Q_i(r_i^\tau, m_i^\tau, r_{i,\tau}) \quad \forall \tau < t \right\} \text{ and} \\ J_i^{A,t} &= \left\{ h_i^{A,t} \in H_i^{A,t} : m_{i,\tau} \in Q_i(r_i^\tau, m_i^\tau, r_{i,\tau}) \quad \forall \tau \leq t \right\}, \end{aligned}$$

let $J_i^{R,t+} = \left\{ \left(h_i^{R,t}, r_{i,t} \right) : h_i^{R,t} \in J_i^{R,t} \right\}$ and $J_i^{A,t+} = \left\{ \left(h_i^{A,t}, a_{i,t} \right) : h_i^{A,t} \in J_i^{A,t} \right\}$, let $P = R$, and let

$$F|_P = \prod_{t=1}^{T} \left\{ f_t : \prod_{\tau=1}^{t} R_\tau \to Q\left(r^t, \left(f_\tau (r^\tau, r_\tau) \right)_{\tau=1}^{t-1}, r_t \right) \right\}.$$

Note that $J = Z|_Q$ (the set of terminal nodes in $G|_Q$) and $F|_P = F|_Q$ (the set of pure mediation strategies in $G|_Q$). Hence, (σ, J) is a player quasi-strategy profile in G, $(\mu, P, F|_P)$ is a quasi-strategy for the mediator in G, and $\bar{\mu}$ is a CPS on $F|_P \times J$ in G. Moreover, the quasi-strategy profile $(\sigma, \mu, J, P, F|_P)$ is valid in G, since (i) histories outside J cannot arise under any mediation strategy in F (so the first condition in the definition of a valid quasi-strategy profile is satisfied) and (ii) every message history in J can arise for some mediation strategy in F (so the second condition is also satisfied). Finally, the conditions for $(\sigma, \mu, \bar{\mu})$ to be a CPPBE in $G|_Q$ are precisely the conditions for $(\sigma, \mu, J, P, F|_P, \bar{\mu})$ to be a quasi-CPPBE in G. Hence, the latter is a quasi-CPPBE in G, so Lemma 1 implies that $\rho^{(\sigma,\mu)} \in CPPBE(G)$.
■

7.3 SCE Implies Quasi-CPPBE

Lemma 3 *For any game G and any SCE $(\mu, Q, \bar{\mu})$, there exists a canonical CPPBE with outcome distribution $\rho^{(\sigma^*,\mu)}$.*

Proof. In the canonical game G^*, let

$$\begin{aligned} J_i^{R,t} &= \left\{ \begin{array}{c} h_i^{R,t} \in H_i^{R,t} : \\ r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau}) \text{ and } m_{i,\tau} \in Q_i(r_i^\tau, m_i^\tau, r_{i,\tau}) \quad \forall \tau < t \end{array} \right\} \text{ and} \\ J_i^{A,t} &= \left\{ \begin{array}{c} h_i^{A,t} \in H_i^{A,t} : \\ r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau}) \text{ and } m_{i,\tau} \in Q_i(r_i^\tau, m_i^\tau, r_{i,\tau}) \quad \forall \tau \leq t \end{array} \right\}, \end{aligned}$$

let

$$J_i^{R,t+} = \left\{\left(h_i^{R,t}, r_{i,t}\right) : h_i^{R,t} \in J_i^{R,t} \text{ and } r_{i,t} = (a_{i,t-1}, s_{i,t})\right\} \text{ and}$$
$$J_i^{A,t+} = \left\{\left(h_i^{A,t}, a_{i,t}\right) : h_i^{A,t} \in J_i^{A,t} \text{ and } a_{i,t} \in A_{i,t}\right\},$$

let $P = R$, and let

$$F|_P = \prod_{t=1}^{T} \left\{ f_t : \prod_{\tau=1}^{t} R_\tau \to Q\left(r^t, \left(f_\tau\left(r^\tau, r_\tau\right)\right)_{\tau=1}^{t-1}, r_t\right) \right\}.$$

Consider the quasi-strategy profile $(\chi, \psi, J, P, F|_P)$ where χ_i is honest and obedient at each $h_i^t \in J_i^t$ (i.e., $\chi_i^{R,t}\left(h_i^{R,t}\right) = (a_{i,t-1}, s_{i,t})$ for all $h_i^{R,t} \in J_i^{R,t}$ and $\chi_i^{A,t}\left(h_i^{A,t}\right) = m_{i,t}$ for all $h_i^{A,t} \in J_i^{A,t}$). This quasi-strategy profile is valid, since (i) histories outside J_i can arise only if player i is dishonest or the mediator uses a strategy outside $F|_P$ and (ii) every message history in J can arise for some mediation strategy in $F|_P$. Moreover, by inspection, $(\chi, \psi, J, P, F|_P, \bar{\mu})$ is a quasi-CPPBE in G^* if and only if $(\mu, Q, \bar{\mu})$ is a SCE. Hence, the former is a quasi-CPPBE in G^*, so Lemma 1 implies that $\rho^{(\chi,\psi)} \in CPPBE\,(C^*)$. ■

7.4 CPPBE and Codominated Actions

We now show that, in any CPPBE, players do not take codominated actions at any history.

Lemma 4 *For any game G and CPPBE $(\sigma, \mu, \bar{\mu})$,* $\operatorname{supp} \sigma_i(h_i^{A,t}) \cap D_i(\mathring{h}_i^{A,t}) = \emptyset$ *for all i and $h_i^{A,t} \in H_i^{A,t}$.*

7.4.1 Proof of Lemma 4

Fix a game G and a CPPBE $(\sigma, \mu, \bar{\mu})$. The sequential rationality condition at history $h_i^{A,t}$ is

$$\begin{aligned} &\sum_{f \in F, h^{A,t} \in H^{A,t}, x \in X} \bar{\mu}(f, h^{A,t} | h_i^{A,t}) \Pr^{\sigma}\left(x | f, h^{A,t}\right) u_i\left(x\right) \\ = &\max_{\sigma_i' \in \Sigma_i} \sum_{f \in F, h^{A,t} \in H^{A,t}, x \in X} \bar{\mu}(f, h^{A,t} | h_i^{A,t}) \Pr^{\sigma_i', \sigma_{-i}}\left(x | f, h^{A,t}\right) u_i\left(x\right). \end{aligned} \tag{11}$$

Let $\Sigma_{i,t}^*$ denote the set of continuation strategies for player i starting with the period t action phase in the canonical game G^* that do not depend on m_i^t. Let F^* denote the set of pure mediation strategies in G^*. We wish to prove the following lemma, which replaces $h^{A,t}$ in (11) with its payoff-relevant component (x^t, s_t):

Lemma 5 *There exists a CPS $\hat{\mu}$ on $F^* \times X^t \times S_t \times A_t$ such that, for each i, $x_i^t \in X_i^t$,*

$s_{i,t} \in S_{i,t}$, $\tilde{m}_{i,t} \in \bigcup_{h_i^{A,t}:\mathring{h}_i^{A,t}=\left(x_i^t,s_{i,t}\right)} \operatorname{supp} \sigma_i^A(h_i^{A,t})$, *and* $\sigma_i' \in \Sigma_{i,t}^*$, *we have*

$$\begin{aligned} & \sum_{\hat{f}\in F^*,(x^t,s_t)\in X^t\times S_t,x\in X} \hat{\mu}(\hat{f},x^t,s_t|x_i^t,s_{i,t},\tilde{m}_{i,t}) \Pr^{\sigma^*}\left(x|\hat{f},x^t,s_t\right) u_i(x) \\ \geq & \sum_{\hat{f}\in F^*,(x^t,s_t)\in X^t\times S_t,x\in X} \hat{\mu}(\hat{f},x^t,s_t|x_i^t,s_{i,t},\tilde{m}_{i,t}) \Pr^{\sigma_i',\sigma_{-i}^*}\left(x|\hat{f},x^t,s_t\right) u_i(x). \qquad (12) \end{aligned}$$

Inequality (12) is equivalent to (7), the sequential rationality of obedience condition in the definition of SCE. Hence, by Myerson's Lemma 1, (12) implies that $\tilde{m}_{i,t} \notin D_i\left(x_i^t, s_{i,t}\right)$. This completes the proof of Lemma 4.

To prove Lemma 5, we define an auxiliary game G^k and construct an equilibrium in G^k for which (11) is the sequential rationality constraint and (12) is a relaxed version of this constraint. The game G^k is equal to the restriction of G^* from periods t to T, with an "initial state" (x^t, s_t) drawn from $\bar{\mu}^k(x^t, s_t)$, where $\left(\bar{\mu}^k\right)_{k\in\mathbb{N}}$ is a sequence of full-support CPS's that converge to $\bar{\mu}$ as $k \to \infty$. When (x^t, s_t) is drawn, each player i observes $(x_i^t, s_{i,t})$, and the mediator observes (x^t, s_t). Finally, the mediator's strategy is fixed as follows: after observing (x^t, s_t), the mediator draws a mediation plan f and a "fictitious history" $h^{A,t}$ according to $\bar{\mu}^k(f, h^{A,t}|x^t, s_t)$. In period t, the mediator recommends $\tilde{m}_{i,t} \in A_{i,t}$ according to $\sigma_i^A(\tilde{m}_{i,t}|h_i^{A,t})$, independently across players. For the rest of the game, the mediator takes some behavioral strategy $\phi_{f,h^{A,t}}$, independent of k.

Lemma 6 *There exists* $(\phi_{f,h^{A,t}})_{f\in F,h^{A,t}\in H^{A,t}}$ *such that, for each* $(x_i^t, s_{i,t})$ *and* $\tilde{m}_{i,t} \in \bigcup_{h_i^{A,t}:\mathring{h}_i^{A,t}=(x_i^t,s_{i,t})} \operatorname{supp} \sigma_i(h_i^{A,t})$, *we have*

$$\begin{aligned} & \sum_{f\in F,h^{A,t}\in H^{A,t},x\in X} \bar{\mu}(f,h^{A,t}|x_i^t,s_{i,t},\tilde{m}_{i,t}) \Pr^{\sigma^*}\left(x|\phi_{f,h^{A,t}},h^{A,t},\tilde{m}_{i,t}\right) u_i(x) \\ = & \max_{\sigma_i'\in\Sigma_{i,t}^*} \sum_{f\in F,h^{A,t}\in H^{A,t},x\in X} \bar{\mu}(f,h^{A,t}|x_i^t,s_{i,t},\tilde{m}_{i,t}) \Pr^{\sigma_i',\sigma_{-i}^*}\left(x|\phi_{f,h^{A,t}},h^{A,t},\tilde{m}_{i,t}\right) u_i(x), \qquad (13) \end{aligned}$$

where σ^* *is the fully canonical strategy and* $\bar{\mu}$ *is the mediator's limit strategy, given by*

$$\bar{\mu}(f,h^{A,t}|\mathring{h}_i^{A,t},\tilde{m}_{i,t}) = \lim_k \frac{\bar{\mu}^k(\mathring{h}^{A,t}|\mathring{h}_i^{A,t})\bar{\mu}^k(f,h^{A,t}|\mathring{h}^{A,t})\sigma_i^A(\tilde{m}_{i,t}|h_i^{A,t})}{\bar{\mu}^k(\mathring{h}_i^{A,t})\bar{\mu}^k(f,h_i^{A,t}|\mathring{h}_i^{A,t})\sigma_i^A(\tilde{m}_{i,t}|h_i^{A,t})}.$$

Proof. Construction of $\phi_{f,h^{A,t}}$: This is similar to the usual revelation principle argument (e.g., the proof of Proposition 3). Denote player i's period $t+1$ report by by $\tilde{r}_{i,t+1} \in A_{i,t}\times S_{i,t+1}$. Given $(h_i^{A,t}, \tilde{r}_{i,t+1})$, the mediator draws a "fictitious report" $r_{i,t+1} \in R_{i,t+1}$ according to $\sigma_{i,t+1}^R(h_i^{A,t}, \tilde{r}_{i,t+1})$.[23] Next, given r_{t+1}, the mediator calculates a vector of "fictitious messages" $m_{t+1} = f(h^{A,t}, r_{t+1})$. Finally, the mediator draws action recommendation $\tilde{m}_{i,t+1} \in A_{i,t+1}$ according to $\sigma_i^A(h_i^{A,t}, \tilde{a}_{i,t}, \tilde{s}_{i,t+1}, r_{i,t+1}, m_{i,t+1})$, independently across players, and sends message $\tilde{m}_{i,t+1}$ to player i.

[23] As in the proof of Proposition 3, $r_{i,\tau}$ and (in what follows) $\tilde{m}_{i,\tau}$ can be chosen arbitrarily if player i's fictitious history is infeasible.

Recursively, for each $t' > t+1$, given player i's reports $\tilde{r}_{i,\tau} = (\tilde{a}_{i,\tau-1}, \tilde{s}_{i,\tau})$ for each $t \leq \tau \leq t'$ and the fictitious reports and messages $(r_{i,\tau}, m_{i,\tau})$ for each $t \leq \tau < t'$, the mediator draws $r_{i,t'} \in R_{i,t'}$ according to $\sigma_i^R(h_i^{A,t}, \tilde{a}_{i,t}, \tilde{s}_{i,t+1}, r_{i,t+1}, m_{i,t+1}, \ldots, r_{i,t'-1}, m_{i,t'-1}, \tilde{a}_{i,t'-1}, \tilde{s}_{i,t'})$. Next, given $r_{t'}$,the mediator calculates $m_{t'} = f(h^{A,t}, r_{t+1}, \ldots, r_{t'})$. Finally, the mediator draws $\tilde{m}_{i,t'} \in A_{i,t'}$ according to $\sigma_i^A(h_i^{A,t}, \tilde{a}_{i,t}, \tilde{s}_{i,t+1}, \ldots, r_{i,t'}, m_{i,t'})$, independently across players, and sends message $\tilde{m}_{i,t'}$ to player i.

Proof of (13): Suppose (13) is violated for some $(x_i^t, s_{i,t})$ and $\tilde{m}_{i,t} \in \bigcup_{h_i^{A,t}:\mathring{h}_i^{A,t}=(x_i^t,s_{i,t})} \operatorname{supp} \sigma_i(h_i^{A,t})$. Then, there exist $h_i^{A,t}$ with $\bar{\mu}(h_i^{A,t}|\mathring{h}_i^{A,t}, \tilde{m}_{i,t}) > 0$ and $\sigma_i' \in \Sigma_{i,t}^*$ such that

$$\begin{aligned} &\sum_{f\in F, h^{A,t}, x\in X} \bar{\mu}(f, h^{A,t}|h_i^{A,t}, \tilde{m}_{i,t}) \Pr^{\sigma^*}\left(x|\phi_{f,h^{A,t}}, h^{A,t}, \tilde{m}_{i,t}\right) u_i(x) \\ < &\sum_{f\in F, h^{A,t}, x\in X} \bar{\mu}(f, h^{A,t}|h_i^{A,t}, \tilde{m}_{i,t}) \Pr^{\sigma'_i,\sigma^*_{-i}}\left(x|\phi_{f,h^{A,t}}, h^{A,t}, \tilde{m}_{i,t}\right) u_i(x). \end{aligned} \tag{14}$$

By definition of $\bar{\mu}$,

$$\bar{\mu}(h_i^{A,t}|\mathring{h}_i^{A,t}, \tilde{m}_{i,t}) = \lim_k \frac{\bar{\mu}^k(h_i^{A,t}|\mathring{h}_i^{A,t})\sigma_i^A(\tilde{m}_{i,t}|h_i^{A,t})}{\sum_{h_i'^{A,t}} \bar{\mu}^k(h_i'^{A,t}|\mathring{h}_i^{A,t})\sigma_i^A(\tilde{m}_{i,t}|h_i'^{A,t})}.$$

Hence, $\bar{\mu}(h_i^{A,t}|\mathring{h}_i^{A,t}, \tilde{m}_{i,t}) > 0$ implies $\tilde{m}_{i,t} \in \operatorname{supp} \sigma_i^A(h_i^{A,t})$.

Note that $\phi_{f,h^{A,t}}$ is constructed so that the conditional distribution of x given f, $h^{A,t}$, and $\tilde{m}_{i,t}$ (when players follow σ^*) is the same as the conditional distribution of x given (f, σ) and realization $a_{i,t} = \tilde{m}_{i,t} \in \operatorname{supp} \sigma_i^A(h_i^{A,t})$. Hence, in order to prove (13), it suffices to show that there exists a strategy in G^* that attains the expected payoff in the second line of (14).

