The Revelation Principle in Multistage Games*

Takuo Sugaya and Alexander Wolitzky

Stanford GSB and MIT

June 29, 2020

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 show that in general the communication revelation principle fails for the solution concept of sequential equilibrium. However, it holds in important classes of games, including single-agent games, games with pure adverse selection, games with pure moral hazard, and a class of social learning games. For general multistage games, we establish that an outcome is implementable in sequential equilibrium if and only if it is implementable in a canonical Nash equilibrium in which players never take codominated actions. We also prove that the communication revelation principle holds for the more permissive solution concept of *conditional probability perfect Bayesian equilibrium*.

*For helpful comments, we thank Dirk Bergemann, Drew Fudenberg, Dino Gerardi, Shengwu Li, George Mailath, Roger Myerson, Alessandro Pavan, Harry Pei, Ludovic Renou, Adam Szeidl, Juuso Toikka, and Bob Wilson; seminar participants at Austin, Bocconi, Collegio Carlo Alberto, Cowles, Harvard-MIT, Kobe, LSE, NYU, UPenn, Princeton, Queen Mary, Stony Brook, UCL, UCLA, and UCSD; and the anonymous referees. Daniel Clark and Hideo Suehiro read the paper carefully and gave us excellent comments. Wolitzky acknowledges financial support from the NSF and the Sloan Foundation.

1 Introduction

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 canonical form: players communicate only their private information to the mediator, the mediator communicates only recommended actions to the players, and in equilibrium players report honestly and obey the mediator's recommendations. 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.

More recently, there has been a surge of interest in designing dynamic mechanisms and information systems.[1] Forges (1986) showed that the communication revelation principle (RP) is valid in multistage games under the solution concept of Nash equilibrium (NE). But NE is usually not a satisfactory solution concept in dynamic games: following Kreps and Wilson (1982), economists prefer solution concepts that require rationality even after off-path events and impose "consistency" restrictions on players' off-path beliefs, such as sequential equilibrium (SE) or various versions of perfect Bayesian equilibrium (PBE). And it is unknown whether the RP 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 current paper resolves this question. We show that in general multistage games the communication RP fails for SE. However, it holds in important classes of games, including single-agent games, games with pure adverse selection, games with pure moral hazard, and a class of social learning games. Our main result establishes that, in general multistage games, an outcome is implementable in SE if and only if it can be implemented in a canonical Nash equilibrium in which players never (on or off path) take *codominated actions*, which are actions that cannot be motivated by any belief compatible with a player's own information

[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), Pavan, Segal, and Toikka (2014), and Battaglini and Lamba (2019). For dynamic information design, see for example Kremer, Mansour, and Perry (2014), Ely, Frankel, and Kamenica (2015), Che and Hörner (2017), Ely (2017), Renault, Solan, and Vieille (2017), Ely and Szydlowski (2019), and Ball (2020).

and the presumption that her opponents will avoid codominated actions in the future. This is an extension to SE of the main result of Myerson (1986), which we review in Section 2.2.[2] We also show that the communication RP holds in general multistage games for the solution concept of *conditional probability perfect Bayesian equilibrium (CPPBE)*, a simple and relatively permissive version of PBE.

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, it suffices to calculate the set of outcomes implementable in *Nash* equilibrium *excluding codominated actions*, using *direct* communication.[3] These two sets are always the same, even though actually implementing some outcomes as sequential equilibria might require a richer communication system (that is, even though the communication RP is generally invalid for SE).

Let us preview the intuition for our key result: any outcome that can be implemented in a NE that excludes codominated actions is also implementable in SE. By definition, any non-codominated action can be motivated by some belief compatible 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 positive probability (along a sequence of strategy profiles converging to the equilibrium) and then honestly report their signals and actions to the mediator, and the mediator appropriately conditions his recommendations on the reports. An obstacle to this construction is that a player who trembles 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 (i.e., with probability converging to 1). Then, the desired belief can be generated by letting players believe that their opponents received promises to disregard their reports, trembled, and then reported truthfully, and that subsequently the mediator did not disregard their reports after all. However, to afford the mediator the ability to make such an advance promise, the communication system must be enriched with an extra message. Note that the mediator's "promise to ignore reports"

[2]Myerson's main result establishes the same characterization for the novel concept of *sequential communication equilibrium*, which is not the same as sequential equilibrium. See Section 2.2.

[3]The set of codominated actions itself can be calculated recursively. See Appendix A.

is made with equilibrium probability 0, so our construction is "canonical on path."[4] At the same time, the need to have this extra message available explains why the communication RP is invalid for SE.

In several important classes of games, the set of non-codominated actions in each period does not depend on the history of players' signals and actions. These include single-agent games, games of pure adverse selection, games of pure moral hazard, and a class of social learning games. In such games, the above obstacle to implementing non-codominated actions does not arise, and the communication RP holds for SE. Furthermore, SE and NE are outcome-equivalent in many of these classes.

By way of further motivation, 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 (1986) and then restrict attention to direct mechanisms without quite claiming that this is without loss of 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, PBE, and SE are outcome-equivalent in pure adverse selection settings like this one.[5]

Our simple positive results for games with one agent, pure adverse selection, or pure moral hazard imply that the subtleties at the heart of our paper are most relevant for multi-agent, multi-stage games with both adverse selection and moral hazard: that is, *multi-agent dynamic information design*. Papers on this topic include Gershkov and Szentes (2009),

[4] As this discussion indicates, it is important that our definition of sequential equilibrium allows the mediator to tremble. Section 5.2 discusses the case where the mediator cannot tremble.

[5] A caveat is that much of the dynamic mechanism design literature assumes continuous type spaces to facilitate the use of the envelope theorem, while we restrict attention to finite games to have a well-defined notion of sequential equilibrium. We discuss this point in Section 5.2. We also assume a finite time horizon.

Aoyagi (2010), Kremer, Mansour, and Perry (2014), Che and Hörner (2017), Halac, Kartik, and Liu (2017), Sugaya and Wolitzky (2017), Ely (2017), Doval and Ely (2020), and Makris and Renou (2020). Some of these papers prove versions of the RP directly, while others appeal to existing results with more or less precision. For example, Kremer, Mansour, and Perry (2014) do not specify a solution concept and state that the RP is established for their setting by Myerson; an implication of our Proposition 4 is that the RP is valid in their model for SE. We hope our results will find application in this emerging literature; to this end, we provide a compact summary at the end of the paper.

1.1 Example

We begin with an example that illustrates how letting the mediator make advance promises to disregard players' reports can expand the set of implementable outcomes.

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 $a_0 \in \{A, B\}$.

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

Player 1's payoff equals $1_{\{a_0 \neq a_1\}} - 1_{\{a_2 = P\}} - 3 \times 1_{\{a_1 = C\}}$, and player 2's payoff equals $-1_{\{(a_1,\theta) \neq (C,p) \wedge a_2 = P\}}$, where $1_{\{\cdot\}}$ denotes the indicator function. 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 signals and actions 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.