The argument is similar to the proof of Proposition 3. Suppose that in game G^k player i can additionally observe the fictitious information $h_i^{A,t}$, $(r_{i,\tau})_{\tau=t+1}^T$, and $(m_{i,\tau})_{\tau=t+1}^T$. We can view any strategy σ_i' in G^k as such an extended strategy, one that simply ignores the fictitious information. Now, as in the proof of Proposition 3, there is another extended strategy σ_i'' that does not condition on $(\tilde{m}_{i,\tau})_{\tau\geq t}$ and satisfies $\rho^{\sigma_i',\sigma_{-i},f} = \rho^{\sigma_i'',\sigma_{-i},f}$ for each f (conditional on $h_i^{A,t}, \tilde{m}_{i,t}$). And, as in the same proof, there is yet another extended strategy $\sigma_i''' \in \bar{\Sigma}_i$ that does also not condition on $(\tilde{r}_{i,\tau})_{\tau\geq t}$ and satisfies $\rho^{\sigma_i',\sigma_{-i},f} = \rho^{\sigma_i''',\sigma_{-i},f}$. We can now view σ_i''' as a strategy in G, and it attains the higher expected payoff in (14). This contradicts (11), and therefore establishes (13). ■

Given (13), we establish (12) by applying Kuhn's theorem and summing over payoff-irrelevant information.

Proof of Lemma 5. By Kuhn's theorem, for each k, there exists a mixed mediation strategy $\mu^k \in \Delta(F^*)$ such that

$$\sum_{f\in F, h^{A,t}\in H^{A,t}} \bar{\mu}^k\left(f, h^{A,t}, \tilde{m}_t|x^t, s_t\right) \Pr^{\sigma}\left(x|\phi_{f,h^{A,t}}, h^{A,t}, \tilde{m}_t\right) = \sum_{\hat{f}\in F^*} \mu^k\left(\hat{f}|x^t, s_t\right) \Pr^{\sigma}\left(x|\hat{f}, x^t, s_t\right)$$

for each (x^t, s_t), $\tilde{m}_t$, x, and $\sigma \in \Sigma^*$.[24] By Bayes rule, this implies

$$\begin{aligned} & \sum_{f,h^{A,t}} \lim_k \bar{\mu}^k(f, h^{A,t}|x_i^t, s_{i,t}, \tilde{m}_{i,t}) \Pr^\sigma\left(x|\phi_{f,h^{A,t}}, h^{A,t}, \tilde{m}_{i,t}\right) \\ = & \sum_{\hat{f}\in F^*} \lim_k \hat{\mu}^k(\hat{f}, x^t, s_t|x_i^t, s_{i,t}, \tilde{m}_{i,t}) \Pr^\sigma\left(x|\hat{f}, x^t, s_t\right), \end{aligned} \tag{15}$$

where $\hat{\mu}^k$ is determined by $\bar{\mu}^k$, μ^k, and Bayes' rule:

$$\hat{\mu}^k(\hat{f}, x^t, s_t|x_i^t, s_{i,t}, \tilde{m}_{i,t}) = \frac{\bar{\mu}^k\left(x^t, s_t\right)\mu^k\left(\hat{f}|x^t, s_t\right)}{\sum_{\substack{\hat{f}', x'^t_{-i}, s'_{-i,t}: \\ \hat{f}'_i\left(x_i^t, x'^t_{-i}, s_{i,t}, s'_{-i,t}\right) = \tilde{m}_{i,t}}} \bar{\mu}^k\left(x_i^t, x'^t_{-i}, s_{i,t}, s'_{-i,t}\right)\mu^k\left(\hat{f}'|x_i^t, x'^t_{-i}, s_{i,t}, s'_{-i,t}\right)}.$$

Define $\hat{\mu} = \lim_{k\to\infty} \hat{\mu}^k$. By (15), for each $\sigma \in \Sigma^*$ we have

$$\begin{aligned} & \sum_{f,h^{A,t},x} \lim_k \bar{\mu}^k(f, h^{A,t}|x_i^t, s_{i,t}, \tilde{m}_{i,t}) \Pr^\sigma\left(x|\phi_{f,h^{A,t}}, h^{A,t}, \tilde{m}_{i,t}\right) u_i(x) \\ = & \sum_{\hat{f}\in F^*(x^t,s_t),x} \lim_k \hat{\mu}^k(\hat{f}, x^t, s_t|x_i^t, s_{i,t}, \tilde{m}_{i,t}) \Pr^\sigma\left(x|\hat{f}, x^t, s_t\right) u_i(x). \end{aligned}$$

Hence, (13) implies (12). ■

7.5 Proof of Propositions 1 and 4

Proposition 1: By Lemma 3, for each SCE μ, there exists $\rho \in CPPBE(C^*)$ with $\rho = \rho^{\sigma^*,\mu}$. Conversely, take $\rho \in CPPBE(C^*)$. By Lemma 4, players do not take codominated actions at any history. Since every CPPBE is a NE, Proposition 3 implies that there exists a fully canonical NE with outcome ρ where players do not take codominated actions at any history. Hence, by Proposition 2, there exists a SCE μ with $\rho = \rho^{\sigma^*,\mu}$.

Proposition 4: Since Lemma 4 does not require that $C = C^*$, the same argument just given implies that, for any C and $\rho \in CPPBE(C)$, there exists a SCE μ with $\rho = \rho^{\sigma^*,\mu}$. Hence, by Lemma 3, there exists a canonical CPPBE with outcome ρ.

8 Proof of Proposition 5

While $\bigcup_{C\in\mathcal{C}} SE(C) \supsetneq SE(C^*)$ implies that the communication RP does not hold for SE, for expositional clarity we prove the proposition in two steps: we first prove that, in the opening example, the outcome distribution $\frac{1}{2}(A, A, N) + \frac{1}{2}(B, B, N)$ is implementable in non-canonical SE but not in canonical SE; and then we extend the example to prove that $\bigcup_{C\in\mathcal{C}} SE(C) \supsetneq SE(C^*)$.

[24] On the right-hand side, we omit the period t recommendation $\tilde{m}_t$, as it equals $\hat{f}(x^t, s_t)$.

8.1 Failure of the Communication RP in the Opening Example

Possibility for Non-Canonical SE By Propositions 1 and 2 and Theorem 7, we need only construct a canonical NE that implements $\frac{1}{2}(A, A, N) + \frac{1}{2}(B, B, N)$ in which players do not take codominated actions at any history. Such a NE is: the mediator recommends $m_1 = A$ and $m_1 = B$ with equal probability, plays $m_0 = m_1$, and recommends $m_2 = N$ if $s = 0$ and $m_2 = P$ if $s = 1$. Note that each $a_1 \in \{A, B\}$ and N are never codominated, and $a_2 = P$ is not codominated after $s = 1$ as P is optimal if $(a_1, \theta) = (C, p)$.

Impossibility for Canonical SE Since $a_1 = C$ is strictly dominated, if a canonical SE implements $\frac{1}{2}(A, A, N) + \frac{1}{2}(B, B, N)$, the mediation range $Q_1(\emptyset)$ must be equal $\{A, B\}$. That is, the mediator can never recommend $m_1 = C$ (even as the result of a tremble).

Note that, for each strategy for the mediator and player 2 and each realization of $\left(m_1, \hat{a}_1, \hat{\theta}, s\right)$, the resulting probability $\Pr\left(m_2 = P | m_1, a_1, \theta, \hat{a}_1, \hat{\theta}, s\right)$ does not depend on (a_1, θ). Now, conditional on reaching history $(m_1, a_1 = C, \theta)$, player 1 chooses her report $\left(\hat{a}_1, \hat{\theta}\right)$ to minimize $\Pr(m_2 = P)$ (since, in a canonical equilibrium, $a_2 = P$ iff $m_2 = P$). Since $a_1 = C$ implies $s = 1$, and when $a_1 = C$ player 1 must be willing to report $\left(\hat{a}_1, \hat{\theta}\right) = (C, \theta)$ for each value of θ, we have

$$\Pr\left(m_2 = P | m_1, \hat{a}_1 = C, \hat{\theta} = n, s = 1\right) = \Pr\left(m_2 = P | m_1, \hat{a}_1 = C, \hat{\theta} = p, s = 1\right).$$

In addition, if a canonical SE implements $\frac{1}{2}(A, A, N) + \frac{1}{2}(B, B, N)$, it must satisfy

$$\Pr(m_2 = P | m_1, \hat{a}_1 = C, s = 1) > 0$$

for each $m_1 \in \{A, B\}$. Otherwise, given that player 2 never plays $a_2 = P$ with positive probability when $s = 0$ (since $s = 0$ implies $a_1 \neq C$), player 1 could guarantee a payoff of $\frac{1}{2}$ by playing A and B with equal probability and reporting $\hat{a}_1 = C$. Hence, for each $m_1 \in \{A, B\}$,

$$\Pr\left(m_2 = P | m_1, \hat{a}_1 = C, \hat{\theta} = n, s = 1\right) = \Pr\left(m_2 = P | m_1, \hat{a}_1 = C, \hat{\theta} = p, s = 1\right) > 0.$$

Since player 1 honestly reports each (a_1, θ) in a canonical SE,

$$\Pr(m_2 = P | m_1, a_1 = C, \theta = n, s = 1) = \Pr(m_2 = P | m_1, a_1 = C, \theta = p, s = 1) > 0.$$

Hence, along any sequence of completely mixed profiles indexed by k converging to the equilibrium,

$$\lim_k \Pr^k(m_2 = P | m_1, a_1 = C, \theta = n, s = 1) = \lim_k \Pr^k(m_2 = P | m_1, a_1 = C, \theta = p, s = 1) > 0. \tag{16}$$

Therefore,

$$\begin{aligned}
& \Pr\left((a_1,\theta)=(C,p)\,|s=1,m_2=P\right) \\
= & \lim_k \frac{\Pr^k\left((a_1,\theta)=(C,p),s=1,m_2=P\right)}{\Pr^k\left(s=1,m_2=P\right)} \\
= & \lim_k \frac{\Pr^k\left(m_2=P|\,(a_1,\theta)=(C,p),s=1\right)\Pr^k\left((a_1,\theta)=(C,p),s=1\right)}{\left(\begin{array}{c}\Pr^k\left(m_2=P|\,(a_1,\theta)=(C,p),s=1\right)\Pr^k\left((a_1,\theta)=(C,p),s=1\right)\\ +\Pr^k\left(m_2=P|\,(a_1,\theta)\neq(C,p),s=1\right)\Pr^k\left((a_1,\theta)\neq(C,p),s=1\right)\end{array}\right)} \\
\leq & \lim_k \frac{\Pr^k\left(m_2=P|\,(a_1,\theta)=(C,p),s=1\right)\Pr^k\left((a_1,\theta)=(C,p),s=1\right)}{\left(\begin{array}{c}\Pr^k\left(m_2=P|\,(a_1,\theta)=(C,p),s=1\right)\Pr^k\left((a_1,\theta)=(C,p),s=1\right)\\ +\Pr^k\left(m_2=P|\,(a_1,\theta)=(C,n),s=1\right)\Pr^k\left((a_1,\theta)=(C,n),s=1\right)\end{array}\right)} \\
= & \lim_k \frac{\Pr^k\left(m_2=P|\,(a_1,\theta)=(C,p),s=1\right)}{\Pr^k\left(m_2=P|\,(a_1,\theta)=(C,p),s=1\right)+\Pr^k\left(m_2=P|\,(a_1,\theta)=(C,n),s=1\right)} \\
= & \frac{1}{2},
\end{aligned}$$

where the second-to-last line follows because $\theta = n$ or p with equal probability, independent of a_1 and s, and the last line follows since (16) holds for each $m_1 \in \{A, B\}$, which are the only possible values for m_1. This implies that player 2 will not follow recommendation $m_2 = P$ when $s = 1$ in any canonical SE. Hence, $a_2 = P$ cannot be played with positive probability at any history in any canonical SE. Given this, player 1 can guarantee a payoff of $\frac{1}{2}$ by mixing A and B, so $\frac{1}{2}(A, A, N) + \frac{1}{2}(B, B, N)$ cannot implemented.

8.2 Extending the Example

In the extended example, there are four players (in addition to the mediator) and four periods, with the following timing:

Period 1. No signals are observed. Player 1 takes an action $a_1 \in \{A_1, B_1\}$.

Period 2. The mediator observes a_1. Player 2 takes $a_2 \in \{A_2, B_2, C_2\}$ and player 3 takes $a_3 \in \{A_3, B_3\}$.

Period 3. Player 2 observes $\theta \in \{n, p\}$ with equal probability. The mediator takes $a_0 \in \{A_0, B_0\}$.

Period 4. The mediator and player 4 observe $s \in \{0, 1\}$, where $s = 0$ iff $a_1 = A_1$ and $a_2 \neq a_0$. Player 4 takes $a_4 \in \{N, P\}$.

Player 1's payoff equals $1_{\{a_1=B_1\}} - 1_{\{a_2=C_2\}}1_{\{a_4=P\}}$. Player 2's payoff is given by

$a_1 = A_1$:

	A_0	B_0
A_2	$0-1_{\{a_4=P\}}$	$1-1_{\{a_4=P\}}$
B_2	$1-1_{\{a_4=P\}}$	$0-1_{\{a_4=P\}}$
C_2	$-3-1_{\{a_4=P\}}$	$-3-1_{\{a_4=P\}}$

$a_1 = B_1$:

	A_0	B_0
A_2	$1-1_{\{a_4=P\}}$	$1-1_{\{a_4=P\}}$
B_2	$1-1_{\{a_4=P\}}$	$1-1_{\{a_4=P\}}$
C_2	0	0

.

Player 3's payoff is constant. Player 4's payoff equals $1_{\{(a_1,a_2,\theta)\neq(A_1,C_2,p)\}}1_{\{a_4=N\}}$.

Note that we have omitted time subscripts on A_i since each player moves only once. We do the same for M_i.

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, $\frac{1}{2}(A_2, A_3, A_0) + \frac{1}{2}(B_2, B_3, B_0)$ is played in periods 2 and 3, (iii) when B_1 is played in period 1, (A_2, A_3, A_0) is played in periods 2 and 3, and (iv) N is played in period 4. We claim that this distribution is implementable in SE, but not with $C = C^*$.

Possibility for $C \neq C^*$ Again, it suffices to implement the target distribution in a canonical NE in which players do not take codominated actions at any history. Consider the following mediator strategy: The mediator draws $m_1 \in \{A_1, B_1\}$ with equal probability.

When $m_1 = a_1 = A_1$, the mediator draws $m_2 \in \{A_2, B_2\}$ with equal probability, and recommends $m_0 = m_2 = m_3$. If $s = 0$, he recommends $m_4 = N$; if $s = 1$, he recommends $m_4 = P$.

When $m_1 = A_1$ but $a_1 = B_1$, the mediator recommends $m_2 = C_2$, $m_3 = A_3$, and $m_4 = P$.

When $m_1 = B_1$ (regardless of a_1), the mediator recommends $m_2 = A_2$, $m_3 = A_3$, $m_0 = A_0$, and $m_4 = N$.

It is straightforward to check that this is a NE. Moreover, no codominated actions are recommended: For player 4, after $s = 0$, P is codominated; and after $s = 1$, no action is codominated, since $(a_1, a_2, \theta) \neq (A_1, C_2, p)$ is feasible. Given this, no action is codominated for player 2, as each $a_2 \in \{A_2, B_2\}$ can be optimal after $a_1 = A_1$, and $a_2 = C_2$ is optimal after $a_1 = B_1$ when $a_4 = P$ is anticipated. Given that $a_2 = C_2$ and $a_4 = P$ are not codominated, neither is $a_1 = A_1$.

Impossibility for $C = C^*$ Suppose towards a contradiction that such a SE exists. In what follows, each fraction p/q should be read as $\lim_{k\to\infty} p^k/q^k$, where $p^k, q^k > 0$ denote probabilities along a sequence of strategy profiles converging to the equilibrium.

For each player i and action a_i that is played with positive probability in the target distribution, assume without loss that a_i is played with positive probability after $m_1 = a_1$. Moreover, since the on-path actions of players 2 and 3 must be perfectly correlated, it is without loss to assume that, for $i \in \{2, 3\}$, $a_i \in \{A_i, B_i\}$ is played with probability 1 after $m_i = a_i$. Further, to deter a deviation by player 1 following $m_1 = A_1$, player 2 must play $a_2 = C_2$ with probability 1 after some message, which without loss we take to be $m_2 = C_2$. Since player 3 is indifferent among all outcomes, we can also let $a_3 = m_3$ with probability 1. Finally, since player 4 moves last, the usual static revelation principle argument implies that we can let $a_4 = m_4$ with probability 1.

Note that C_2 is strictly dominated conditional on $a_1 = A_1$ and weakly dominated conditional on $a_1 = B_1$. Since player 2 is willing to take C_2 after $m_2 = C_2$, we have $\Pr(a_1 = B_1 | m_2 = C_2) = 1$. Since a_1 and a_2 are independent conditional on m_2, this implies $\Pr(a_1 = B_1 | m_2 = C_2, a_2 = A_2) = 1$. Hence, if player 2 trembles to $a_2 = A_2$ after $m_2 = C_2$, she believes that $a_1 = B_1$ with probability 1, and she therefore chooses her report $\left(\hat{a}_2, \hat{\theta}\right)$ to minimize the probability that $a_4 = P$. Since player 2 can always report as if she took

$a_2 = C_2$, this implies that

$$\begin{aligned} & \Pr\left(a_1 = B_1, a_4 = P | m_2 = C_2, a_2 = A_2\right) \\ \leq \; & \Pr\left(a_1 = B_1, a_4 = P | m_2 = C_2, a_2 = C_2\right). \end{aligned}$$

Note that if $\Pr\left(a_1 = B_1, a_4 = P | m_2 = C_2, a_2 = A_2\right) < 1$ then player 2 would deviate to A_2 after $m_2 = C_2$. So this probability must equal 1, and hence

$$\Pr\left(a_1 = B_1, a_4 = P | m_2 = C_2, a_2 = C_2\right) = 1.$$

Since $a_1 = B_1$ implies $s = 1$, we have

$$\Pr\left(a_1 = B_1, m_4 = P, s = 1 | m_2 = C_2, a_2 = C_2\right) = 1.$$

Finally, since $a_2 = C_2$ with probability 1 after $m_2 = C_2$, we have

$$\Pr\left(a_1 = B_1, a_2 = C_2, m_4 = P, s = 1 | m_2 = C_2\right) = 1. \tag{17}$$

On the other hand, since player 4 to be willing to take N after $s = 1$ and $m_4 = P$, we have $\Pr\left(a_1 = A_1, a_2 = C, \theta = p | s = 1, m_4 = P\right) = 1$. In particular,

$$1 = \frac{\left(\begin{array}{c} \Pr\left(m_2 = C_2\right) \Pr\left(a_1 = A_1, a_2 = C_2, \theta = p, m_4 = P, s = 1 | m_2 = C_2\right) \\ + \sum_{m_2 \neq C_2} \Pr\left(m_2\right) \Pr\left(a_1 = A_1, a_2 = C_2, \theta = p, m_4 = P, s = 1 | m_2\right) \end{array} \right)}{\left(\begin{array}{c} \Pr\left(m_2 = C_2\right) \sum_{a_1, a_2} \Pr\left(a_1, a_2, m_4 = P, s = 1 | m_2 = C_2\right) \\ + \sum_{m_2 \neq C_2} \Pr\left(m_2\right) \sum_{a_1, a_2} \Pr\left(a_1, a_2, m_4 = P, s = 1 | m_2\right) \end{array} \right)}.$$

Since $(a + c) / (b + d) < (a/b) + (c/d)$ for all positive numbers a, b, c, d, the right-hand side is no more than

$$\frac{\Pr\left(a_1 = A_1, a_2 = C_2, \theta = p, m_4 = P, s = 1, | m_2 = C_2\right)}{\sum_{a_1, a_2} \Pr\left(a_1, a_2, m_4 = P, s = 1 | m_2 = C_2\right)} + \sum_{m_2 \neq C_2} \frac{\Pr\left(a_1 = A_1, a_2 = C_2, \theta = p, m_4 = P, s = 1 | m_2\right)}{\sum_{a_1, a_2} \Pr\left(a_1, a_2, m_4 = P, s = 1 | m_2\right)}.$$

Note that

$$\begin{aligned} & \Pr\left(a_1 = A_1, a_2 = C_2, \theta = p, m_4 = P, s = 1, | m_2 = C_2\right) \\ \leq \; & \Pr\left(a_1 = A_1, a_2 = C_2, m_4 = P, s = 1, | m_2 = C_2\right) = 0 \quad \text{(by (17))}. \end{aligned}$$

Hence,

$$
\begin{aligned}
1 &= \sum_{m_2 \neq C_2} \frac{\Pr(a_1 = A_1, a_2 = C_2, \theta = p, m_4 = P, s = 1 | m_2)}{\sum_{a_1, a_2} \Pr(a_1, a_2, m_4 = P, s = 1 | m_2)} \\
&= \sum_{m_2 \neq C_2} \frac{\Pr(a_1 = A_1, a_2 = C_2, \theta = p, m_4 = P, s = 1 | m_2)}{\Pr(a_1 = A_1, a_2 = C_2, m_4 = P, s = 1 | m_2)} \\
&= \sum_{m_2 \neq C_2} \frac{\Pr(\theta = p, m_4 = P | a_1 = A_1, m_2, a_2 = C_2)}{\Pr(m_4 = P | a_1 = A_1, m_2, a_2 = C_2)}, \qquad (18)
\end{aligned}
$$

where the second line drops the event $(a_1, a_2) \neq (A_1, C_2)$ from the denominator (giving a necessary condition) and the third line uses the fact that $a_2 = C_2$ implies $s = 1$.

Now, after $a_2 = C_2$, player 2 is strictly better off if player 4 takes N if $a_1 = A_1$, and player 2 is indifferent if $a_1 = B_1$. Moreover, $\Pr(a_1 = A_1 | m_2) > 0$ for each $m_2 \neq C_2$. Hence, each $m_2 \neq C_2$, after $(m_2, a_2 = C_2)$ player 2 chooses her report $\left(\hat{a}_2, \hat{\theta}\right)$ to minimize the probability that $a_4 = P$, and hence the probability $m_4 = P$ (since $a_4 = m_4$ with probability 1), independent of θ. Therefore, for each $m_2 \neq C_2$,

$$
\Pr(\theta = P, m_4 = P | a_1 = A_1, m_2, a_2 = C_2) = \Pr(\theta = P) \Pr(m_4 = P | a_1 = A_1, m_2, a_2 = C_2),
$$

and thus

$$
\sum_{m_2 \neq C_2} \frac{\Pr(\theta = p, m_4 = P | a_1 = A_1, m_2, a_2 = C_2)}{\Pr(m_4 = P | a_1 = A_1, m_2, a_2 = C_2)} = \frac{1}{2}.
$$

This contradicts (18).

9 Results for SE

This section contains our analysis of SE, culminating in the proofs of Propositions 6 and 7.

9.1 Quasi-Strategies and Quasi-SE

As we did for CPPBE, we begin by introducing notions of quasi-strategy and quasi-equilibrium. For each player i, the definition of a quasi-strategy (χ_i, J_i) is the same as in Section 7.1. For the mediator, a *quasi-strategy* (ψ, K) consists of

1. A set of histories $K = \prod_{t=1}^{T} K^t \cup K^{t+}$ with $K^t \subseteq R^{t+1} \times M^t$ and $K^{t+} \subseteq R^{t+1} \times M^{t+1}$ such that (i) for each $(r^{t+1}, m^t) \in K^t$, there exists $\left(r^{T+1}, m^T\right) \in K^T$ that coincides with (r^{t+1}, m^t) up to period t, (ii) for every $(r^{T+1}, m^T) \in K^T$ and every (r^{t+1}, m^t) that coincides with (r^{T+1}, m^T) up to period t, we have $(r^{t+1}, m^t) \in K^t$, and (iii) the same conditions hold for K^{t+}.

2. A function $\psi = \left(\psi^t\right)_{t=1}^{T}$, where $\psi^t : K^t \to \Delta(M_t)$.

We write $h^t \in (J, K)$ if $h_i^t \in J_i^t$ and $(r^t, m^t) \in K^{t-1,+}$ (or, $(r^{t+1}, m^t) \in K^t$ if h^t includes r_t, or $(r^{t+1}, m^{t+1}) \in K^{t+}$ if h^t includes m_t), where h_i^t and (r^{t+1}, m^t) and are the relevant elements of h^t. A strategy profile (σ, ϕ) has *full support on* (J, K) if (i) for each $(r^{t+1}, m^t) \in K^t$, $\psi(m_t|r^{t+1}, m^t) > 0$ if and only if $(r^{t+1}, m^{t+1}) \in K^{t+}$, (ii) for each $h_i^{R,t} \in J_i^{R,t}$, $\sigma_i^R(r_{i,t}|h_i^{R,t}) > 0$ if and only if $(h_i^{R,t}, r_{i,t}) \in J_i^{R,t+}$, and (iii) for each $h_i^{A,t} \in J_i^{A,t}$, $\sigma_i^A(a_{i,t}|h_i^{A,t}) > 0$ if and only if $(h_i^{A,t}, a_{i,t}) \in J_i^{A,t+}$.

We say a quasi-strategy profile (χ, ψ, J, K) is *valid* if

1. For each $h^t \in (J, K)$, $i \neq 0$, σ_i, and $h_{-i}^{T+1} \notin (J_{-i}, K)$, we have $\Pr^{(\sigma_i, \chi_{-i}, \psi)}\left(h_{-i}^{T+1}|h^t\right) = 0$.
2. For each $i \neq 0$, $h_i^t \in J_i^t$, and full-support strategy profile (σ, ϕ) over (J, K), we have $\Pr^{(\sigma,\phi)}\left(h_i^t\right) > 0$.
3. For each full-support strategy profile (σ, ϕ) over (J, K), we have $\Pr^{(\sigma,\phi)}\left(h^{T+1} \in (J, K)\right) = 1$.

Finally, a *quasi-SE* $(\chi, \psi, J, K, \beta)$ is a valid quasi-strategy profile (χ, ψ, J, K) and a belief system β such that:

1. *[Sequential rationality of reports]* For all $i \neq 0, t, \sigma_i' \in \Sigma_i$, and $h_i^{R,t} \in J_i^{R,t}$, we have
$$\sum_{h^{R,t} \in (J,K)} \beta\left(h^{R,t}|h_i^{R,t}\right) \bar{u}_i\left(\chi, \psi|h^{R,t}\right) \geq \sum_{h^{R,t} \in (J,K)} \beta\left(h^{R,t}|h_i^{R,t}\right) \bar{u}_i\left(\sigma_i', \chi_{-i}, \psi|h^{R,t}\right).$$
2. *[Sequential rationality of actions]* For all $i \neq 0, t, \sigma_i' \in \Sigma_i$, and $h_i^{A,t} \in J_i^{A,t}$, we have
$$\sum_{h^{R,t} \in (J,K)} \beta\left(h^{A,t}|h_i^{A,t}\right) \bar{u}_i\left(\chi, \psi|h^{A,t}\right) \geq \sum_{h^{A,t} \in (J,K)} \beta\left(h^{A,t}|h_i^{A,t}\right) \bar{u}_i\left(\sigma_i', \chi_{-i}, \psi|h^{A,t}\right).$$
3. *[Kreps-Wilson consistency]* There exists a sequence of strategy profiles $\left(\sigma^k, \phi^k\right)_{k=1}^{\infty}$ such that $\left(\sigma^k, \phi^k\right)$ has full support over (J, K) for each k, $\lim_{k\to\infty}\left(\sigma^k, \phi^k\right)\left(h^t\right) = (\chi, \psi)$ for each $h^t \in (J, K)$, and
$$\beta\left(h^t|h_i^t\right) = \lim_{k\to\infty} \frac{\Pr^{\sigma^k,\phi^k}\left(h^t\right)}{\Pr^{\sigma^k,\phi^k}\left(h_i^t\right)}$$
for each $h_i^t \in J_i^t$ and $h^t \in (J, K)$. Note that validity of (χ, ψ, J, K) implies $\Pr^{\left(\sigma^k,\phi^k\right)}\left(h_i^t\right) > 0$ for each $h_i^t \in J_i^t$.

As for CPPBE, it is without loss to consider quasi-SE rather than fully specified SE.

Lemma 7 *For any game $G = (\Gamma, C)$ and any outcome distribution $\rho \in \Delta(X)$, we have $\rho \in SE(C)$ if and only if $\rho = \rho^{(\chi,\psi)}$ for some quasi-SE profile $(\chi, \psi, J, K, \beta)$.*

Proof. The proof is similar to the proof of Lemma 1.

Fix a game G. One direction is immediate: If (σ, ϕ, β) is a SE in G, then define $(\chi, \psi) = (\sigma, \phi)$ and $J_i^t = H_i^t$, $K^t = R^t \times M^{t-1}$, and $K^{t+} = R^t \times M^t$. Then, $(\chi, \psi, J, K, \beta)$ is quasi-SE with $\rho^{(\chi,\psi)} = \rho^{(\sigma,\phi)}$.

For the converse, fix a quasi-SE $(\chi, \psi, J, K, \beta)$ and a converging sequence $\left(\tilde{\sigma}^k, \tilde{\phi}^k\right)_k$ with full support over (J, K). For each k,

$$\varepsilon_k = \min_{\left(h_i^{R,t}, r_{i,t}\right) \in J_i^{R,t+}, \left(h_i^{A,t}, a_{i,t}\right) \in J_i^{A,t+}, \left(r^{t+1}, m^{t+1}\right) \in K^{t+}} \min\{\tilde{\sigma}_i^{R,k}(r_t|h_i^{R,t}), \tilde{\sigma}_i^{A,k}(a_{i,t}|h_i^{A,t}), \tilde{\phi}^k(m_t|r^{t+1}, m^t)\}.$$

Consider the following auxiliary game $\left(\Gamma^k, C\right)$:

1. The mediator chooses probability distributions $\phi^k(\cdot|r^{t+1}, m^t) \in \Delta(M_t)$ for each $(r^{t+1}, m^t) \in (R^{t+1} \times M^t) \setminus K^t$. At histories $(r^{t+1}, m^t) \in K^t$, the mediator has no choice to make, and we set $\phi^k(\cdot|r^{t+1}, m^t) = \tilde{\phi}^k(\cdot|r^{t+1}, m^t)$.
2. Each player i chooses probability distributions $\sigma_i^{R,k}(\cdot|h_i^{R,t}) \in \Delta(R_{i,t})$ and $\sigma_i^{A,k}(\cdot|h_i^{A,t}) \in \Delta(A_{i,t})$ for each $t, h_i^{R,t} \in H_i^{R,t} \backslash J_i^{R,t}, h_i^{A,t} \in H_i^{A,t} \backslash J_i^{A,t}$. At histories $h_i^t \in J_i^t$, player i has no choice to make, and we set $\sigma_i^{R,k}(\cdot|h_i^{R,t}) = \tilde{\sigma}_i^{R,k}\left(\cdot|h_i^{R,t}\right)$ and $\sigma_i^{A,k}(\cdot|h_i^{A,t}) = \tilde{\sigma}_i^{A,k}\left(\cdot|h_i^{A,t}\right)$.
3. Given $\left(\sigma^k, \phi^k\right)$, the distribution of H^{T+1} is determined recursively as follows: Given $h^{R,t} \in H^{R,t}$, each $r_{i,t} \in R_{i,t}$ is drawn independently across players with probability

$$\begin{cases} \left(1 - \frac{\varepsilon_k}{k}\left|\left\{\tilde{r}_{i,t} \in R_{i,t} : \left(h_i^{R,t}, \tilde{r}_{i,t}\right) \notin J_i^{R,t+}\right\}\right|\right) \tilde{\sigma}_i^{R,k}(r_{i,t}|h_i^{R,t}) & \text{if } h_i^{R,t} \in J_i^{R,t} \wedge \left(h_i^{R,t}, r_{i,t}\right) \in J_i^{R,t+}, \\ \frac{\varepsilon_k}{k} & \text{if } h_i^{R,t} \in J_i^{R,t} \wedge \left(h_i^{R,t}, r_{i,t}\right) \notin J_i^{R,t+}, \\ \left(1 - \frac{\varepsilon_k}{k}\left|R_{i,t}\right|\right) \sigma_i^{R,k}(r_{i,t}|h_i^{R,t}) + \frac{\varepsilon_k}{k} & \text{if } h_i^{R,t} \notin J_i^{R,t}. \end{cases}$$