Our main result shows that the outcome of any NE that excludes codominated actions is implementable in SE. In this example, the action $a_2 = P$ is not codominated at the history following signal $s = 1$, because $a_2 = P$ is an optimal action for player 2 if $(a_1, \theta) = (C, p)$. Our result thus implies that $\frac{1}{2}(A, A, N) + \frac{1}{2}(B, B, N)$ is a SE outcome for some communication system. 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.[6]

Non-implementability in canonical SE Throughout the paper, by a "tremble" to a particular action a we mean a sequence of strategies converging to equilibrium along which a player (or the mediator) takes action a with positive probability converging to 0. In a canonical equilibrium, players who have not previously lied to the mediator obey all recommendations from the mediator, even those that the mediator sends only as the result of a tremble. Since action C is strictly dominated for player 1, this implies that action C can never be recommended in a canonical equilibrium, even as the result of a tremble.[7] Hence, the mediator can only ever recommend $m_1 \in \{A, B\}$. If player 1 trembles to 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 $\Pr\left((a_1, \theta) = (C, p) \,|s = 1, m_2 = P\right) \leq \frac{1}{2}$ (in the limit where trembles vanish). Hence, player 2 will not follow the recommendation $m_2 = P$ when $s = 1$, so the desired outcome is not implementable in a canonical SE.

Implementability in 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. This makes player 1 willing to truthfully report any pair (a_1, θ) after getting a free pass. 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

[6] For the details, see the proof of Proposition 3.

[7] That is, action C must lie outside the *mediation range* in any canonical equilibrium. See Section 2.2.

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 $a_0 = m_1$; 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 in the $k \to \infty$ limit: if $m_1 \in \{A, B\}$, then any deviation from $a_1 = m_1$ leads to $a_2 = P$ with limit probability 1 regardless of player 1's report; while if $m_1 = \star$, then $a_2 = N$ with limit 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, in the $k \to \infty$ limit, 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.

Remark 1 *The non-implementability result established above uses the fact that the mediator cannot recommend C in a canonical equilibrium, since player 1 will not obey such a recommendation. However, the target outcome $\frac{1}{2}(A, A, N) + \frac{1}{2}(B, B, N)$ can be implemented in a non-canonical equilibrium of the direct mechanism where $m_1 \in \{A, B, C\}$ (without introducing the extra message $\star$), by using message C as a stand-in for message $\star$. This trick always works when each player has a strictly dominated (more generally, codominated) action at every information set, but not more generally: in the proof of Proposition 3, we give an ex-*

ample where restricting attention to direct mechanisms (even without additionally requiring honesty and obedience) is with loss of generality.

The remainder of the paper is organized as follows. Section 2 describes the model and reviews some background theory, including the notion of codominated actions. Section 3 presents our results for SE: the communication RP fails in general but holds in some important special classes of games; and in general games an outcome is SE-implementable if and only if it is implementable in a canonical NE that excludes codominated actions. Section 4 defines CPPBE and shows that the communication RP holds for this more permissive solution concept. Section 5 summarizes our results and discusses possible extensions. All proofs are deferred to the print or online 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), and (iii) "trembles" by the mediator may be treated differently than trembles by the other players.[8] 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 $M_{i,t}$. For each i and t, let $S_i^t = \prod_{\tau=1}^{t-1} S_{i,\tau}$, let $S_t = \prod_{i=0}^{N} S_{i,t}$, and let $S^t = \prod_{\tau=1}^{t-1} S_\tau$, and analogously define A_i^t, A_t, A^t, R_i^t, R_t, R^t, M_i^t, M_t, and M^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

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

for the players). Note also the artificial assumption that the mediator "communicates with himself," which simplifies notation.

The timing within each period t is as follows:

1. A signal $s_t \in S_t$ is drawn with probability $p\left(s_t|s^t, a^t\right)$, where $(s^t, a^t) \in S^t \times A^t$ is the vector of past signals and actions. Player i observes $s_{i,t} \in S_{i,t}$, 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 $a_{i,t} \in A_{i,t}$.

For each t, denote the set of possible histories of signals and actions ("payoff-relevant histories") at the beginning of period t by

$$X^t = \left\{\left(s^t, a^t\right) \in S^t \times A^t : p\left(s_\tau|s^{\tau-1}, a^{\tau-1}\right) > 0 \ \ \forall \tau \leq t-1\right\}.$$

Similarly, denote the set of possible payoff-relevant histories after the period t signal realization s_t by

$$Y^t = \left\{\left(s^{t+1}, a^t\right) \in S^{t+1} \times A^t : p\left(s_\tau|s^{\tau-1}, a^{\tau-1}\right) > 0 \ \ \forall \tau \leq t\right\}.$$

For each i, let X_i^t and Y_i^t denote the projections of X^t and Y^t on $S_i^t \times A_i^t$ and $S_i^{t+1} \times A_i^t$, respectively. Note that, typically, $X^t \neq \prod_{i=0}^N X_i^t$ and $Y^t \neq \prod_{i=0}^N Y_i^t$. Assume without loss of generality that $S_{i,t} = \bigcup_{x^t \in X^t} \operatorname{supp} p_i\left(\cdot|x^t\right)$ for all i and t, where p_i denotes the marginal distribution of p.

Let $H^t = X^t \times R^t \times M^t$ denote the set of possible histories of signals, actions, reports, and messages ("complete histories") at the beginning of period t, with $H^1 = \emptyset$. Let $Z = H^{T+1}$ denote the set of terminal histories of the game. Given a complete history $h^t = (x^t, r^t, m^t) \in H^t$, let $\mathring{h}^t = x^t$ denote the projection of h^t onto X^t, the payoff-relevant component of H^t; and let $\dot{h}^t = (r^t, m^t)$ denote the projection of h^t onto $R^t \times M^t$, the payoff-irrelevant component of H^t. Let $X = X^{T+1}$ denote the set of payoff-relevant pure outcomes of the game. Let $u_i : X \to \mathbb{R}$ denote player i's payoff function, where u_0 is a constant function.

We refer to the tuple $\Gamma := (N, T, S, A, p, u)$ as the *base game* and refer to the pair $\mathfrak{C} := (R, M)$ as the *communication system*. The *implementation problem* asks, for a given base game Γ and a given equilibrium concept, which outcomes $\rho \in \Delta(X)$ arise in some equilibrium of the multistage game $G = (\Gamma, \mathfrak{C})$ for some communication system $\mathfrak{C}$? Such outcomes ρ are *implementable*.

We now introduce histories, strategies, and beliefs. For each i and t, let $H_i^t = X_i^t \times R_i^t \times M_i^t$ denote the set of player i's possible histories of signals, actions, reports, and messages at the beginning of period t. When a complete history $h^t \in H^t$ is understood, we let $h_i^t = (x_i^t, r_i^t, m_i^t)$ denote the projection of h^t onto H_i^t; that is, h_i^t is player i's information set. Conversely, let $H^t[h_i^t]$ denote the set of histories $h^t \in H^t$ with i-component h_i^t. Note that, typically, $H^t[h_i^t] \neq \prod_{j \neq i} H_j^t$. Let $\mathring{h}_i^t = x_i^t$ denote the payoff-relevant component of h_i^t, and let $\dot{h}_i^t = (r_i^t, m_i^t)$ denote the payoff-irrelevant component of h_i^t. We also let $H_i^{R,t} = Y_i^t \times R_i^t \times M_i^t$ and $H_i^{A,t} = Y_i^t \times R_i^{t+1} \times M_i^{t+1}$ denote the sets of *reporting* and *acting* histories for player i, respectively. $H^{R,t}[h_i^{R,t}]$, $H^{A,t}[h_i^{A,t}]$, $\mathring{h}_i^{R,t}$, $\mathring{h}_i^{A,t}$, $\dot{h}_i^{R,t}$, and $\dot{h}_i^{A,t}$ are similarly defined.

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 $\beta_{i,t}^A : H_i^{A,t} \to \Delta\left(H^{A,t}\right)$.[9] We write $\sigma_{i,t}^R\left(r_{i,t} | h_i^{R,t}\right)$ for $\sigma_{i,t}^R\left(h_i^{R,t}\right)(r_{i,t})$, and similarly for $\sigma_{i,t}^A$, $\beta_{i,t}^R$, and $\beta_{i,t}^A$. 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 $h_i^{R,t}$ or $h_i^{A,t}$ as its argument.

A *mediation plan* is a function $f = (f_t)_{t=1}^T$, 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.[10] A *mixed mediation plan* is a distribution $\mu \in \Delta(F)$, where F denotes the set of (pure) mediation plans. A *behavioral mediation plan* is a function $\phi = (\phi_t)_{t=1}^T$, where $\phi_t : R^{t+1} \times M^t \to \Delta(M_t)$ maps past reports and messages to current messages. Since the mediator can receive signals and

[9] 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.

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

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 plan depends on the solution concept under consideration. All three perspectives will be used in this paper. In contrast, we always view players as choosing behavioral strategies $(\sigma_i)_{i=1}^N$, and similarly view the mediator's report/action strategy σ_0 as a behavioral strategy.

Denote the probability distribution on Z induced by behavioral strategy profile $\sigma = (\sigma_i)_{i=0}^N$ and mediation plan f by $\Pr^{\sigma,f}$, and denote the corresponding distribution for a mixed or behavioral mediation plan by $\Pr^{\sigma,\mu}$ or $\Pr^{\sigma,\phi}$, respectively. Denote the corresponding probability distribution on X (the "outcome") 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 analogously. In contrast, we avoid the "bad" notation $\bar{u}_i(\sigma, \mu|h^t)$, which is not well-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 $i \neq 0$ and $\sigma_i' \in \Sigma_i$.[11] (Or put ϕ in place of μ; the definitions are equivalent by Kuhn's theorem, which implies that for any μ, there exists ϕ such that $\Pr^{\sigma,\mu} = \Pr^{\sigma,\phi}$ for all σ, and vice versa.)

[11] In the context of games with communication, a NE is also called a *communication equilibrium* (Forges, 1986).

2.2 Theoretical Background

Our results build on four concepts introduced by Myerson (1986), which we briefly review here: mediation ranges, conditional probability systems, sequential communication equilibria, and codomination.

A *mediation range* $Q = (Q_{i,t})_{i\neq 0,t}$ specifies a set of possible messages $Q_{i,t}\left(r_i^t, m_i^t, r_{i,t}\right) \subseteq M_{i,t}$ 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 compatible with mediation range Q by

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

Say that a reporting history $h_i^{R,t} \in H_i^{R,t}$ is compatible with mediation range Q if $m_{i,\tau} \in Q_{i,\tau}\left(r_i^\tau, m_i^\tau, r_{i,\tau}\right)$ for all $\tau < t$; similarly, an acting history $h_i^{A,t} \in H_i^{A,t}$ is compatible with mediation range Q if $m_{i,\tau} \in Q_{i,\tau}\left(r_i^\tau, m_i^\tau, r_{i,\tau}\right)$ for all $\tau \leq t$.

Given a base game $\Gamma = (N, T, S, A, p, u)$, the *direct communication system* $\mathfrak{C}^* = (R^*, M^*)$ 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. In a game with direct communication $G^* = (\Gamma, \mathfrak{C}^*)$, player i is *honest* at reporting history $h_i^{R,t}$ if she reports $r_{i,t} = (a_{i,t-1}, s_{i,t})$, and she is *obedient* at acting history $h_i^{A,t}$ if she plays $a_{i,t} = m_{i,t}$. Let Σ_i^* denote the set of strategies for player i in game G^*. Let $G^*|_Q$ denote the game where, at each history for the mediator $Y_0^t \times R^{t+1} \times M^t$, the mediator is restricted to sending messages $m_{i,t} \in Q_{i,t}\left(r_i^t, m_i^t, r_{i,t}\right)$ for each i.

Note that players will obey the mediator's recommendations even after trembles by the mediator only if the possibility that the mediator trembles to recommending unmotivatable actions is excluded. This can be achieved by restricting the mediation range. Given a game with direct communication G^* and a mediation range Q, the *fully canonical strategy profile* in $G^*|_Q$, denoted σ^*, is defined by letting players behave honestly and obediently at all histories compatible with Q. 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. 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 $\mathfrak{C}_0 = \mathfrak{C}_0^*$ and $\sigma_0 = \sigma_0^*$.

With direct communication, given a mediation plan f and a payoff-relevant history $x^t = (s^t, a^t)$, the unique complete history h^t compatible with the mediator following f and all players reporting honestly in every period $\tau < t$ satisfies $r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$ and $m_{i,\tau} = f_{i,\tau}(r^{\tau+1})$ for all i and $\tau < t$. Denote this history by $\bar{h}(f, x^t)$. Similarly, we denote by $\bar{h}(x^t)$ the unique complete history h^t compatible with all players reporting honestly and acting obediently in every period $\tau < t$.

Denote the set of terminal histories in G^* compatible with mediation range Q and honest behavior by the players by

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

Analogously define $\tilde{Z}^{R,t}|_Q$ and $\tilde{Z}^{A,t}|_Q$. Denote the set of pairs $(f, z) \in F \times \tilde{Z}|_Q$ such that terminal history z is compatible with mediation plan f by

$$\tilde{\mathcal{Z}}|_Q = \left\{ (f, z) \in F \times \tilde{Z}|_Q : m_t = f_t\left(r^{t+1}\right) \quad \forall t \right\}.$$

Analogously define $\tilde{\mathcal{Z}}^{R,t}|_Q$, and $\tilde{\mathcal{Z}}^{A,t}|_Q$. Denote the subset of $\tilde{\mathcal{Z}}|_Q$ with period-t reporting history $h^{R,t}$ by $\tilde{\mathcal{Z}}\left[h^{R,t}\right]|_Q$. $\tilde{\mathcal{Z}}[h^{A,t}]|_Q$, $\tilde{\mathcal{Z}}[h_i^{R,t}]|_Q$, and $\tilde{\mathcal{Z}}[h_i^{A,t}]|_Q$ are similarly defined.