Given $(r^t, m^{t-1}) \in R^t \times M^{t-1}$, each m_t is drawn with probability

$$\begin{cases} (1 - \frac{\varepsilon_k}{k}|\{\tilde{m}_t \in M_t \colon (r^{t+1}, m^t, \tilde{m}_t) \notin K^{t+}\}|)\tilde{\phi}^k(m_t|r^{t+1}, m^t) & \text{if } (r^{t+1}, m^t) \in K^t \wedge (r^{t+1}, m^{t+1}) \in K^{t+}, \\ \frac{\varepsilon_k}{k} & \text{if } (r^{t+1}, m^t) \in K^t \wedge (r^{t+1}, m^{t+1}) \notin K^{t+}, \\ \left(1 - \frac{\varepsilon_k}{k}\left|M_t\right|\right) \phi^k(m_t|r^{t+1}, m^t) + \frac{\varepsilon_k}{k} & \text{if } (r^{t+1}, m^t) \notin K^t. \end{cases}$$

Given $h^{A,t} \in H^{A,t}$, each $a_{i,t} \in A_{i,t}$ is drawn independently across players with probability

$$\begin{cases} \left(1 - \frac{\varepsilon_k}{k}\left|\left\{\tilde{a}_{i,t} \in A_{i,t} : \left(h_i^{A,t}, \tilde{a}_{i,t}\right) \notin J_i^{A,t+}\right\}\right|\right) \tilde{\sigma}_i^{A,k}(a_{i,t}|h_i^{A,t}) & \text{if } h_i^{A,t} \in J_i^{A,t} \wedge \left(h_i^{A,t}, a_{i,t}\right) \in J_i^{A,t+}, \\ \frac{\varepsilon_k}{k} & \text{if } h_i^{A,t} \in J_i^{A,t} \wedge \left(h_i^{A,t}, a_{i,t}\right) \notin J_i^{A,t+}, \\ \left(1 - \frac{\varepsilon_k}{k}\left|A_{i,t}\right|\right) \sigma_i^{A,k}(a_{i,t}|h_i^{A,t}) + \frac{\varepsilon_k}{k} & \text{if } h_i^{A,t} \notin J_i^{A,t}. \end{cases}$$

Then, given $a_t \in A_t$, each $s_{t+1} \in S_{t+1}$ is drawn with probability $p\left(s_{t+1}|\mathring{h}^{A,t}, a_t\right)$.

The rest of the proof is the same as the proof of Lemma 1. ■

In addition, as for CPPBE, restricting the mediator's messages to lie in some mediation range does not expand the outcome set.

Lemma 8 *For any game G, we have $SE(G) = \bigcup_{Q \in \mathcal{Q}(G)} SE(G|_Q)$.*

Proof. The same as the proof of Lemma 2. ■

Finally, in the proofs of Propositions 6 and 7, it will be convenient to describe the mediator's strategy in terms of first choosing a period-t "state" $\theta_t \in \Theta_t$ as a function of the mediator's history (r^t, m^t) and the past states $\theta^t = (\theta_1, \ldots, \theta_{t-1})$, and then choosing period-t messages m_t as a function of the vector $\left(\theta^t, r^t, m^t, \theta_t, r_t\right)$. When convenient, we will include these states as part of the mediator's history. In such cases, the realizations of past θ^t are included in K.

9.2 Proof of Proposition 6

We first slightly strengthen Proposition 3. Fixing a base game Γ, let $\Sigma(C)$ denote the strategy set in game $G = (\Gamma, C)$.

Lemma 9 *For any game $G = (\Gamma, C)$, if (σ, ϕ) is a NE then $\rho^{(\sigma,\phi)}$ is the outcome distribution of a canonical NE $\left(\tilde{\sigma}, \tilde{\phi}\right)$ such that, for all $i \neq 0$, $\bigcup_{\sigma'_{-i} \in \Sigma_{-i}} \operatorname{supp} \rho_i^{\left(\sigma_i, \sigma'_{-i}, \phi\right)} \supseteq \bigcup_{\tilde{\sigma}'_{-i} \in \Sigma_{-i}(C^*)} \operatorname{supp} \rho_i^{\left(\tilde{\sigma}_i, \tilde{\sigma}'_{-i}, \tilde{\phi}\right)}$.*

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

Proof. We show that, for the strategy profile $\left(\tilde{\sigma}, \tilde{\phi}\right)$ constructed in the proof of Proposition 3, $\bigcup_{\sigma'_{-i} \in \Sigma_{-i}} \operatorname{supp} \rho_i^{\sigma_i, \sigma'_{-i}, \phi} \supseteq \bigcup_{\tilde{\sigma}'_{-i} \in \Sigma_{-i}(C^*)} \operatorname{supp} \rho_i^{\tilde{\sigma}_i, \tilde{\sigma}'_{-i}, \tilde{\phi}}$. This follows because, for any $\tilde{\sigma}' \in \Sigma(C^*)$, one can construct a strategy profile $\sigma' \in \Sigma$ such that $\rho_i^{\sigma_i, \sigma'_{-i}, \phi} = \rho_i^{\tilde{\sigma}_i, \tilde{\sigma}'_{-i}, \tilde{\phi}}$ 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 $\tilde{\sigma}_i'^A(s_{i,1}, \tilde{r}_{i,1}, \tilde{m}_{i,1})$.

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 $\sigma_i^R((\tilde{r}_{i,\tau}, r_{i,\tau}, m_{i,\tau}), \tilde{r}_{i,t})_{\tau=1}^{t-1}$.[25] 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 honest and obedient strategy $\tilde{\sigma}_i$ is mapped to the original equilibrium strategy σ_i, so $\bigcup_{\sigma'_{-i} \in \Sigma_{-i}} \rho_i^{\sigma_i, \sigma'_{-i}, \phi} \supseteq \bigcup_{\tilde{\sigma}'_{-i} \in \Sigma_{-i}(C^*)} \rho_i^{\tilde{\sigma}_i, \tilde{\sigma}'_{-i}, \tilde{\phi}}$ and hence $\bigcup_{\sigma'_{-i} \in \Sigma_{-i}} \operatorname{supp} \rho_i^{\sigma_i, \sigma'_{-i}, \phi} \supseteq \bigcup_{\tilde{\sigma}'_{-i} \in \Sigma_{-i}(C^*)} \operatorname{supp} \rho_i^{\tilde{\sigma}_i, \tilde{\sigma}'_{-i}, \tilde{\phi}}$. ■

[25] If $(\tilde{r}_{i,\tau})_{\tau=1}^{t}$ is not a possible history, player i draws $r_{i,t}$ arbitrarily, and similarly for $\tilde{m}_{i,t}$ and $a_{i,t}$ in what follows.

We now prove Claim 1. Consider the fully canonical strategy profile (σ^*, ϕ^*) constructed in the proof of Proposition 3. Let $\chi = \sigma^*$, let

$$J_i^t = \left\{ h_i^t \in H_i^t : \mathring{h}_i^t \in \operatorname{supp} \rho_i^{\sigma^*, \phi^*} \wedge r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau}) \ \forall \tau \leq t-1 \right\} \text{ for all } i, t,$$

let $\psi = \phi^*$, and let

$$K^t = \{ (r^{t+1}, m^t) \in R^{t+1} \times M^t : m_\tau \in \operatorname{supp} \phi^*(r^{\tau+1}, m^\tau) \ \forall \tau \leq t-1 \}.$$

By Lemma 9, $h_i^t \in J_i^t$ iff $\mathring{h}_i^t \in \bigcup_{\tilde{\sigma}'_{-i} \in C^*(\Sigma_{-i})} \operatorname{supp} \rho_i^{(\sigma_i^*, \tilde{\sigma}'_{-i}, \phi^*)}$ and $r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau}) \ \forall \tau \leq t-1$. Therefore, (χ, ψ, J, K) is valid. Moreover, Proposition 3 implies that, for strategy profile (χ, ψ), (3) and (4) are satisfied at all histories $h_i^{R,t}, h_i^{A,t} \in J_i^t$: this follows since these are on-path histories, so ex ante optimality implies sequential rationality. Hence, Claim 1 follows from Lemma 7.

To prove Claim 2, note that the condition $\operatorname{supp} \rho_i^\sigma = \bigcup_{\sigma'_{-i} \in \Sigma_{-i}} \operatorname{supp} \rho_i^{\sigma_i, \sigma'_{-i}}$ is vacuous when $N = 1$. Hence, the result follows from Claim 1.

9.3 Proof of Proposition 7

We must show that $CPPBE(C^*) \subseteq SE(C^{**})$. By Propositions 1 and 2, it suffices to show that $SE(C^{**})$ contains all outcomes that are implementable in a canonical NE in which codominated actions are never recommended at any history.

Under quasi-direct communication, we say that player i is *faithful* at history $h_i^t = (s_i^t, r_i^t, m_i^t, a_i^t)$ if $r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$ for each $\tau < t$ and $a_{i,\tau} = m_{i,\tau}$ for each $\tau < t$ with $m_{i,\tau} \in A_{i,\tau}$ (i.e., with $m_{i,\tau} \neq \star$). That is, player i is faithful at history h_i^t if thus far she has been honest and has followed all action recommendations.

Full-Support Nash Equilibrium Let $(\varepsilon_k)_{k \in \mathbb{N}}$ satisfy $\varepsilon_k \to 0$ and $k (\varepsilon_k)^{NT} / |A^T| \to \infty$. For each k, let σ^k be a NE in the unmediated, ε_k-perturbed game (i.e., the game where each player is constrained to play each action with probability at least ε_k at each information set), such that the sequence $(\sigma^k)_{k \in \mathbb{N}}$ converges to a NE $\hat{\sigma}$ in the unperturbed game. For each $(x_i^t, s_{i,t})$, let $\hat{B}_i(x_i^t, s_{i,t}) = A_{i,t} \setminus \operatorname{supp} \hat{\sigma}_{i,t}(x_i^t, s_{i,t})$ denote the set of actions that are taken at $(x_i^t, s_{i,t})$ only when player i trembles.

We can implement each full-support strategy profile σ^k in the mediated, unperturbed game where the mediator trembles but players are faithful. Suppose the mediator independently performs all randomizations in the limit NE $\hat{\sigma}$ for each player i at the beginning of the game. Denote the resulting mixed strategy for the mediator by $\hat{\mu}$. Note that, since $\hat{\sigma}$ does not have full support, neither does $\hat{\mu}$: that is, we have not yet introduced mediator trembles. It will be convenient to write $f^{<t} \in \operatorname{supp} \hat{\mu}^{<t}$ if there exists $f^{\geq t}$ such that $(f^{<t}, f^{\geq t}) \in \operatorname{supp} \hat{\mu}$, where $f^{<t}$ and $f^{\geq t}$ denote mediation plans before and after period t. Define $f_i \in \operatorname{supp} \hat{\mu}_i$ and $f_i^{<t} \in \operatorname{supp} \hat{\mu}_i^{<t}$ analogously, focusing on the mediation plan for player i.

Now let $(\hat{\sigma}^k, \hat{\mu})$ denote the strategy profile in the mediated game where players are honest and obedient, while trembling uniformly over actions with probability ε_k. Let $\Pr^{\hat{\sigma}^k, \hat{\mu}}$ denote the resulting probability distribution. Note that $\hat{\sigma}^k$ converges to the fully canonical strategy

σ^*. By construction, $(\sigma^*, \hat{\mu}) = \lim_{k\to\infty}\left(\hat{\sigma}^k, \hat{\mu}\right)$ is a quasi-SE, together with $\chi = \sigma^*$, $\psi = \hat{\mu}$, $(J, K) = \{(f, z) : \Pr^{\hat{\sigma}^k,\hat{\mu}}(f, z) > 0\}$,[26] and belief system

$$\beta^{\sigma^*,\hat{\mu}}\left(f, \mathring{h}^{R,t}|h_i^{R,t}\right) = \lim_{k\to\infty} \frac{\hat{\mu}(f)\Pr^{\hat{\sigma}^k,\hat{\mu}}\left(\mathring{h}^{R,t}|f\right)}{\Pr^{\hat{\sigma}^k,\hat{\mu}}\left(h_i^{R,t}\right)}.$$

Since $\hat{\sigma}$ is a NE, for each $h_i^{R,t}$ with $\Pr^{\hat{\sigma}^k,\hat{\mu}}\left(h_i^{R,t}\right) > 0$ and any continuation strategy σ_i' from period t, we have

$$\sum_{f\in\text{supp}\,\hat{\mu},\mathring{h}^{R,t}} \beta^{\sigma^*,\hat{\mu}}\left(f, \mathring{h}^{R,t}|h_i^{R,t}\right)\bar{u}_i\left(\sigma^*|f, \mathring{h}^{R,t}\right) \geq \sum_{f\in\text{supp}\,\hat{\mu},\mathring{h}^{R,t}} \beta^{\sigma^*,\hat{\mu}}\left(f, \mathring{h}^{R,t}|h_i^{R,t}\right)\bar{u}_i\left(\sigma_i', \sigma_{-i}^*|f, \mathring{h}^{R,t}\right), \tag{19}$$

where $\bar{u}_i\left(\sigma^*|f, \mathring{h}^{R,t}\right)$ is defined as $\bar{u}_i\left(\sigma^*|f, h^{R,t}\right)$ with the history $h^{R,t}$ corresponding to the payoff-relevant history $\mathring{h}^{R,t}$ with honest reports and messages given by f.

Rationalizing Non-Codominated Actions Let $F^* = \left\{f \in F : f_i(r^{t+1}) \in A_{i,t} \setminus B_i(r_i^{t+1})\ \forall i, t, r^{t+1}\right\}$, and let $F^{*\geq t} = \left\{f^{\geq t} \in F^{\geq t} : f_i(r^{\tau+1}) \in A_{i,\tau} \setminus B_i(r_i^{\tau+1})\ \forall i, \tau \geq t, r^{\tau+1}\right\}$. As in Myerson's Lemma 3, the contrapositive of the definition of codomination requires that there exist distributions $(\pi_t)_{t=1}^T$ with $\pi_t \in \Delta\left(F^{*\geq t} \times X^t \times S_t\right)$ for each t such that the following conditions hold:[27]

1. Each non-codominated action is recommended: for each i, t, and $(x_i^t, s_{i,t})$, letting $\text{supp}\left(x_i^t, s_{i,t}\right) = \{m_{i,t} \in A_{i,t} : \exists\left(f^{\geq t}, x_{-i}^t, s_{-i,t}\right) \text{ with } \pi_t\left(f^{\geq t}, x^t, s_t\right) > 0 \text{ and } f_i^{\geq t}(x^t, s_t) = m_{i,t}\}$ denote the support of $f_i^{\geq t}$ at $(x_i^t, s_{i,t})$ under π_t, we have
$$\text{supp}(x_i^t, s_{i,t}) = A_{i,t} \setminus B_i(\mathring{h}_i^{R,t}).$$

2. Honesty and obedience is optimal under π_t on path: for each i and t, letting $\Pr^{\sigma^*,\pi_t}(f^{\geq t}, \mathring{h}^{T+1}, m^{t:T}) = \pi_t\left(f^{\geq t}, \mathring{h}^{R,t}\right)\Pr^{\sigma^*,f^{\geq t}}(\mathring{h}^{T+1}, m^{t:T}|\mathring{h}^{R,t})$ with $m^{t:T} = (m_t, \ldots, m_T)$, for each $\tau \geq t$, each $(\mathring{h}_i^{R,\tau}, m_i^{t:\tau})$ with $\Pr^{\sigma^*,\pi_t}(\mathring{h}_i^{R,\tau}, m_i^{t:\tau}) > 0$, and each continuation strategy σ_i' from period τ, we have
$$\begin{aligned}&\sum_{f^{\geq t}\in F^{*\geq t},\mathring{h}^{R,\tau}} \beta^{\sigma^*,\pi_t}\left(f^{\geq t}, \mathring{h}^{R,\tau}|\mathring{h}_i^{R,\tau}, m_i^{t:\tau}\right)\bar{u}_i\left(\sigma^*|f^{\geq t}, \mathring{h}^{R,\tau}\right)\\ \geq &\sum_{f^{\geq t}\in F^{*\geq t},\mathring{h}^{R,\tau}} \beta^{\sigma^*,\pi_t}\left(f^{\geq t}, \mathring{h}^{R,\tau}|\mathring{h}_i^{R,\tau}, m_i^{t:\tau}\right)\bar{u}_i\left(\sigma_i', \sigma_{-i}^*|f^{\geq t}, \mathring{h}^{R,\tau}\right),\end{aligned} \tag{20}$$

[26] Here we view f as a mediator state drawn at the start of the game, and include it in K.