A *conditional probability system (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 $\mu(A|B)\,\mu(B|C) = \mu(A|C)$. Theorem 1 of Myerson (1986) shows that $\mu(\cdot|\cdot)$ is a CPS on Ω if and only if it is the limit of conditional probabilities derived by Bayes' rule along a sequence of completely mixed probability distributions on Ω (see also Rènyi, 1955). Given a CPS $\bar{\mu}$ on $\tilde{\mathcal{Z}}|_Q$, $f \in F$, $\{f\} \times Y \subset \tilde{\mathcal{Z}}|_Q$, and $\{f\} \times Y' \subset \tilde{\mathcal{Z}}|_Q$, we write $\bar{\mu}(f) = \sum_{(f,z) \in \tilde{\mathcal{Z}}|_Q} \bar{\mu}(f, z)$ and $\bar{\mu}(Y|f, Y') = \sum_{y \in Y} \bar{\mu}(y|f, Y')$. A *sequential communication equilibrium (SCE)* is then a mixed mediation plan $\mu \in \Delta(F)$ in a direct-communication game G^* together with a mediation range Q and a CPS $\bar{\mu}$ on $\tilde{\mathcal{Z}}|_Q$ such that

- *[CPS consistency]* For all $f \in F|_Q$, t, $h^{R,t} = (s^{t+1}, r^t, m^t, a^t) \in \tilde{Z}^{R,t}|_Q$, $h^{A,t} = (s^{t+1}, r^{t+1}, m^{t+1}, a^t) \in \tilde{Z}^{A,t}|_Q$, m_t, a_t, and s_{t+1} such that $\left(f, h^{R,t}\right) \in \tilde{\mathcal{Z}}^{R,t}|_Q$ and $\left(f, h^{A,t}\right) \in \tilde{\mathcal{Z}}^{A,t}|_Q$, we have

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

(Here the first argument of $\bar{\mu}(\cdot|\cdot)$ must be read as a subset of $\tilde{\mathcal{Z}}|_Q$. For example, $\bar{\mu}\left(r_t | f, h^{R,t}\right) = \sum_{z' : (f,z') \in \tilde{\mathcal{Z}}[h^{R,t}, r_t]|_Q} \bar{\mu}\left((f, z') | f, h^{R,t}\right)$.)

- *[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 $m_{i,\tau} \in Q_{i,\tau}(r_i^\tau, m_i^\tau, r_{i,\tau})$ for all $\tau < t$, we have

$$\sum_{(f, h^{R,t}) \in \tilde{\mathcal{Z}}[h_i^{R,t}]|_Q} \bar{\mu}\left(f, h^{R,t} | h_i^{R,t}\right) \bar{u}_i\left(\sigma^*, f | h^{R,t}\right) \geq \sum_{(f, h^{R,t}) \in \tilde{\mathcal{Z}}[h_i^{R,t}]|_Q} \bar{\mu}\left(f, h^{R,t} | h_i^{R,t}\right) \bar{u}_i\left(\sigma_i', \sigma_{-i}^*, f | h^{R,t}\right). \tag{1}$$

- *[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 $m_{i,\tau} \in Q_{i,\tau}(r_i^\tau, m_i^\tau, r_{i,\tau})$ for all $\tau \leq t$, we have

$$\sum_{(f, h^{A,t}) \in \tilde{\mathcal{Z}}[h_i^{A,t}]|_Q} \bar{\mu}\left(f, h^{A,t} | h_i^{A,t}\right) \bar{u}_i\left(\sigma^*, f | h^{A,t}\right) \geq \sum_{(f, h^{A,t}) \in \tilde{\mathcal{Z}}[h_i^{A,t}]|_Q} \bar{\mu}\left(f, h^{A,t} | h_i^{A,t}\right) \bar{u}_i\left(\sigma_i', \sigma_{-i}^*, f | h^{A,t}\right). \tag{2}$$

This definition of SCE is identical to Myerson's, except that Myerson defines the CPS $\bar{\mu}$ over $F|_Q \times X$ rather than $F|_Q \times \tilde{Z}|_Q$. To see that this difference is immaterial, note that SCE imposes sequential rationality only for players who have not lied to the mediator, and $\tilde{Z}|_Q \subset X \times R^{T+1} \times M^{T+1}$ is the set of history profiles at which all players have not lied to the mediator. Therefore, for any $(f, h^t) \in F|_Q \times \tilde{Z}|_Q$ such that $\bar{\mu}(f, h^t | h_i^t) > 0$ for some h_i^t where sequential rationality is imposed, the complete history h^t is uniquely determined by f and the payoff-relevant history $x = \mathring{h}^t$ by specifying that all players have always reported truthfully: that is, $h^t = \bar{h}(f, x)$. Hence, the two definitions are equivalent.

Remark 2 *Note that a CPS over $F|_Q \times \tilde{Z}|_Q$ is equivalent to a CPS over terminal nodes in the tree of an alternative game where first the mediator chooses a (pure) mediation plan f and a copy of the original game G follows each choice of f, where paths inconsistent with the mediator's initial choice of f are deleted and players' information sets with the same history of messages from the mediator are merged. The reader may ask why Myerson considers CPS's over $F|_Q \times \tilde{Z}|_Q$ (or equivalently $F|_Q \times X$) rather than only $\tilde{Z}|_Q$. The reason is that specifying a CPS over $F|_Q \times \tilde{Z}|_Q$ implicitly lets the mediator tremble over strategies rather than actions (i.e., in normal form rather than independently at each information set), which allows a wider range of off-path expectations of future mediator behavior. Without this additional flexibility, Myerson's characterization of SCE outcomes would not be valid.*[12]

Myerson characterizes SCE in terms of codominated actions. The set of codominated actions for player i at payoff-relevant history $y_i^t \in Y_i^t$, denoted $\mathfrak{D}_{i,t}(y_i^t) \subset A_{i,t}$, can be given either an recursive or a fixed point definition. Here we give the fixed point definition, which is more concise. We give the recursive definition, which may be more useful for calculating the correspondence $\mathfrak{D}$ in applications, in Appendix A.

Fix a direct-communication game G^*. For any correspondence $\mathfrak{B}$ that specifies a set of actions $\mathfrak{B}_{i,t}(y_i^t) \subset A_{i,t}$ for each $i \neq 0$, t, and $y_i^t \in Y_i^t$, let $E^t(\mathfrak{B}) = \{f \in F : f_{i,\tau}(r^{\tau+1}) \notin \mathfrak{B}_{i,\tau}\left(r_i^{\tau+1}\right) \ \forall i, \tau > t, r^{\tau+1} \in R^{\tau+1}\}$ be the set of mediation plans that avoid actions in $\mathfrak{B}$ after period t, with the convention that $E^T(\mathfrak{B}) = F$,[13] and let $\phi^t(\mathfrak{B}) = \{(f, y^t) \in F \times Y^t : \exists i \neq 0$ s.t. $f_{i,t}(y^t) \in \mathfrak{B}_{i,t}(y_i^t)\}$ be the set of mediation plans f and payoff-relevant history profiles y^t such that f recommends an action in $\mathfrak{B}_{i,t}(y_i^t)$ to some player i at y^t. Such a correspondence $\mathfrak{B}$ is a *codomination correspondence* if, for every period t and every probability distribution $\pi \in \Delta(F \times Y^t)$ satisfying (i) $\pi(E^t(\mathfrak{B}) \times Y^t) = 1$ and (ii) $\pi(f, y^t) > 0$ for some $(f, y^t) \in$

[12]Roughly speaking, when CPS's are defined over $F|_Q \times \tilde{Z}|_Q$, a player can believe that her own past deviations are inherently correlated with future deviations by the mediator. This is not possible when CPS's are defined only over $\tilde{Z}|_Q$. Indeed, if the SCE definition were strengthened by defining CPS's only over $\tilde{Z}|_Q$, the proof of Proposition 3 could be adopted to give a counterexample to the claim that every SE-implementable outcome (and hence, every outcome of a NE in which players avoid codominated actions) arises in a SCE.