[27] Myerson (1986) relies on the existence of a sequence $\pi_t^1, ..., \pi_t^L$ for each t such that $\bigcup_{l=1}^L \text{supp}^l(x_i^t, s_{i,t}) = A_{i,t} \setminus B_i(\mathring{h}_i^{R,t})$ and (20) holds for each π_t^l. Given such a sequence, the distribution $\pi_t := \frac{1}{L}\sum_{l=1}^L \pi_t^l$ satisfies both of our conditions, so we work with a single distribution π_t for each t.

where $\beta^{\sigma^*,\pi_t}\left(f^{\geq t},\mathring{h}^{R,\tau}|\mathring{h}_i^{R,\tau},m_i^{t:\tau}\right)=\Pr^{\sigma^*,\pi_t}(f^{\geq t},\mathring{h}^{R,\tau},m_i^{t:\tau})/\Pr^{\sigma^*,\pi_t}(\mathring{h}_i^{R,\tau},m_i^{t:\tau})$.

"Motivating Equilibrium" Construction We define a sequence of quasi-strategy profiles $\left(\sigma^k,\phi^k\right)_k$, where the quasi-SE strategies for the desired "motivating equilibrium" are $(\sigma,\phi):=\lim_k\left(\sigma^k,\phi^k\right)$.

Players' strategies σ^k: Each player i is faithful: she plays $a_{i,t}=m_{i,t}$ after each $m_{i,t}\in A_{i,t}$ and always reports $r_{i,t}=(a_{i,t-1},s_{i,t})$. After receiving $m_{i,t}=\star$, with probability $1-\sqrt{\varepsilon_k}$ player i takes $a_{i,t}$ according to $\hat{\sigma}_i(x_i^t,s_{i,t})$, and with probability $\sqrt{\varepsilon_k}$ she plays all actions with equal probability.

Mediator's strategy ϕ^k: At the beginning of the game, the mediator draws the following three variables: First, for each player i and each period t, independently across i and t, he draws $\theta_{i,t}\in\{0,1\}$ with $\Pr\left(\theta_{i,t}=0\right)=1-\sqrt{\varepsilon_k}$. Second, again independently for each i and t, he draws $\zeta_{i,t}\in\{0,1\}$ with $\Pr\left(\zeta_{i,t}=0\right)=1-\left(\frac{1}{k}\right)^{T+1}$. Third, independently for each i, he draws $\hat{f}_i$ from $\hat{\mu}_i$. Given a vector $\zeta=\zeta^{T+1}$, let $|\zeta|=\sum_{i,t}\zeta_{i,t}$ be the l_1-norm of ζ.

In each period t, the mediator has a state

$$\omega_t\in\bigcup\nolimits_{0\leq t^*\leq T,f\in F^{*,\geq t^*}}\left(t^*,f\right),$$

with initial state $\omega_0=\left(0,\hat{f}\right)$. Let $\omega=\omega^{T+1}$. Given $\theta=\theta^{T+1}$ and ζ, for each period t, the mediator recursively calculates the state ω_t and recommends $m_{i,t}\in A_{i,t}\cup\{\star\}$ as follows:

- **Notation:** Denote the number of tuples $\left(f_i^{<t},\theta_i^t,m_i^t\right)$ such that the mediator sends $m_{i,\tau}$ according to $f_i^{<t}\in\operatorname{supp}\hat{\mu}_i^{<t}$ if $\theta_{i,\tau}=0$ and sends $m_{i,\tau}=\star$ if $\theta_{i,\tau}=1$ by

$$\#M_i(x_i^t)=\left|\left\{\begin{array}{c}\left(f_i^{<t},\theta_i^t,m_i^t\right)\in\operatorname{supp}\hat{\mu}_i^{<t}\times\{0,1\}^{t-1}\times\prod_{\tau=1}^{t-1}\left(A_{i,\tau}\cup\{\star\}\right)\\ :m_{i,\tau}=f_i^{<t}(x_i^\tau,s_{i,\tau})\ \forall\tau\leq t-1\text{ s.t. }\theta_{i,\tau}=0,\text{ and }m_{i,\tau}=\star\ \forall\tau\leq t-1\text{ s.t. }\theta_{i,\tau}=1\end{array}\right\}\right| \tag{21}$$

  Let $\#M(x^t)=\prod_{i=0}^N\#M_i(x_i^t)$. In addition, for $f^{<t}\in\operatorname{supp}\hat{\mu}^{<t}$, let $\#(f^{<t})=|\{\tilde{f}\in\operatorname{supp}\hat{\mu}:\tilde{f}^{<t}=f^{<t}\}|$ be the number of recommendation strategies $\tilde{f}\in\operatorname{supp}\hat{\mu}$ which coincide with $f^{<t}$ for the first $t-1$ period.

- **Calculation of ω_t:** We define the distribution of ω_t given ω_{t-1}, r^{t+1}, θ, and ζ. If $\omega_{t-1}\neq\omega_0$ then $\omega_t=\omega_{t-1}$ for sure. If $\omega_{t-1}=\omega_0$ then the mediator calculates the probability of $(\omega_{t-1}=\omega_0,\theta,\zeta,r^{t+1},m^t)$ given σ^k and the construction of ϕ^k for up to period t. Denote this probability by $p_k\left(\omega_0,\theta,\zeta,r^{t+1},m^t\right)$.

  If $p_k(\omega_0,\theta,\zeta,r^{t+1},m^t)=0$ (that is, the mediator "knows" some player was unfaithful) then $\omega_t=\omega_{t-1}$ for sure. If $p_k(\omega_0,\theta,\zeta,r^{t+1},m^t)>0$ then, for each $f^{\geq t}\in F^{*\geq t}$, the mediator draws $\omega_t=\left(t,f^{\geq t}\right)$ with probability

$$q_k\left(\omega_t|\omega_0,\theta,\zeta,r^{t+1},m^t\right)=\left(\frac{1}{k}\right)^{t+(T+1)|\zeta|}\times\frac{1}{p_k(\omega_0,\theta,\zeta,r^{t+1},m^t)}\times\frac{\pi_t(f^{\geq t},r^{t+1})}{\#(\hat{f}^{<t})\#M(r^t)},$$

  and $\omega_t=\omega_0$ with the remaining probability. (With a slight abuse of notation, we write $\#M(r^t)$ since r^t includes the report about x^t.)

As we will see, this tremble probability for ω_t ensures that, at any history, conditional on the mediator trembling to $\omega_t = (t, f^{\geq t})$, the mediation plan and the payoff-relevant history are distributed according to $\pi_t(f^{\geq t}, x^t, s_t)$ given honesty $(r^{t+1} = (x^t, s_t))$, since p_k, $\#(f^{<t})$, and $\#M(r^t)$ cancel out the probability of reaching history $(f^{<t}, x^t)$. Moreover, trembles for larger t are less likely.

- **Calculation of m_t:** If $\omega_t = \left(0, \hat{f}\right)$, then the mediator recommends $m_{i,t} = \hat{f}_i(r_i^t)$ if $\zeta_{i,t} = \theta_{i,t} = 0$, recommends $m_{i,t} = \star$ if $\zeta_{i,t} = 0$ and $\theta_{i,t} = 1$, and recommends all non-codominated actions $A_{i,t} \setminus B_i(r_i^{t+1})$ with equal probability if $\zeta_{i,t} = 1$. If $\omega_t = (t, f^{\geq t})$ with $t \geq 1$, then the mediator recommends $m_{i,t} = f_i^{\geq t}(r^{t+1})$.

(Intuitively, $\omega_t = \left(0, \hat{f}\right)$ means the mediator intends to implement the mediated, full-support NE; $\omega_t = (t, f^{\geq t})$ with $t \geq 1$ means the mediator switched in period t to implementing mediation plan $f^{\geq t}$; $\theta_{i,t} = 1$ indicates a tremble signalling player i to play the unmediated NE in period t; and $\zeta_{i,t} = 1$ indicates a tremble recommending all non-codominated actions with positive probability while still intending to implement the mediated NE.)

Joint Distribution of Histories and Mediator States Let $\delta_k(h^{T+1}, \theta, \zeta, \omega)$ denote the probability of $(h^{T+1}, \theta, \zeta, \omega)$ under quasi-strategy profile (σ^k, ϕ^k). Note that, for each $(\omega_0, \theta, \zeta, r^{t+1}, m^t)$ such that $p_k(\omega_0, \theta, \zeta, r^{t+1}, m^t) > 0$, we have

$$p_k(\omega_0, \theta, \zeta, r^{t+1}, m^t) \geq \hat{\mu}\left(\hat{f}\right) \times \frac{(\varepsilon_k)^{NT}}{|A^T|} \times \left(\frac{1}{k}\right)^{(T+1)|\zeta|}.$$

Hence, for each $(\omega_0, \theta, \zeta, r^{t+1}, m^t)$ and $\omega_t \neq \omega_0$, we have

$$q_k\left(\omega_t | \omega_0, \theta, \zeta, r^{t+1}, m^t\right) \leq \left(\frac{1}{k}\right)^t \times \frac{|A^T|}{(\varepsilon_k)^{NT}} \times \frac{1}{\hat{\mu}\left(\hat{f}\right)} \times \frac{\pi_t(f^{\geq t}, r^{t+1})}{\#(\hat{f}^{<t})\#M(r^t)}.$$

Since $k\,(\varepsilon_k)^{NT} / |A^T| \to \infty$ as $k \to \infty$, this implies

$$\lim_{k \to \infty} q_k\left(\omega_t | \omega_0, \theta, \zeta, r^{t+1}, m^t\right) = 0. \tag{22}$$

Given x^t, $f^{<t} \in \operatorname{supp} \hat{\mu}^{<t}$, θ^t, and ζ^t, let $M^t(x^t, f^{<t}, \theta^t, \zeta^t)$ denote the set of m^t such that, for each i and $\tau = 1, ..., t-1$, (i) $m_{i,\tau} = f_i^{<t}(x_i^\tau, s_{i,\tau})$ if $\zeta_{i,\tau} = \theta_{i,\tau} = 0$, (ii) $m_{i,\tau} = \star$ if $\zeta_{i,\tau} = 0$ and $\theta_{i,\tau} = 1$, and (iii) $m_{i,\tau} \in A_{i,\tau} \setminus B_i(x_i^\tau, s_{i,\tau})$ if $\zeta_{i,\tau} = 1$.

If $\omega_0 = \cdots = \omega_{t-1} = \left(0, \hat{f}\right)$ and $\omega_t = \cdots = \omega_T = \left(t, \tilde{f}^{\geq t}\right)$, we define $t^*(\omega) = t$ and

$f(\omega) = \left(\hat{f}^{<t}, \tilde{f}^{\geq t}\right)$. We have[28]

$$\begin{aligned}
& \delta_k(h^{T+1}, \theta, \zeta, \omega) \\
= \; & 1_{\{t^*(\omega)=T+1\}} \times \Pr^{\phi^k}(\omega_t = \omega_0 \forall t) \times \Pr^{\phi^k}(\theta, \zeta) \times \hat{\mu}(f(\omega)) \times \Pr^{\sigma^k}\left(h^{T+1} | f(\omega), \theta, \zeta\right) \\
& + \sum_{t=1}^{T} 1_{\{t^*(\omega)=t\}} \sum_{\tilde{f}: \tilde{f}^{<t} = f^{<t}(\omega)} \left(\begin{array}{c} p_k(\omega_0, \theta, \zeta, \mathring{h}^{R,t}, m^t) \left(\frac{1}{k}\right)^{t+(T+1)|\zeta|} \\ \times \frac{1}{p_k(\omega_0, \theta^t, \zeta^t, \mathring{h}^{R,t}, m^t)} \frac{\pi_t(f^{\geq t}(\omega), \mathring{h}^{R,t})}{\#(f^{<t}(\omega)) \# M(\mathring{h}^t)} \\ \times \Pr^{\sigma^*}\left(\mathring{h}^{T+1} | \mathring{h}^{R,t}, f^{\geq t}(\omega)\right) \\ \times 1_{\left\{m^t \in M^t(\mathring{h}^t, f^{<t}(\omega), \theta^t, \zeta^t) \text{ and } m_\tau = f^{\geq t}(\omega)(\mathring{h}^{R,\tau}) \; \forall \tau \geq t\right\}} \end{array} \right).
\end{aligned}$$

Canceling out p_k and $\frac{1}{p_k}$, and $\sum_{\tilde{f}: \tilde{f}^{<t} = f^{<t}(\omega)}$ and $\frac{1}{\#(f^{<t}(\omega))}$, we have

$$\begin{aligned}
& \delta_k(h^{T+1}, \theta, \zeta, \omega) \\
= \; & 1_{\{t^*(\omega)=T+1\}} \times \Pr^{\phi^k}(\omega_t = \omega_0 \forall t) \times \Pr^{\phi^k}(\theta, \zeta) \times \hat{\mu}(f(\omega)) \times \Pr^{\sigma^k}\left(h^{T+1} | f(\omega), \theta, \zeta\right) \\
& + \sum_{t=1}^{T} 1_{\{t^*(\omega)=t\}} \left(\begin{array}{c} \left(\frac{1}{k}\right)^{t+(T+1)|\zeta|} \frac{\pi_t(f^{\geq t}(\omega), \mathring{h}^{R,t})}{\# M(\mathring{h}^t)} \Pr^{\sigma^*}\left(\mathring{h}^{T+1} | \mathring{h}^{R,t}, f^{\geq t}(\omega)\right) \\ \times 1_{\left\{m^t \in M^t(\mathring{h}^t, f^{<t}(\omega), \theta^t, \zeta^t) \text{ and } m_\tau = f^{\geq t}(\omega)(\mathring{h}^{R,\tau}) \; \forall \tau \geq t\right\}} \end{array} \right).
\end{aligned}$$

Quasi-Sequential Equilibrium and its Validity Let K be the set of mediator histories $(r, \theta, \zeta, \omega)$ consistent with ϕ^k. Let J_i be the set of player i's histories h_i^{T+1} such that (i) $m_{i,t} \in A_{i,t} \cup \{\star\} \setminus B_i(r_i^t) \; \forall t$, (ii) $r_{i,t} = (a_{i,t-1}, s_{i,t}) \; \forall t$, and (iii) $a_{i,t} = m_{i,t} \; \forall t$ with $m_{i,t} \in A_{i,t} \setminus B_i(r_i^{t+1})$.

We claim that $(\sigma, \phi, J, K, \beta)$ is valid (for any consistent beliefs β). Clearly, no faithful strategies lead the mediator's history out of K. Moreover, no unfaithful strategy of players $-i$ leads player i's history out of J_i, since J_i is determined by the mediation range and player i's own behavior. Finally, we say history h_i^{T+1} is *compatible* with a pure strategy σ_i if, for all t, we have $m_{i,t} \in A_{i,t} \cup \{\star\} \setminus B_i(r_i^{t+1})$, $r_{i,t} = \sigma_i^R(h_i^{R,t})$, and $a_{i,t} = \sigma_i^A(h_i^{A,t})$. To establish validity, it suffices to show that, for any faithful pure strategy σ_i, each h_i^{T+1} that is compatible with σ_i occurs with positive probability at profile $\left(\sigma_i, \sigma_{-i}^k, \phi^k\right)$. Take (i) $f = \hat{f} \in \operatorname{supp} \hat{\mu}$ such that $m_{i,t} = \hat{f}_i(r_i^{t+1})$ for each t with $m_{i,t} \in A_i \setminus \hat{B}_i(r_i^{t+1})$, (ii) $\omega_t = (0, \hat{f})$ for each t, (iii) $\zeta_{i,t} = 0$ for each t with $m_{i,t} \in A_{i,t} \cup \{\star\} \setminus \hat{B}_i(r_i^{t+1})$ (and $\zeta_{i,t} = 1$ otherwise), (iv) $\theta_{i,t} = 1$ for each t with $m_{i,t} = \star$ (and $\theta_{i,t} = 0$ otherwise), (v) $\zeta_{j,t} = 0$ and $\theta_{j,t} = 1$ for each t and $j \neq i$. By definition of δ_k, (i)–(v) occurs with positive probability. Given (i)–(iv), m_i^{T+1} occurs with positive probability; and given (v), each player j takes each action with probability at least $\sqrt{\varepsilon_k}$. Hence, h_i^{T+1} occurs with a positive probability.