[13]In this definition, let $\mathfrak{B}_{i,\tau}\left(r_i^{\tau+1}\right) = \emptyset$ if $r_i^{\tau+1} \in R_i^{\tau+1} \setminus Y_i^{\tau}$.

$\phi^t(\mathfrak{B})$, there exists $i \neq 0$, y_i^t, $a_{i,t} \in \mathfrak{B}_{i,t}(y_i^t)$, and $\sigma_i' \in \Sigma_i$ such that

$$\sum_{\substack{(f,y^t)\in F\times Y^t,\\ f_{i,t}(y^t)=a_{i,t}}} \pi(f,y^t)\,\bar{u}_i\left(\sigma^*, f|\bar{h}(f,y^t)\right) < \sum_{\substack{(f,y^t)\in F\times Y^t,\\ f_{i,t}(y^t)=a_{i,t}}} \pi(f,y^t)\,\bar{u}_i\left(\sigma_i', \sigma_{-i}^*, f|\bar{h}(f,y^t)\right).$$

That is, if there is positive probability that some player is recommended a codominated action in period t, but zero probability that any player will be recommended a codominated action after period t, then some (possibly other) player has a profitable deviation in the event that she is recommended a codominated action in period t. The correspondence $\mathfrak{D}$ is then defined as the union of all codomination correspondences.[14]

Myerson's main result is that an outcome arises in a SCE if and only if it arises in a fully canonical NE in which players never take codominated actions.

Proposition 1 (Myerson (1986; Theorem 2, Lemma 1)) *For any base game Γ and any outcome $\rho \in \Delta(X)$, there exists a SCE $(\mu, Q, \bar{\mu})$ satisfying $\rho = \rho^{\sigma^*,\mu}$ if and only if there exist a NE (σ^*, μ) in $G^*|_Q$ satisfying $\rho = \rho^{\sigma^*,\mu}$ and $Q_{i,t}\left(r_i^{t+1}, m_i^t\right) \cap \mathfrak{D}_{i,t}(r_i^{t+1}) = \emptyset$ for all $i \neq 0$, t, r_i^{t+1}, and m_i^t.*

The set of SCE outcomes can be calculated as follows: For every player i and payoff-relevant history y_i^t, calculate the set of codominated actions $\mathfrak{D}_{i,t}(y_i^t)$. Delete these actions from the game tree. Then calculate the set of canonical NE (i.e., communication equilibrium) outcomes in the resulting game: as is well-known, this set is a compact polyhedron, defined by a finite set of linear inequalities (the incentive constraints).[15]

[14] We occasionally extend the domain of $\mathfrak{D}_{i,t}$ from Y_i^t to R_i^{t+1} by letting $\mathfrak{D}_{i,t}\left(\tilde{r}_i^{t+1}\right) = \emptyset$ for all $\tilde{r}_i^{t+1} \in R_i^{t+1} \setminus Y_i^t$.

[15] Establishing the validity of this algorithm requires one fact beyond Proposition 1: the set of NE in $G^*|_Q$ in which codominated actions are never played equals the set of NE in the game tree where codominated actions are deleted. Since the latter set obviously includes the former, this amounts to showing that deleting codominated actions never relaxes an incentive constraint: that is, to verify incentive compatibility in $G^*|_Q$, it suffices to consider only deviant strategies that never take codominated actions. This intuitive fact is established by Myerson (1986; Theorem 3).

2.3 Communication Revelation Principle: Definition

In a direct-communication game G^*, 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]* $\sigma_{i,t}^R\left(h_i^{R,t}\right) = (a_{i,t-1}, s_{i,t})$ for all $h_i^{R,t} \in H_i^{R,t}$ such that $r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$ and $m_{i,\tau} \in Q_{i,\tau}\left(r_i^\tau, m_i^\tau, r_{i,\tau}\right)$ for all $\tau < t$.

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

The communication RP states that it is without loss to restrict attention to direct communication systems, and furthermore to restrict attention to canonical strategies.[16]

Communication Revelation Principle Fix an equilibrium concept. For any game $(\Gamma, \mathfrak{C})$, any outcome $\rho \in \Delta(X)$ that arises in any equilibrium of $(\Gamma, \mathfrak{C})$ also arises in a canonical equilibrium of $(\Gamma, \mathfrak{C}^*)|_Q$ for some mediation range Q.

Forges (1986) established the communication RP for NE. For this result, it is not necessary to restrict the mediator's messages via a mediation range; it suffices to consider the *unrestricted mediation range* Q^U given by $Q_{i,t}^U\left(r_i^{t+1}, m_i^t\right) = A_{i,t}$ for all $i \neq 0$, t, r_i^{t+1}, and m_i^t.

Proposition 2 (Forges (1986; Proposition 1)) *The communication RP holds for NE, with the unrestricted mediation range.*

As we build on this result, we give a proof in Appendix B.[17] The intuition is that, in any game $(\Gamma, \mathfrak{C})$, 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 $\mathfrak{C}$, and then recommends actions to the player. Each player may

[16] 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). Our results show that enriching the communication system in this way does not expand the set of implementable outcomes. Townsend's motivation was to formulate incentive constraints in terms of one-shot deviations. In contrast, we follow Myerson in considering multi-shot deviations, as in inequalities (1) and (2).

[17] Forges's proof is convincing but informal, as are all other proofs of this result that we are aware of (e.g., pp. 106–107 of Mertens, Sorin, and Zamir, 2015).

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 direct-communication game $(\Gamma, \mathfrak{C}^*)$, where player i's personal mediator now automatically executes its equilibrium communication strategy from game $(\Gamma, \mathfrak{C})$. Then it remains optimal for each player to be honest and obedient, as a player has access to fewer deviations when she cannot directly control her personal mediator.

3 Sequential Equilibrium

3.1 Definition

Our definition of sequential equilibrium in a multistage 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, σ_i', and $h_i^{R,t}$, we have

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

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

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

- *[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 $i \neq 0$, t, $h_i^{R,t} \in H_i^{R,t}$, and $h^{R,t} \in H^{R,t}[h_i^{R,t}]$; 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 $i \neq 0$, t, $h_i^{A,t} \in H_i^{A,t}$, and $h^{A,t} \in H^{A,t}[h_i^{A,t}]$.

In this definition, the mediator takes a behavioral strategy and trembles independently at every information set, just like each of the players.[18]

3.2 Failure of Communication Revelation Principle

Our first substantive result is a negative one: the communication RP is generally invalid for SE, and even restricting attention to direct communication systems (without necessarily also restricting attention to canonical strategies) is generally with loss of generality.

Proposition 3 *The communication RP does not hold for SE. Furthermore, there exists a game* $(\Gamma, \mathfrak{C})$ *and an outcome* $\rho \in \Delta(X)$ *that arises in a SE of* $(\Gamma, \mathfrak{C})$ *but not in any SE of* $(\Gamma, \mathfrak{C}^*)$.