By Lemma 7, it suffices to show that, for $\beta(Z|Z') = \lim_{k \to \infty} \frac{\delta_k(Z \cap Z')}{\delta_k(Z')}$,

[28] Here $\mathring{h}^t$ is the projection of h^{T+1} on X^t.

1. *[Sequential rationality of reports]* For all $i \neq 0, t, \sigma_i' \in \Sigma_i$ and all $h_i^{R,t} \in J_i^{R,t}$,

$$\sum_{\omega_t,\theta,\zeta,h^{R,t}\in(J,K)} \beta\left(\omega_t,\theta,\zeta,h^{R,t}|h_i^{R,t}\right)\bar{u}_i\left(\sigma,\phi|\omega_t,\theta,\zeta,h^{R,t}\right)$$
$$\geq \sum_{\omega_t,\theta,\zeta,h^{R,t}\in(J,K)} \beta\left(\omega_t,\theta,\zeta,h^{R,t}|h_i^{R,t}\right)\bar{u}_i\left(\sigma_i',\sigma_{-i},\phi|\omega_t,\theta,\zeta,h^{R,t}\right). \qquad (23)$$

2. *[Sequential rationality of actions]* For all $i \neq 0, t, \sigma_i' \in \Sigma_i$ and all $h_i^{A,t} \in J_i^{A,t}$,

$$\sum_{\omega_t,\theta,\zeta,h^{R,t}\in(J,K)} \beta\left(\omega_t,\theta,\zeta,h^{A,t}|h_i^{A,t}\right)\bar{u}_i\left(\sigma,\phi|\omega_t,\theta,\zeta,h^{A,t}\right)$$
$$\geq \sum_{\omega_t,\theta,\zeta,h^{R,t}\in(J,K)} \beta\left(\omega_t,\theta,\zeta,h^{A,t}|h_i^{A,t}\right)\bar{u}_i\left(\sigma_i',\sigma_{-i},\phi|\omega_t,\theta,\zeta,h^{A,t}\right). \qquad (24)$$

Mediator Trembles that Explain a Faithful History Given a faithful history $h_i^{R,t}$ for some i,t, we say $(0,\zeta)$ *explains* $h_i^{R,t}$ if there exist $\hat{f} \in \operatorname{supp}\hat{\mu}$, θ, and $h_{-i}^{R,t}$ such that, for each j and $\tau = 1,...,t-1$, (i) $m_{j,\tau} = a_{j,\tau} = \hat{f}_j(\mathring{h}_j^{R,\tau})$ if $\zeta_{j,\tau} = \theta_{j,\tau} = 0$, (ii) $m_{j,\tau} = \star$ if $\zeta_{j,\tau} = 0$ and $\theta_{j,\tau} = 1$, and (iii) $m_{j,\tau} = a_{j,\tau} \in A_{j,\tau} \setminus B_j(\mathring{h}_j^{R,\tau})$ if $\zeta_{j,\tau} = 1$, and also for each $\tau = 0,\ldots,t-1$, (iv) $p\left(s_{\tau+1}|s^\tau, a^\tau\right) > 0$. We say *(i)–(iii) hold for* $t = t_0$ if the first three conditions hold with $t = t_0$.

Given faithful $h_i^{R,t}$, we say (t^*,ζ) *with* $t^* \geq 1$ *explains* $h_i^{R,t}$ if there exist $\hat{f}^{<t^*} \in \operatorname{supp}\hat{\mu}^{<t^*}$, $f^{\geq t^*}$, θ, and $h_{-i}^{R,t}$ such that (i)–(iii) hold for $t = t^*$, (iv) $p\left(s_{\tau+1}|s^\tau, a^\tau\right) > 0$ for each $\tau = 0,...,t-1$, (a) $\pi_{t^*}\left(f^{\geq t^*}, \mathring{h}^{R,t^*}\right) > 0$, (b) $m_\tau = a_\tau = f^{\geq t^*}(\mathring{h}^{R,\tau})$ for each $\tau = t^*,...,t-1$, and (c) $\Pr^{\sigma^k}\left(\mathring{h}^{R,t}|\mathring{h}^{R,t^*}, f^{\geq t^*}\right) > 0$.

Similarly, given a faithful history $h_i^{A,t}$, we say $(0,\zeta)$ *explains* $h_i^{A,t}$ if (i)–(ii) hold for $\tau = 1,...,t-1$, (iii) holds for $\tau = 1,...,t$, and (iv) holds for $\tau = 0,...,t-1$; and (t^*,ζ) *explains* $h_i^{A,t}$ if the above conditions are satisfied with "for each $\tau = t^*,...,t-1$" in condition (b) replaced with "for each $\tau = t^*,...,t$."

Let

$$\Xi = \bigcup_{0\leq t^*\leq T}\bigcup_{\zeta\in\{0,1\}^{NT}}(t^*,\zeta).$$

Order the elements of Ξ such that $(t^*,\zeta) < (\tilde{t}^*,\tilde{\zeta})$ if (i) $|\zeta| < \left|\tilde{\zeta}\right|$ or (ii) $|\zeta| = \left|\tilde{\zeta}\right|$ and $t^* < \tilde{t}^*$. That is, $(t^*,\zeta) < (\tilde{t}^*,\tilde{\zeta})$ if a tremble to π_{t^*} with $\zeta_{j,\tau} = 1$ for $|\zeta|$ values of j,τ is more likely than a tremble to $\pi_{\tilde{t}^*}$ with $\zeta_{j,\tau} = 1$ for $\left|\tilde{\zeta}\right|$ values of j,τ.

Given the specified order on Ξ, let $\xi(h_i^{R,t})$ and $\xi(h_i^{A,t})$ be the smallest triples (t^*,ζ) that explain $h_i^{R,t}$ and $h_i^{A,t}$, respectively. Since Ξ is a finite set and the distribution over player i's compatible histories has full support, these are well-defined. As the order reflects the likelihood of trembles, the following lemma holds:

For any $\zeta_i^t \in \{0,1\}^{t-1}$, let $\left(0,\zeta_i^t\right) := \left(0,\tilde{\zeta}^t\right)$ with $\tilde{\zeta}_i^t = \zeta_i^t$, $\tilde{\zeta}_{i,\tau} = 0$ $\forall\tau \geq t$, and $\tilde{\zeta}_{j,\tau} = 0$ $\forall j \neq i,\tau$. Define $\left(t^*,\zeta_i^{t^*}\right)$ similarly.

Lemma 10 *For each faithful $h_i^{R,t}$ and $h_i^{A,t}$, the following three claims hold:*

1. *$\xi(h_i^{R,t})$ and $\xi(h_i^{A,t})$ satisfy $\zeta_{j,t} = 0$ for all $j \neq i$ and t.*

2. *Either $\xi(h_i^{R,t}) = \left(0, \zeta_i^t\right)$ for some ζ_i^t, or $\xi(h_i^{R,t}) = \left(t^*, \zeta_i^{t^*}\right)$ for some $t^* \leq t$ and $\zeta_i^{t^*}$. Likewise, either $\xi(h_i^{A,t}) = \left(0, \zeta_i^t\right)$ for some ζ_i^t, or $\xi(h_i^{A,t}) = \left(t^*, \zeta_i^{t^*}\right)$ for some $t^* \leq t$ and $\zeta_i^{t^*}$.*

3. *We have*

$$\lim_{k\to\infty} \delta_k \left(\xi = \xi(h_i^{R,t}) | h_i^{R,t}\right) = 1, \tag{25}$$

$$\lim_{k\to\infty} \delta_k \left(\xi = \xi(h_i^{A,t}) | h_i^{A,t}\right) = 1. \tag{26}$$

Proof. Claim 1: Whenever (t, ζ) with $\zeta_{j,\tau} = 1$ for some $j \neq i$ and τ explains $h_i^{R,t}$ or $h_i^{A,t}$, so does $\left(t, \tilde{\zeta}\right)$ where $\tilde{\zeta}_{j,\tau} = 0$ and $\tilde{\zeta}_{j',\tau'} = \zeta_{j',\tau'}$ for all $(j', \tau') \neq (j, \tau)$, as given $a_{j,\tau}$ and [$t = 0$ or $\tau < t^*$],[29] we can take $\zeta_{j,\tau} = 0$, $\theta_{j,\tau} = 1$, and $m_{j,\tau} = \star$, rather than $\zeta_{j,\tau} = 1$ and $m_{j,\tau} = a_{j,\tau}$. The claim follows as $\left(t, \tilde{\zeta}\right) < (t, \zeta)$.

Claim 2: By definition, whenever $\left(0, \zeta_i^{t'}\right)$ with $t' > t$ explains $h_i^{R,t}$, so does $\left(0, \zeta_i^t\right)$. In addition, whenever $(t', \zeta_i^{t''})$ with $t' > t$ explains $h_i^{R,t}$, so does $(0, \tilde{\zeta}_i^t)$ satisfying $\tilde{\zeta}_{i,\tau} = \zeta_{i,\tau}$ for $\tau \leq \min\{t'', t\}$ and $\tilde{\zeta}_{i,\tau} = 0$ for $\tau > \min\{t'', t\}$. Finally, whenever $\left(t', \zeta_i^{t''}\right)$ with $t'' > t'$ explains $h_i^{R,t}$, so does $\left(t', \zeta_i^{t'}\right)$. Hence, the conclusion for $\xi(h_i^{R,t})$ holds.

The proof for $\xi(h_i^{A,t})$ is the same, except that we also show $\xi(h_i^{A,t}) \neq \left(0, \zeta_i^{t+1}\right)$ with $\zeta_{i,t} = 1$. To see why this new condition holds, whenever $\left(0, \zeta_i^{t+1}\right)$ with $\zeta_{i,t} = 1$ explains $h_i^{A,t}$, so does some $\left(t^*, \zeta_i^{t^*}\right)$ with $t^* = t$. This is because, given $t^* = t$, for each r_i^{t+1}, each $m_{i,t} \in A_{i,t} \setminus B_i(r_i^{t+1})$, and each $\hat{f}^{<t} \in \operatorname{supp} \hat{\mu}^{<t}$, we have $m_{i,t} \in \operatorname{supp}_i(r_i^{t+1})$. Given the order on ξ, we have $\xi(h_i^{A,t}) = \left(t^*, \zeta_i^{t^*}\right)$.

Claim 3: We prove (25); the proof of (26) is analogous. Suppose (t^*, ζ^*) explains $h_i^{R,t}$. Given Claim 1, we can take $\zeta_{j,t}^* = 0$ for each $j \neq i$ and t. Since (i) any action profile $a_{j,t}$ is taken with probability at least $\sqrt{\varepsilon_k} / |A_{j,t}|$ after each $\hat{f} \in \operatorname{supp} \hat{\mu}$ given $\zeta_{j,t} = 0$ and $\theta_{j,t} = 1$, and (ii) $\theta_{j,t} = 1$ occurs with probability $\sqrt{\varepsilon_k}$, we have

$$\delta_k \left(t^*, \zeta, h_i^{R,t}\right) \geq \left(\frac{1}{k}\right)^{(T+1)|\zeta^*|} \left(1 - \left(\frac{1}{k}\right)^Z\right)^{NT} \frac{(\varepsilon_k)^{NT}}{|A^T|} \left(\frac{1}{k}\right)^{t^*} \underline{\varepsilon},$$

where we take $\underline{\varepsilon} > 0$ such that

$$\pi_t \left(f^{\geq t}, x^t, s_t\right) \Pr^{\sigma^*} \left(x | x^t, s_t, f^{\geq t}\right) \geq \underline{\varepsilon} \;\forall \left(t, x, f^{\geq t}\right) \text{ s.t. } \pi_t \left(f^{\geq t}, x^t, s_t\right) \Pr^{\sigma^*} \left(x | f^{\geq t}, x^t, s_t\right) > 0.$$

[29] Note that, given $\tau \geq t^*$, $m_{j,t}$ is independent of $\zeta_{j,t}$.

For each $(t,\zeta) \neq (t^*, \zeta^*)$, if it does not explain $h_i^{R,t}$, then $\delta_k\left(t,\zeta,h_i^{R,t}\right) = 0$. If it does, then

$$\delta_k\left(t,\zeta,h_i^{R,t}\right) \leq \left(\frac{1}{k}\right)^{(T+1)|\zeta|}\left(\frac{1}{k}\right)^t.$$

Since either $|\zeta^*| < |\zeta|$ or $[|\zeta^*| = |\zeta|$ and $t^* < t]$, we have

$$\frac{\delta_k\left(t,\zeta,h_i^{R,t}\right)}{\delta_k\left(t^*,\zeta,h_i^{R,t}\right)} \leq \frac{\frac{1}{k}}{\left(1-\left(\frac{1}{k}\right)^{T+1}\right)^{NT}\frac{(\varepsilon_k)^{NT}}{|A^T|}\underline{\varepsilon}}.$$

Since $k\left(\varepsilon_k\right)^{NT} / \left|A^T\right| \to \infty$, this ratio converges to 0, which implies (25). ■

Incentive Compatibility We now establish (23) and (24). By Lemma 10, there are two cases:

Case 1: $\xi(h_i^{R,t}) = \left(0,\zeta_i^t\right)$ or $\xi(h_i^{A,t}) = \left(0,\zeta_i^t\right)$.

Let $\Omega_0 = \bigcup_{\hat{f}\in \operatorname{supp}\hat{\mu}}\left(0,\hat{f}\right)$. Since $\Pr^{\phi^k}\left(\omega_t = \omega_0 \forall t\right) \to 1$, we see that $\xi(h_i^{R,t}) = \left(0,\zeta_i^t\right)$ implies $\delta_k(\omega_T \in \Omega_0 | h_i^{R,t}) = 1$, and $\xi(h_i^{A,t}) = \left(0,\zeta_i^t\right)$ implies $\delta_k(\omega_T \in \Omega_0 | h_i^{A,t}) = 1$.

For each i, t, and $(x_i^t, s_{i,t})$, arbitrarily fix some action $m_{i,t}^*(x_i^t, s_{i,t}) \in A_{i,t} \setminus \hat{B}_i(x_i^t, s_{i,t})$. With a slight abuse of notation, we write

$$m_{i,t}^*(h_i^{R,t+1}) = \begin{cases} m_{i,t} & \text{if } m_{i,t} \in A_{i,t} \setminus \hat{B}_i(\mathring{h}_i^{R,t+1}) \\ m_{i,t}^*(\mathring{h}_i^{R,t+1}) & \text{if } m_{i,t} \notin A_{i,t} \setminus \hat{B}_i(\mathring{h}_i^{R,t+1}) \end{cases},$$

where $m_{i,t}$ is the corresponding element of $h_i^{R,t+1}$. For each i and each faithful $h_i^{R,t}$ with $\delta_k\left(h_i^{R,t}\right) > 0$, let $\lambda(h_i^{R,t})$ be the history where each message $m_{i,\tau}$ is replaced by $m_{i,\tau}^*(\mathring{h}_i^{R,\tau+1}) \in A_{i,\tau} \setminus \hat{B}_i(\mathring{h}_i^{R,\tau+1})$ for each $\tau \leq t-1$. That is, we replace each action recommendation outside the support of $\hat{\mu}$ with some fixed recommendation within the support. Note that $\Pr^{\hat{\sigma}^k,\hat{\mu}}\left(\lambda(h_i^{R,t})\right) > 0$ whenever $\delta_k\left(h_i^{R,t}\right) > 0$. Define $\lambda(h_i^{A,t})$ analogously.