The failure of the RP for SE was previewed in the introduction. The stronger result that restricting attention to direct communication systems is with loss is proved by extending the opening example so 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 $\star$, so the set of possible messages must be expanded.

3.3 Special Classes of Games

While the communication RP is invalid for SE in general multistage games, we show that it does hold in several leading classes of games. Moreover, NE and SE are outcome-equivalent in many of these classes.

[18] An alternative definition, where the mediator cannot tremble at all, yields a more restrictive version of sequential equilibrium, for which our main results do not hold. We discuss this possibility in Section 5.2.

First, the communication RP holds and NE and SE are outcome-equivalent under a *full support* condition: any NE outcome distribution under which no player can perfectly detect another's unilateral deviation is a canonical SE outcome distribution. This result is not very surprising, but the formal proof is not completely straightforward.

Second, the communication RP holds and NE and SE are outcome-equivalent in *single-agent settings*. This is 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 and NE and SE are outcome-equivalent in the following class of *social learning games*: A state $\omega \in \Omega$ is drawn with probability $\hat{p}(\omega)$ at the beginning of the game. Given the state ω and a payoff-relevant history $x^t = (s^t, a^t)$, period-t signals s_t are drawn with probability $\hat{p}(s_t|\omega, x^t)$. We assume that, for each player i, there is a period t_i such that $|A_{i,t}| = 1$ for all $t \neq t_i$ and $|S_{i,t}| = 1$ for all $t < t_i$: that is, each player is "active" in only a single period. Player i's final payoff $\hat{u}_i(\omega, a_{i,t_i})$ depends only on the state ω and her own action. Such games are included in our model by letting $p(s_t|x^t) = \sum_\omega \hat{p}(\omega)\hat{p}(s_t|\omega, x^t)$ and $u_i(x) = \sum_\omega \Pr(\omega|x)\hat{u}_i(\omega, a_{i,t_i})$ (where a_{i,t_i} is i's action at outcome x).[19] The model of Kremer, Mansour, and Perry (2014) lies in this class.

Fourth, the communication RP holds and NE and SE are outcome-equivalent in games of *pure adverse selection*: $|A_{i,t}| = 1$ 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).

Fifth, the communication RP holds (but NE and SE are not outcome-equivalent) in games of *pure moral hazard*: for each $i \neq 0$, there exist $(\tilde{u}_{i,t})_{t=1}^T$ and $(\tilde{p}_t)_{t=1}^T$ such that $u_i(x) = \sum_{t=1}^T \tilde{u}_{i,t}(a_t)$, and $p(s_t|y^{t-1}, a_{t-1}) = \tilde{p}_t(s_t|a_{t-1})$. In a pure moral hazard game, payoffs are additively separable across periods, and signals are payoff-irrelevant and time-separable. This implies that the distribution of future payoff-relevant outcomes is independent of the realization of past payoff-relevant outcomes, conditional on the path of future actions. If these assumptions were violated, a player's payoff-relevant history would consti-

[19] Here the dependence of $u_i(x)$ on ω is accommodated by allowing $|S_{i,t}| > 1$ for $t > t_i$.

tute a "hidden state" that the mediator might need to elicit, possibly leading to a failure of the communication RP. Pure moral hazard games include finitely repeated games with complete information.

Proposition 4 *The following hold:*

1. *For any game, if* (σ, ϕ) *is a NE and* $\operatorname{supp} \rho_i^{\sigma,\phi} = \bigcup_{j \neq i,0} \bigcup_{\sigma'_j \in \Sigma_j} \operatorname{supp} \rho_i^{\sigma'_j, \sigma_{-j}, \phi}$ *for all* $i \neq 0$*, then* $\rho^{\sigma,\phi}$ *is a canonical SE outcome.*[20]

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

3. *In social learning games, any NE outcome is a canonical SE outcome.*

4. *In games of pure adverse selection, any NE outcome is a canonical SE outcome.*

5. *In games of pure moral hazard, the communication RP holds for SE.*

The logic of these results is as follows:

Parts 1 and 2 are intuitive. In any NE, each player's strategy is sequentially rational at on-path histories, and each player's strategy at off-path histories that follow her own deviation can be changed to a sequentially rational strategy without affecting other players' on-path incentives. Under the full-support condition, every history is either on path or follows a player's own deviation. Hence, any NE can be transformed into an SE by changing each player's strategies at histories that follow her own deviation.

In social learning games, since each player moves once and there are no payoff externalities, changing a player's off-path strategy never affects other players' on-path incentives. So again NE and SE are outcome-equivalent.

Part 4 follows from noting that the construction in Proposition 5 (in the next subsection) is canonical in pure adverse selection games.

Finally, in pure moral hazard games, the set of codominated actions in a given period t does not depend on the payoff-relevant history y^t. Hence, the mediator does not need to elicit information about y^t to motivate all non-codominated actions in period t. Under this condition, we show that the communication RP is valid for SE.[21]

[20] Here, $\rho_i^{\sigma,\phi}$ denotes the marginal distribution of $\rho^{\sigma,\phi}$ on X_i.

[21] The set of codominated actions is also independent of the payoff-relevant history for social learning

3.4 Characterization of SE-Implementable Outcomes

Our main result is that an outcome is implementable in SE if and only if it arises in an SCE, or equivalently if and only if it arises in a canonical NE that excludes codominated actions.

Define the *pseudo-direct communication system* $\mathfrak{C}^{**} = (R^*, M^{**})$ 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. Under pseudo-direct communication, in every period a single extra message from the mediator to each player is permitted.

Proposition 5 *For any base game, an outcome is SE-implementable if and only if it arises in a SCE. In addition, every such outcome arises in a SE with pseudo-direct communication.*

The "easy" implication of Proposition 5 is that every SE-implementable outcome arises in a SCE: we abbreviate this statement as $SE \subset SCE$. This follows from Propositions 6 through 8 in Section 4.

The "hard" implication is that every SCE outcome is SE-implementable (and in particular can be implemented with pseudo-direct communication): that is, $SCE \subset SE$. 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.[22]

Message $\star$ corresponds to the "free pass" in the opening example. As in that example, the role of message $\star$ is to cause a player to 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

games and pure adverse selection games. Thus, the proof of Part 5 of Proposition 4 also implies that the communication RP holds for SE in social learning games and pure adverse selection games. However, Parts 3 and 4 of Proposition 4 establish the stronger result that NE and SE are outcome-equivalent in such games.

[22] 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, a player cannot believe that her own past deviations are inherently correlated with past or future deviations by other players or the mediator. This makes it hard to ensure that previously disobedient players are honest.

beliefs are those required to motivate that action. For instance, in the example, when player 2 receives recommendation $m_2 = P$, he believes that the mediator trembled first to $m_1 = \star$ and then to $m_2 = P$ following $(\hat{a}_1, \hat{\theta}) = (C, p)$, which generates the belief required 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.

We end this section by sketching the proof of Proposition 5.

It is useful to first briefly review Myerson's proof that the outcome of every NE that excludes codominated actions is a SCE, as we build on this proof. Myerson first shows that every CPS is generated as the limit of beliefs induced from a sequence of full-support probability distributions over moves.[23] He then 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. Finally, he constructs another equilibrium where the mediator mixes this "motivating" 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 motivating SCE 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 trembling-hand perfect equilibrium (PE) in the unmediated game: that is, the limit as $\varepsilon \to 0$ of a sequence of NE in the unmediated, ε-constrained game where each player is required to take each action at each history with independent probability at least ε. We let the convergence of ε to 0 be slow in comparison to other trembles we will introduce: that is, action trembles in the PE are relatively likely. In the SE we construct in the mediated game, the mediator uses the off-path message $\star$ to signal to a player that the PE is being played. Since the PE is an equilibrium in the unmediated game, a player who receives message $\star$ believes that her future reports are almost-surely

[23] This differs from Kreps-Wilson consistency in that the move distributions may not be strategies. For example, some CPS's can be generated only by supposing that a player takes different actions at nodes in the same information set. This gap between SCE and SE has been noted before. See, for example, Kreps and Ramey (1987) and Fudenberg and Tirole (1991).

inconsequential, and thus reports honestly.[24] Specifically, when the mediator implements the PE, he recommends $m_{i,t} \in A_{i,t}$ according to the PE strategy of player i with probability $1-\sqrt{\varepsilon}$ and recommends $m_{i,t} = \star$ with probability $\sqrt{\varepsilon}$, independently across players. Player i obeys each recommendation $m_{i,t} \in A_{i,t}$ (with negligible trembling probability) and, and after message $\star$, takes $a_{i,t}$ according to her PE strategy but trembles with probability $\sqrt{\varepsilon}$. Since the mediator's tremble to $m_{i,t} = \star$ is independent across players, from the other players' perspectives, it is as if player i plays her PE strategy while trembling with probability $\sqrt{\varepsilon} \times \sqrt{\varepsilon} = \varepsilon$.

In order to provide on-path incentives, the mediator must also be able to recommend specific, non-codominated punishment actions off path. 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 showing that, since trembles in the PE 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 compatible 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 PE. These trembles are less likely than both the action trembles within the PE and the mediator trembles to mediation plans that rationalize non-codominated actions. There-

[24] This step is absent in Myerson's proof, as the SCE solution concept allows mediator trembles to be inherently correlated with player trembles about which the mediator has no information, so the mediator does not need to elicit information from players about their trembles.

fore, 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 PE (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'.[25]

To complete the construction, this "motivating equilibrium" (the mixture of the PE, 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 hence obey all non-codominated recommendations. Since the target NE excludes codominated actions, all on- and off-path recommendations are incentive compatible.

4 Conditional Probability Perfect Bayesian Equilibrium

Denote the set of terminal histories in G compatible with mediation range Q by

$$Z|_Q = \left\{ z \in Z : m_{i,t} \in Q_{i,t}\left(r_i^t, m_i^t, r_{i,t}\right) \quad \forall i,t \right\}.$$

Note that, in contrast to the set $\tilde{Z}|_Q$ defined in Section 2.2, the set $Z|_Q$ is defined for any communication system.[26] Analogously define $Z^{R,t}|_Q$ and $Z^{A,t}|_Q$.

Denote the set of pairs $(f,z) \in F|_Q \times Z|_Q$ such that terminal history z is compatible

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

[26] The set $\tilde{Z}|_Q$ also contained only histories at which players have been honest. This restriction is not well-defined with indirect communication and does not appear in the definition of $Z|_Q$.

with mediation plan f by

$$\mathcal{Z}|_Q = \left\{(f,z) \in F|_Q \times Z|_Q : m_t = f_t\left(r^{t+1}\right) \quad \forall t\right\}.$$

Analogously define $\mathcal{Z}^{R,t}|_Q$ and $\mathcal{Z}^{A,t}|_Q$. Denote the subset of $\mathcal{Z}|_Q$ with period-t reporting history $h^{R,t}$ by $\mathcal{Z}^{R,t}\left[h^{R,t}\right]|_Q$. $\mathcal{Z}[h^{A,t}]|_Q$, $\mathcal{Z}[h_i^{R,t}]|_Q$, and $\mathcal{Z}[h_i^{A,t}]|_Q$ are similarly defined.

We consider perfect Bayesian equilibria in which beliefs are derived from a common CPS on $\mathcal{Z}|_Q$. A *conditional probability perfect Bayesian equilibrium (CPPBE)* is a profile (σ,μ) together with a mediation range Q and a CPS $\bar{\mu}$ on $\mathcal{Z}|_Q$ such that

- *[CPS Consistency]* For all $f \in F|_Q$, t, $h^{R,t} = (s^{t+1}, r^t, m^t, a^t) \in Z^{R,t}|_Q$, $h^{A,t} = (s^{t+1}, r^{t+1}, m^{t+1}, a^t) \in Z^{A,t}|_Q$, m_t, a_t, and s_{t+1} such that $\left(f, h^{R,t}\right) \in \mathcal{Z}^{R,t}|_Q$ and $\left(f, h^{A,t}\right) \in \mathcal{Z}^{A,t}|_Q$, we have

$$\begin{array}{cc} \bar{\mu}(f) = \mu(f), & \bar{\mu}\left(r_t|f,h^{R,t}\right) = \prod_{i=0}^N \sigma_{i,t}^R\left(r_{i,t}|h_i^{R,t}\right), \\ \bar{\mu}\left(a_t|f,h^{A,t}\right) = \prod_{i=0}^N \sigma_{i,t}^A\left(a_{i,t}|h_i^{A,t}\right), & \bar{\mu}\left(s_{t+1}|f,h^{A,t},a_t\right) = p\left(s_{t+1}|\mathring{h}^{A,t},a_t\right), \\ \bar{\mu}\left(m_t|f,h^{R,t},r_t\right) = 1_{\{m_t = f_t(r^t,r_t)\}}. & \end{array} \tag{3}$$

- *[Sequential rationality of reports]* 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 $m_{i,\tau} \in Q_{i,\tau}\left(r_i^\tau, m_i^\tau, r_{i,\tau}\right)$ for all $\tau < t$, we have

$$\sum_{(f,h^{R,t}) \in \mathcal{Z}[h_i^{R,t}]|_Q} \bar{\mu}\left(f,h^{R,t}|h_i^{R,t}\right) \bar{u}_i\left(\sigma,f|h^{R,t}\right) \geq \sum_{(f,h^{R,t}) \in \mathcal{Z}[h_i^{R,t}]|_Q} \bar{\mu}\left(f,h^{R,t}|h_i^{R,t}\right) \bar{u}_i\left(\sigma_i',\sigma_{-i},f|h^{R,t}\right). \tag{4}$$

- *[Sequential rationality of actions]* 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 $m_{i,\tau} \in Q_{i,\tau}\left(r_i^\tau, m_i^\tau, r_{i,\tau}\right)$ for all $\tau \leq t$, we have

$$\sum_{(f,h^{A,t}) \in \mathcal{Z}^{A,t}[h_i^{A,t}]|_Q} \bar{\mu}\left(f,h^{A,t}|h_i^{A,t}\right) \bar{u}_i\left(\sigma,f|h^{A,t}\right) \geq \sum_{(f,h^{A,t}) \in \mathcal{Z}^{A,t}[h_i^{A,t}]|_Q} \bar{\mu}\left(f,h^{A,t}|h_i^{A,t}\right) \bar{u}_i\left(\sigma_i',\sigma_{-i},f|h^{A,t}\right). \tag{5}$$