Given $\xi = \left(0,\zeta_i^{t-1}\right)$, all trembles are independent across players. Each player i can then safely ignore the possibility that $\zeta_{j,t} = 1$ for any j, since $\left(\frac{1}{k}\right)^{T+1}$ (the probability that $\zeta_{j,t} = 1$) is much less than ε_k (the probability that $m_{j,t} = \star$ and player j trembled). This suggests that each player's beliefs in the constructed quasi-SE coincide with those in the original full-support NE, as confirmed by following lemma:

Lemma 11 *The following two claims hold:*

1. *For each $h_i^{R,t}$ and ζ_i^t satisfying $\xi(h_i^{R,t}) = \left(0,\zeta_i^t\right)$ and each $\mathring{h}_{-i}^{R,t} \in X_{-i}^t \times S_{-i,t}$, we have*

$$\lim_k \delta_k\left(\mathring{h}_{-i}^{R,t} | h_i^{R,t}\right) = \lim_k \Pr^{\hat{\sigma}^k,\hat{\mu}}\left(\mathring{h}_{-i}^{R,t} | \lambda(h_i^{R,t})\right). \tag{27}$$

2. *For each $h_i^{A,t}$ and ζ_i^t satisfying $\xi(h_i^{A,t}) = \left(0, \zeta_i^t\right)$ and each $\mathring{h}_{-i}^{A,t} \in X_{-i}^t \times S_{-i,t}$, we have*

$$\lim_k \delta_k \left(\mathring{h}_{-i}^{A,t}, m_{-i,t} | h_i^{A,t}\right) = \lim_k \Pr^{\hat{\sigma}^k, \hat{\mu}} \left(\mathring{h}_{-i}^{A,t}, m_{-i,t} | \lambda(h_i^{A,t})\right). \tag{28}$$

Lemma 11 follows from applying Bayes' rule inductively on t. We relegate the proof to the end of the appendix.

We now simplify (23). $\xi(h_i^{R,t}) = \left(0, \zeta_i^t\right)$ implies $\delta_k(\omega_T \in \Omega_0 | h_i^{R,t}) = 1$. Hence, we can replace ω_t with $\hat{f}$ drawn from $\hat{\mu}$. Since θ and ζ are independent across periods, $\beta\left(\theta_\tau = \zeta_\tau = 0 \ \forall \tau \geq t | h_i^{R,t}\right) = 1$. Hence, (23) simplifies to

$$\sum_{\hat{f} \in \operatorname{supp} \hat{\mu}, \mathring{h}^{R,t}} \beta\left(\hat{f}, \mathring{h}^{R,t} | h_i^{R,t}\right) \bar{u}_i\left(\sigma^* | \hat{f}, \mathring{h}^{R,t}\right) \geq \sum_{\hat{f} \in \operatorname{supp} \hat{\mu}, \mathring{h}^{R,t}} \beta\left(\hat{f}, \mathring{h}^{R,t} | h_i^{R,t}\right) \bar{u}_i\left(\sigma_i', \sigma_{-i}^* | \hat{f}, \mathring{h}^{R,t}\right).$$

By Lemma 11, this is equivalent to

$$\sum_{\hat{f} \in \operatorname{supp} \hat{\mu}, \mathring{h}^{R,t}} \beta^{\sigma^*, \hat{\mu}}\left(\hat{f}, \mathring{h}^{R,t} | \lambda(h_i^{R,t})\right) \bar{u}_i\left(\sigma^* | \hat{f}, \mathring{h}^{R,t}\right) \geq \sum_{\hat{f} \in \operatorname{supp} \hat{\mu}, \mathring{h}^{R,t}} \beta^{\sigma^*, \hat{\mu}}\left(\hat{f}, \mathring{h}^{R,t} | \lambda(h_i^{R,t})\right) \bar{u}_i\left(\sigma_i', \sigma_{-i}^* | \hat{f}, \mathring{h}^{R,t}\right),$$

which follows from (19). The proof for (24) is the same.

Case 2: $\xi(h_i^{R,t}) = \left(t^*, \zeta_i^{t^*-1}\right)$ or $\xi(h_i^{A,t}) = \left(t^*, \zeta_i^{t^*-1}\right)$.

In the next lemma, we abbreviate the event $\{\omega : t^*(\omega) = t \text{ and } f^{\geq t}(\omega) = f^{\geq t}\}$ by $f^{\geq t}$.

Lemma 12 *The following two claims hold:*

1. *For each $h_i^{R,t}$ satisfying $\xi(h_i^{R,t}) = \left(t^*, \zeta_i^{t^*}\right)$, each $f^{\geq t^*} \in F^{\geq t*}$, and each $\mathring{h}_{-i}^{R,t} \in X_{-i}^t \times S_{-i,t}$, we have*

$$\lim_k \delta_k\left(f^{\geq t^*}, \mathring{h}_{-i}^{R,t} | h_i^{R,t}, t^*, \zeta_i^{t^*}\right) = \frac{\pi_{t^*}(f^{\geq t^*}, \mathring{h}^{R,t^*}) \Pr^\sigma\left(\mathring{h}^{R,t} | f^{\geq t^*}, \mathring{h}^{R,t^*}\right)}{\sum_{\mathring{h}_{-i}'^{R,t}, f'^{\geq t}} \pi_{t^*}\left(f'^{\geq t^*} \mathring{h}_i^{R,t^*}, \mathring{h}_{-i}'^{R,t^*}\right) \Pr^\sigma\left(\mathring{h}_i^{R,t}, \mathring{h}_{-i}'^{R,t} | f'^{\geq t^*}, \mathring{h}_i^{R,t^*}, \mathring{h}_{-i}'^{R,t^*}\right)}. \tag{29}$$

2. *For each $h_i^{A,t}$ satisfying $\xi(h_i^{A,t}) = \left(t^*, \zeta_i^{t^*}\right)$, each $f^{\geq t^*} \in F^{\geq t*}$, and each $\mathring{h}_{-i}^{R,t} \in X_{-i}^t \times S_{-i,t}$, we have*

$$\lim_k \delta_k\left(f^{\geq t^*}, \mathring{h}_{-i}^{A,t} | h_i^{A,t}, t^*, \zeta_i^{t^*}\right) = \frac{\pi_{t^*}(f^{\geq t^*}, \mathring{h}^{A,t^*}) \Pr^\sigma\left(\mathring{h}^{A,t} | f^{\geq t^*}, \mathring{h}^{A,t^*}\right)}{\sum_{\mathring{h}_{-i}'^{A,t}, f'^{\geq t}} \pi_{t^*}\left(f'^{\geq t^*}, \mathring{h}_i^{A,t^*}, \mathring{h}_{-i}'^{A,t^*}\right) \Pr^\sigma\left(\mathring{h}_i^{A,t}, \mathring{h}_{-i}'^{A,t} | f'^{\geq t^*}, \mathring{h}_i^{A,t^*}, \mathring{h}_{-i}'^{A,t^*}\right)}. \tag{30}$$

Lemma 12 follows from another application of Bayes' rule. The proof is relegated to the end of the appendix.

Given $\xi(h_i^{R,t}) = \left(t^*, \zeta_i^{t^*-1}\right)$, player i believes that the mediator and players $-i$ do not tremble after period t^*, and that recommendations are independent of θ and ζ after period t^*. Hence, by Lemma 12, (23) is equivalent to (20), and therefore follows from the definition of π_{t^*}. The proof for (24) is analogous.

Final Construction Take any canonical NE in which codominated actions are never recommended. Let π^* be the distribution of the mediator's pure strategy in this target equilibrium. At the beginning of the game, the mediator draws $f \in F^*$ according to π^* with probability $1 - \frac{1}{k}$, and the mediator follows the motivating equilibrium strategy constructed above with probability $\frac{1}{k}$. Players are faithful, and after receiving $m_{i,t} = \star$, with probability $1 - \sqrt{\varepsilon_k}$ player i takes $a_{i,t}$ according to $\hat{\sigma}_i(\mathring{h}_i^{R,t})$, and with probability $\sqrt{\varepsilon_k}$ she takes all actions with equal probability.

Since $f \in F^*$ does not recommend codominated actions and each player's history has full support, this quasi-strategy is valid. Moreover, beliefs are defined as the limit of

$$\tilde{\delta}_k(z, \theta, \zeta, \omega) = \left(1 - \frac{1}{k}\right)\pi^*(f, z) + \frac{1}{k}\delta_k(z, \theta, \zeta, \omega),$$

where when the mediator follows π^* we define $\theta = \zeta = 0$ and $\omega = (0, f)$, where f is drawn from π^*, and $\pi^*(f, z)$ is the probability of (f, z) when the mediator draws f from π^*.

As $k \to \infty$, until player i observes an off-path recommendation or signal given π^*, she believes the equilibrium follows $\pi^*(f, z)$. Faithfulness is optimal since π^* is a NE. Once player i observes an off-path recommendation or signal, she believes the equilibrium follows $\delta_k(f, z)$. In this case, faithfulness is optimal by Lemma 7, (23), and (24).

Proof of Lemma 11 We prove (27); the proof of (28) is analogous. We will prove the following: for each i, t, each faithful $h_i^{R,t}$ with $\delta_k\left(h_i^{R,t}\right) > 0$, each ζ_i^t, and each $\mathring{h}_{-i}^{R,t}$, there exist numbers $\varphi_k^R(h_i^{R,t}, \zeta_i^t) \geq 0$ and $e_k^R\left(h_i^{R,t}, \zeta_i^t, \mathring{h}_{-i}^{R,t}\right) \geq 0$ such that

$$\begin{gathered}\delta_k(h_i^{R,t}, \zeta_i^t, \mathring{h}_{-i}^{R,t} | \omega_T \in \Omega_0) = \varphi_k^R(h_i^{R,t}, \zeta_i^t)\left(\Pr^{\hat{\sigma}^k, \hat{\mu}}(\lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t}) + e_k^R\left(h_i^{R,t}, \zeta_i^t, \mathring{h}_{-i}^{R,t}\right)\right), \\ \lim_{k\to\infty} \frac{e_k^R\left(h_i^{R,t}, \zeta_i^t, \mathring{h}_{-i}^{R,t}\right)}{\left(\frac{1}{k}\right)^{T+1}} \leq t.\end{gathered} \tag{31}$$

(31) is sufficient for (27), since the former implies, for each ζ_i^t,

$$\begin{aligned} & \lim_k \delta_k\left(\mathring{h}_{-i}^{R,t} | \zeta_i^t, h_i^{R,t}\right) \\ = & \lim_k \delta_k\left(\mathring{h}_{-i}^{R,t} | \zeta_i^t, h_i^{R,t}, \omega_T \in \Omega_0\right) \quad (\text{since } \xi(h_i^{R,t}) = (0, \zeta_i^t)\,) \\ = & \lim_k \frac{\varphi_k^R(h_i^{R,t}, \zeta_i^t)\left(\Pr^{\hat{\sigma}^k, \hat{\mu}}(\lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t}) + e_k^R\left(h_i^{R,t}, \zeta_i^t, \mathring{h}_{-i}^{R,t}\right)\right)}{\sum_{\mathring{h}_{-i}^{R,t\prime}} \varphi_k^R(h_i^{R,t}, \zeta_i^t)\left(\Pr^{\hat{\sigma}^k, \hat{\mu}}(\lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t\prime}) + e_k^R\left(h_i^{R,t}, \zeta_i^t, \mathring{h}_{-i}^{R,t\prime}\right)\right)} \quad (\text{by } (31)) \\ = & \lim_k \frac{\Pr^{\hat{\sigma}^k, \hat{\mu}}(\lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t}) + e_k^R\left(h_i^{R,t}, \zeta_i^t, \mathring{h}_{-i}^{R,t}\right)}{\sum_{\mathring{h}_{-i}^{R,t\prime}} \Pr^{\hat{\sigma}^k, \hat{\mu}}(\lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t\prime}) + \sum_{\mathring{h}_{-i}^{R,t\prime}} e_k^R\left(h_i^{R,t}, \zeta_i^t, \mathring{h}_{-i}^{R,t\prime}\right)} \\ = & \lim_k \frac{\Pr^{\hat{\sigma}^k, \hat{\mu}}(\lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t})}{\sum_{\mathring{h}_{-i}^{R,t\prime}} \Pr^{\hat{\sigma}^k, \hat{\mu}}(\lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t\prime})}, \end{aligned}$$

where the last equality follows as $\Pr^{\hat{\sigma}^k,\hat{\mu}}(\lambda(h_i^{R,t}),\mathring{h}_{-i}^{R,t}) \geq (\varepsilon_k)^{NT}/|A|^T$, $e_k^R\left(h_i^{R,t},\zeta_i^t,\mathring{h}_{-i}^{R,t}\right) \leq \left(\frac{1}{k}\right)^{T+1}T$, and $k(\varepsilon_k)^{NT}/|A|^T \to \infty$.

We prove (31) by induction on t. Taking $\varphi_k^R(h_i^{R,1},\zeta_i^1) = e_k^R\left(h_i^{R,1},\zeta_i^1,\mathring{h}_{-i}^{R,t}\right) = 1$, (31) holds for $t=1$. Suppose it holds for t. Since θ, ζ, and randomizations under $\left(\hat{\sigma}^k,\hat{\mu}\right)$ are independent across players, we have

$$\begin{aligned}
& \delta_k(h_i^{R,t+1},\zeta_i^t,\mathring{h}_{-i}^{R,t+1},\theta_t,\zeta_t|\omega_T\in\Omega_0) \\
= \; & \delta_k(h_i^{R,t},\zeta_i^t,\mathring{h}_{-i}^{R,t}|\omega_T\in\Omega_0) \\
& \times\left(\begin{array}{l}
\prod_{j:\theta_{j,t}=0,\zeta_{j,t}=0}\left(1-\sqrt{\varepsilon_k}\right)\left(1-\left(\frac{1}{k}\right)^{T+1}\right)\hat{\sigma}_j(\mathring{h}_j^{R,t})(a_{j,t}) \\
\times\prod_{j:\theta_{j,t}=1,\zeta_{j,t}=0}\sqrt{\varepsilon_k}\left(1-\left(\frac{1}{k}\right)^{T+1}\right)\left(\left(1-\sqrt{\varepsilon_k}\right)\hat{\sigma}_j(\mathring{h}_j^{R,t})(a_{j,t})+\frac{\sqrt{\varepsilon_k}}{|A_{j,t}|}\right) \\
\times\prod_{j:\zeta_{j,t}=1}\left(\frac{1}{k}\right)^{T+1}\frac{1}{|A_{j,t}|-\left|B_j(\mathring{h}_j^{R,t})\right|}
\end{array}\right) \\
& \times p(s_{t+1}|\mathring{h}^{R,t},a_t).
\end{aligned} \tag{32}$$

By the inductive hypothesis, the first line of (32) equals

$$\varphi_k^R(h_i^{R,t},\zeta_i^t)\left(\Pr^{\hat{\sigma}^k,\hat{\mu}}(\lambda(h_i^{R,t}),\mathring{h}_{-i}^{R,t})+e_k^R\left(h_i^{R,t},\zeta_i^t,\mathring{h}_{-i}^{R,t}\right)\right).$$

Note that

$$\Pr^{\hat{\sigma}^k,\hat{\mu}}(a_{-i,t}|\lambda(h_i^{R,t}),\mathring{h}_{-i}^{R,t}) = \prod_{j\neq i}\left((1-\varepsilon_k)\,\hat{\sigma}_j(\mathring{h}_j^{R,t})(a_{j,t})+\frac{\varepsilon_k}{|A_{j,t}|}\right).$$

Defining

$$\tilde{\varphi}_k\left(\mathring{h}_i^{R,t},m_{i,t},a_{i,t},\zeta_{i,t}\right) = \left(\begin{array}{c}
1_{\{\zeta_{i,t}=0,m_{i,t}=a_{i,t}\}}\left(1-\sqrt{\varepsilon_k}\right)\left(1-\left(\frac{1}{k}\right)^{T+1}\right)\hat{\sigma}_i(\mathring{h}_i^{R,t})(a_{i,t}) \\
+1_{\{\zeta_{i,t}=0,m_{i,t}=\star\}}\sqrt{\varepsilon_k}\left(1-\left(\frac{1}{k}\right)^{T+1}\right)\left(\left(1-\sqrt{\varepsilon_k}\right)\hat{\sigma}_i(\mathring{h}_i^{R,t})(a_{i,t})+\frac{\sqrt{\varepsilon_k}}{|A_{i,t}|}\right) \\
+1_{\{\zeta_{i,t}=1,m_{i,t}=a_{i,t}\}}\left(\frac{1}{k}\right)^{T+1}\frac{1}{|A_{i,t}|-\left|B_i(\mathring{h}_i^{R,t})\right|}
\end{array}\right)$$

and

$$\begin{aligned}
\tilde{e}_k\left(\mathring{h}_{-i}^{R,t+1}\right) = & \sum_{\theta_{-i,t},\zeta_{-i,t}}\left(\begin{array}{l}
\prod_{j\neq i:\theta_{j,t}=0,\zeta_{j,t}=0}\left(1-\sqrt{\varepsilon_k}\right)\left(1-\left(\frac{1}{k}\right)^{T+1}\right)\hat{\sigma}_j(\mathring{h}_j^{R,t})(a_{j,t}) \\
\times\prod_{j\neq i:\theta_{j,t}=1,\zeta_{j,t}=0}\sqrt{\varepsilon_k}\left(1-\left(\frac{1}{k}\right)^{T+1}\right)\left(\left(1-\sqrt{\varepsilon_k}\right)\hat{\sigma}_j(\mathring{h}_j^{R,t})(a_{j,t})+\frac{\sqrt{\varepsilon_k}}{|A_{j,t}|}\right) \\
\times\prod_{j\neq i:\zeta_{j,t}=1}\left(\frac{1}{k}\right)^{T+1}\frac{1}{|A_{j,t}|-\left|B_j(\mathring{h}_j^{R,t})\right|}
\end{array}\right) \\
& -\Pr^{\hat{\sigma}^k,\hat{\mu}}(a_{-i,t}|\lambda(h_i^{R,t}),\mathring{h}_{-i}^{R,t}) \\
= & \left(\frac{1}{k}\right)^{T+1}\prod_{j\neq i}\left(\frac{1}{|A_{j,t}|-\left|B_j(\mathring{h}_j^{R,t})\right|}-(1-\varepsilon_k)\,\hat{\sigma}_j(\mathring{h}_j^{R,t})(a_{j,t})-\frac{\varepsilon_k}{|A_{j,t}|}\right),
\end{aligned}$$

and summing (32) over $(\theta_{-i,t}, \zeta_{-i,t})$, we have

$$\begin{aligned} & \delta_k(h_i^{R,t+1}, \zeta_i^t, \mathring{h}_{-i}^{R,t+1} | \omega_T \in \Omega_0) \\ = \; & \varphi_k^R(h_i^{R,t}, \zeta_i^t) \left(\Pr^{\hat{\sigma}^k, \hat{\mu}}(\lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t}) + e_k^R \left(h_i^{R,t}, \zeta_i^t, \mathring{h}_{-i}^{R,t} \right) \right) \\ & \times \left(\Pr^{\hat{\sigma}^k, \hat{\mu}}(a_{-i,t} | \lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t}) + \tilde{e}_k \left(\mathring{h}_{-i}^{R,t+1} \right) \right) \\ & \times \tilde{\varphi}_k \left(\mathring{h}_i^{R,t}, m_{i,t}, a_{i,t}, \zeta_{i,t} \right) \times p(s_{t+1} | \mathring{h}^{R,t}, a_t). \end{aligned} \tag{33}$$

Next, define

$$\varphi_k^R(h_i^{R,t+1}, \zeta_i^{t+1}) = \varphi_k^R(h_i^{R,t}, \zeta_i^t) \times \frac{\tilde{\varphi}_k \left(\mathring{h}_i^{R,t}, m_{i,t}, a_{i,t}, \zeta_{i,t} \right)}{\Pr^{\hat{\sigma}^k, \hat{\mu}} \left(m_{i,t}^*(h_i^{R,t+1}), a_{i,t} | \mathring{h}_i^{R,t} \right)}.$$

We can write

$$\begin{aligned} & \delta_k(h_i^{R,t+1}, \zeta_i^{t+1}, \mathring{h}_{-i}^{R,t+1}) \\ = \; & \varphi_k^R(h_i^{R,t+1}, \zeta_i^{t+1}) \begin{pmatrix} \Pr^{\hat{\sigma}^k, \hat{\mu}}(\lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t}) \times \Pr^{\hat{\sigma}^k, \hat{\mu}}(a_{-i,t} | \lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t}) \\ \times \Pr^{\hat{\sigma}^k, \hat{\mu}} \left(m_{i,t}^*(h_i^{R,t+1}), a_{i,t} | \mathring{h}_i^t \right) \times p(s_{t+1} | \mathring{h}_{-i}^{R,t}, a_t) \\ + e_k^R \left(h_i^{R,t+1}, \zeta_i^t, \mathring{h}_{-i}^{R,t+1} \right) \end{pmatrix}, \end{aligned} \tag{34}$$

where $e_k^R \left(h_i^{R,t+1}, \zeta_i^{t+1}, \mathring{h}_{-i}^{R,t+1} \right)$ is defined to satisfy this equality given (33): that is,

$$\begin{aligned} & e_k^R \left(h_i^{R,t+1}, \zeta_i^{t+1}, \mathring{h}_{-i}^{R,t+1} \right) \\ = & \left(\begin{array}{c} \tilde{e}_k \left(\mathring{h}_{-i}^{R,t+1} \right) \Pr^{\hat{\sigma}^k, \hat{\mu}}(\lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t}) \\ + e_k^R \left(h_i^{R,t}, \zeta_i^t, \mathring{h}_{-i}^{R,t} \right) \left(\Pr^{\hat{\sigma}^k, \hat{\mu}}(a_{-i,t} | \lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t}) + \tilde{e}_k \left(\mathring{h}_{-i}^{R,t+1} \right) \right) \end{array} \right) \\ & \times \Pr^{\hat{\sigma}^k, \hat{\mu}} \left(m_{i,t}^*(h_i^{R,t+1}), a_{i,t} | \mathring{h}_i^{R,t} \right) \times p(s_{t+1} | \mathring{h}^{R,t}, a_t). \end{aligned}$$

Since the distribution of player i's message and action is determined by her own history,

$$\Pr^{\hat{\sigma}^k, \hat{\mu}} \left(m_{i,t}^*(h_i^{R,t+1}), a_{i,t} | \mathring{h}_i^{R,t} \right) = \Pr^{\hat{\sigma}^k, \hat{\mu}}(m_{i,t}^*(h_i^{R,t+1}), a_{i,t} | \lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t}, a_{-i,t}).$$

Hence,

$$\begin{aligned} & \Pr^{\hat{\sigma}^k, \hat{\mu}}(\lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t}) \times \Pr^{\hat{\sigma}^k, \hat{\mu}}(a_{-i,t} | \lambda(h_i^{R,t}), \mathring{h}_{-i}^{R,t}) \times \Pr^{\hat{\sigma}^k, \hat{\mu}} \left(m_{i,t}^*(h_i^{R,t+1}), a_{i,t} | \mathring{h}_i^{R,t} \right) \times p(s_{t+1} | \mathring{h}^{R,t}, a_t) \\ = \; & \Pr^{\hat{\sigma}^k, \hat{\mu}}(\lambda(h_i^{R,t+1}), \mathring{h}_{-i}^{R,t+1}). \end{aligned}$$

Substituting this into (34), we have

$$\delta_k(h_i^{R,t+1},\zeta_i^{t+1},\mathring{h}_{-i}^{R,t+1}) = \varphi_k^R(h_i^{R,t+1},\zeta_i^{t+1})\left(\Pr^{\hat{\sigma}^k,\hat{\mu}}(\lambda(h_i^{R,t+1}),\mathring{h}_{-i}^{R,t+1}) + e_k^R\left(h_i^{R,t+1},\zeta_i^t,\mathring{h}_{-i}^{R,t+1}\right)\right).$$

Finally, we have

$$\begin{aligned}\lim_k \frac{e_k^R\left(h_i^{R,t+1},\zeta_i^{t+1},\mathring{h}_{-i}^{R,t+1}\right)}{\left(\frac{1}{k}\right)^{T+1}} &\leq \lim_k \frac{\tilde{e}_k\left(\mathring{h}_{-i}^{R,t+1}\right)}{\left(\frac{1}{k}\right)^{T+1}} + \frac{e_k^R\left(h_i^{R,t},\zeta_i^t,\mathring{h}_{-i}^{R,t}\right)}{\left(\frac{1}{k}\right)^{T+1}}\left(1+\tilde{e}_k\left(\mathring{h}_{-i}^{R,t+1}\right)\right)\\ &\leq 1+t,\end{aligned}$$

where the last line uses $\lim_k \frac{\tilde{e}_k\left(\mathring{h}_{-i}^{R,t+1}\right)}{\left(\frac{1}{k}\right)^{T+1}} \leq 1$ (and hence $\tilde{e}_k\left(\mathring{h}_{-i}^{R,t+1}\right) \to 0$) and the inductive hypothesis that $\lim_k \frac{e_k^R\left(h_i^{R,t},\zeta_i^t,\mathring{h}_{-i}^{R,t}\right)}{\left(\frac{1}{k}\right)^{T+1}} \leq t$. Hence, (31) holds for $t+1$, as desired.

Proof of Lemma 12 We prove (29); the proof of (30) is analogous. From the definition of δ_k, $\delta_k\left(f^{\geq t^*},\mathring{h}_{-i}^{R,t}|h_i^{R,t},t^*,\zeta_i^{t^*}\right)$ equals

$$\frac{A\sum_{f_i^{<t^*},\theta_i^{t^*},\left(f_j^{<t^*},\theta_j^{t^*},m_j^{t^*}\right)_{j\neq i}}\pi_{t^*}\left(f^{\geq t^*},\mathring{h}^{R,t^*}\right)\Pr^\sigma\left(\mathring{h}^{R,t}|f^{\geq t^*},\mathring{h}^{R,t^*}\right)B_i\prod_{j\neq i}C_jD}{A\sum_{\mathring{h}_{-i}^{\prime R,t},f^{\prime\geq t}}\sum_{f_i^{<t^*},\theta_i^{t^*},\left(f_j^{<t^*},\theta_j^{t^*},m_j^{t^*}\right)_{j\neq i}}\pi_{t^*}\left(f^{\prime\geq t^*}\mathring{h}_i^{R,t^*},\mathring{h}_{-i}^{\prime R,t^*}\right)\Pr^\sigma\left(\mathring{h}_i^{R,t},\mathring{h}_{-i}^{\prime R,t}|f^{\prime\geq t^*},\mathring{h}_i^{R,t^*},\mathring{h}_{-i}^{\prime R,t^*}\right)B_i\prod_{j\neq i}C_j'D'}, \tag{35}$$

where the summation is taken over $f_i^{<t^*} \in \operatorname{supp}\hat{\mu}_i^{<t^*}$, $\theta_i^{t^*} \in \{0,1\}^{t^*-1}$, and $(f_j^{<t^*},\theta_j^{t^*},m_j^{t^*}) \in \operatorname{supp}\hat{\mu}_j^{<t^*} \times \{0,1\}^{t^*-1} \times \prod_{\tau=1}^{t^*-1}(A_{j,\tau}\cup\{\star\})\ \forall j$, and we define

$$\begin{aligned}&A = \left(\tfrac{1}{k}\right)^{t^*+(T+1)\left|\zeta_i^{t^*}\right|}, && B_i = \tfrac{1}{\#M_i(\mathring{h}_i^{t^*})},\\ &C_j = \tfrac{1}{\#M_j(\mathring{h}_j^{t^*})}, && C_j' = \tfrac{1}{\#M_j(\mathring{h}_j^{\prime t^*})},\\ &D = 1_{\left\{m^{t^*}\in M^{t^*}(\mathring{h}^{t^*},f^{<t^*},\theta^{t^*},\zeta_i^{t^*})\right\}}, && D' = 1_{\left\{m^{t^*}\in M^{t^*}(\mathring{h}_i^{t^*},\mathring{h}_{-i}^{\prime t^*},f^{<t^*},\theta^{t^*},\zeta_i^{t^*})\right\}}.\end{aligned}$$

Note that A and B_i cancel in (35). Moreover, we have

$$\begin{aligned}D &= D_i^{00}\times D_i^{01}\times D_i^{?1}\times\prod_{j\neq i}\left(D_j^0\times D_j^1\right),\\ D' &= D_i^{00}\times D_i^{01}\times D_i^{?1}\times\prod_{j\neq i}\left(D_j^{\prime 0}\times D_j^{\prime 1}\right),\end{aligned}$$

where

$$\begin{aligned}&D_i^{00} = 1_{\left\{m_{i,\tau}=f_{i,\tau}^{<t^*}(\mathring{h}_i^{R,\tau})\ \forall\tau\leq t^*-1\text{ s.t. }\zeta_{i,\tau}=\theta_{i,\tau}=0\right\}}, && D_i^{01} = 1_{\left\{m_{i,\tau}=\star\ \forall\tau\leq t^*-1\text{ s.t. }\zeta_{i,\tau}=0\text{ and }\theta_{i,\tau}=1\right\}},\\ &D_i^{?1} = 1_{\left\{m_{i,\tau}\in A_{i,\tau}\setminus B(\mathring{h}_i^{R,\tau})\ \forall\tau\leq t^*-1\text{ s.t. }\zeta_{i,\tau}=1\right\}}, && D_j^0 = 1_{\left\{m_{j,\tau}=f_{j,\tau}^{<t^*}(\mathring{h}_j^{R,\tau})\ \forall\tau\leq t^*-1\text{ s.t. }\theta_{j,\tau}=0\right\}},\\ &D_j^1 = 1_{\{m_{j,\tau}=\star\ \forall\tau\leq t^*-1\text{ s.t. }\theta_{j,\tau}=1\}}, && D_j^{\prime 0} = 1_{\left\{m_{j,\tau}=f_{j,\tau}^{<t^*}(\mathring{h}_j^{\prime R,\tau})\ \forall\tau\leq t^*-1\text{ s.t. }\theta_{j,\tau}=0\right\}},\\ &D_j^{\prime 1} = 1_{\{m_{j,\tau}=\star\ \forall\tau\leq t^*-1\text{ s.t. }\theta_{j,\tau}=1\}}.\end{aligned}$$

Since (i) $D_i^{00} D_i^{01} D_i^{?1}$ cancels in (35), (ii) these terms are independent of $f^{\geq t}$, $\mathring{h}_{-i}^{R,t}$, and $\left(f_j^{<t^*}, \theta_j^{t^*}, m_j^{t^*}\right)_{j \neq i}$, and (iii) these are the only terms that depend on $(f_i^{<t^*}, \theta_i^{t^*})$, we can simplify (35) by ignoring $D_i^{00} D_i^{01} D_i^{?1}$ and $\sum_{f_i^{<t^*}, \theta_i^{t^*-1}}$, obtaining

$$\frac{\sum_{\left(f_j^{<t^*}, \theta_j^{t^*}, m_j^{t^*}\right)_{j \neq i}} \pi_{t^*}\left(f^{\geq t^*}, \mathring{h}^{R,t^*}\right) \Pr^\sigma\left(\mathring{h}^{R,t} | f^{\geq t^*}, \mathring{h}^{R,t^*}\right) \prod_{j \neq i}\left(C_j D_j^0 D_j^1\right)}{\sum_{\mathring{h}'^t_{-i}, f'^{\geq t}} \sum_{\left(f_j^{<t^*}, \theta_j^{t^*}, m_j^{t^*}\right)_{j \neq i}} \pi_{t^*}\left(f'^{\geq t^*}, \mathring{h}_i^{R,t^*}, \mathring{h}'^{R,t^*}_{-i}\right) \Pr^\sigma\left(\mathring{h}_i^{R,t}, \mathring{h}'^{R,t}_{-i} | f'^{\geq t^*}, \mathring{h}_i^{R,t^*}, \mathring{h}'^{R,t^*}_{-i}\right) \prod_{j \neq i}\left(C'_j D'^0_j D'^1_j\right)}. \quad (36)$$

Since $D_j^0 D_j^1$ is the only term that depends on $\left(f_j^{<t^*}, \theta_j^{t^*}, m_j^{t^*}\right)$ and we have

$$C_j \sum_{f_j^{<t^*}, \theta_j^{t^*}, m_j^{t^*}} D_j^0 D_j^1 = 1$$

for each $h_j^{R,t}$ by (21), (36) equals (29).