To understand this definition, note that the conditional probabilities $\bar{\mu}\left(f,h^{R,t}|h_i^{R,t}\right)$ and $\bar{\mu}\left(f,h^{A,t}|h_i^{A,t}\right)$ in the sequential rationality conditions (4) and (5) correspond to player i's

beliefs in the alternative game discussed in Remark 2, where the mediator first chooses a pure mediation plan f and a copy of the original game G follows each choice of f. In the language of mechanism design, this corresponds to the designer first unobservably committing to a deterministic dynamic mechanism, and the players then updating their beliefs about the mechanism as they play it. 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) (4) disappears, and (iv) (5) reduces to the usual definition of sequential rationality. In the context of unmediated games, the CPPBE concept is not 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 SE in unmediated games. In contrast, we will see that CPPBE and SE are always outcome-equivalent in mediated games. The basic reason why independence conditions are not required to obtain equivalence with SE in mediated games is that the correlation allowed by CPPBE can be replicated through correlation in the mediator's messages.[27]

We first verify that CPPBE is more permissive than SE.[28]

Proposition 6 *Every SE-implementable outcome $\rho \in \Delta(X)$ is also CPPBE-implementable: that is, $SE \subset CPPBE$.*

Proof. Fix a SE (σ, ϕ, β), and let $\left(\sigma^k, \phi^k\right)_{k=1}^{\infty}$ be a sequence of full-support behavioral strategy profiles that converge to (σ, ϕ) and induce conditional probabilities that converge to β. By Kuhn's theorem, there exists an equivalent sequence of full-support profiles $\left(\sigma^k, \mu^k\right)_{k=1}^{\infty}$ converging to (σ, μ), where the mediator is now viewed as playing a mixed strategy. By Theorem 1 of Myerson (1986), the limit of the sequence of conditional probabilities on $\mathcal{Z}|_{Q^U}$ derived from $\left(\sigma^k, \mu^k\right)_{k=1}^{\infty}$ by Bayes' rule gives a CPS $\bar{\mu}$ on $\mathcal{Z}|_{Q^U}$. Since

[27] 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.

[28] As the proof shows, this holds even if the mediation range in the definition of CPPBE is required to be unrestricted, $Q = Q^U$. In fact, Lemma 8 in the online appendix shows that if an outcome ρ is implementable in CPPBE with some mediation range Q, it is also implementable in CPPBE with mediation range Q^U. It therefore would have been without loss to require $Q = Q^U$ in the definition of CPPBE; however, the current definition makes the connection between CPPBE and SCE more transparent.

$\bar{\mu}\left(f, h^{R,t}|h_i^{R,t}\right) = \beta\left(f, h^{R,t}|h_i^{R,t}\right)$ and $\bar{\mu}\left(f, h^{A,t}|h_i^{A,t}\right) = \beta\left(f, h^{A,t}|h_i^{A,t}\right)$, sequential rationality of (σ, ϕ, β) implies sequential rationality of $\left(\sigma, \mu, Q^U, \bar{\mu}\right)$. Hence, $\left(\sigma, \mu, Q^U, \bar{\mu}\right)$ is a CPPBE and induces the same outcome as (σ, ϕ, β). ■

The relationship between CPPBE and SCE is more subtle. The definition of a CPPBE *in the special case where the communication system is direct*, $\mathfrak{C} = \mathfrak{C}^*$, is similar to the definition of a SCE. (Of course, CPPBE is defined for arbitrary communication systems $\mathfrak{C}$). There are three 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 only for players who have not previously lied to the mediator.

3. SCE requires that a player who has not previously lied to the mediator believes with probability 1 that her opponents have also not previously lied to the mediator.

These properties of SCE were already noted by Myerson (1986), who argued informally that they should be without loss of generality.[29] The following result verifies this conjecture, by showing that the set of SCE outcomes (equivalently, the set of outcomes of canonical NE in which players avoid codominated actions) equals the set of outcomes implementable in a canonical CPPBE.

Proposition 7 *For any base game Γ, outcome $\rho \in \Delta(X)$, and mediation range Q, there exists a SCE $(\mu, Q, \bar{\mu})$ satisfying $\rho^{\sigma^*,\mu} = \rho$ if and only if there exist a canonical strategy profile σ and CPS $\bar{\mu}'$ such that $(\sigma, \mu, Q, \bar{\mu}')$ is a CPPBE in $(\Gamma, \mathfrak{C}^*)$ satisfying $\rho^{\sigma,\mu} = \rho$.*

Our final result establishes the communication RP for CPPBE.

[29] For example, he writes, "... there is nothing to prevent us from assuming that every player always assigns probability zero to the event that any other players have lied to the mediator... This begs the question of whether we could get a larger set of sequentially rational communication equilibria if we allowed players to assign positive probability to the event that others have lied to the mediator. Fortunately, by the revelation principle, this set would not be any larger. Given any mechanism in which a player lies to the mediator with positive probability after some event, there is an equivalent mechansim in which the player does not lie and the mediator makes recommendations exactly as if the player had lied in the given mechanism," (p. 342).

Proposition 8 *The communication RP holds for CPPBE, with mediation range equal to the set of all non-codominated actions:* $Q_{i,t}\left(r_i^{t+1}, m_i^t\right) = A_{i,t} \backslash \mathfrak{D}_{i,t}\left(r_i^{t+1}\right)$ *for all* i, t, r_i^{t+1}, *and* m_i^t.

To prove Propositions 7 and 8, we first establish that every SCE outcome is implementable in a canonical CPPBE. To show this, 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 result saves us from having to specify what a player does after she lies to the mediator, and it also lets us assume that a previously honest player always believes her opponents have also been honest. Given this simplification, every SCE can be viewed as a canonical quasi-CPPBE.[30]

We next establish that every (possibly non-canonical) CPPBE outcome is an SCE outcome: that is, $CPPBE \subset SCE$. This completes the proofs of both Propositions 7 and 8. Since every CPPBE is a NE, by Proposition 1 it suffices to show that codominated actions are never played in any CPPBE. We prove this as Lemma 10 in Online Appendix I. The logic of this result is that if a player is willing to take a certain action in a CPPBE, this action must be motivatable for some belief derived from a CPS, which implies that it is not codominated.

Combining the inclusion $CPPBE \subset SCE$ with Proposition 6, we see that every SE-implementable outcome arises in SCE: that is, $SE \subset CPPBE \subset SCE$. This proves the "easy" direction of Proposition 5.

In total, Propositions 5, 6, 7, and 8 show that $SCE \subset SE \subset CPPBE \subset SCE$. This implies that the characterization of SE-implementable outcomes in Proposition 5 applies equally to 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

[30] Quasi-strategies are also useful in proving Proposition 5.

(2017) "mutual PBE."[31]

5 Conclusion

5.1 Summary

Our main result is that to calculate the set of outcomes implementable in sequential equilibrium by any communication system in a multistage game, it suffices to calculate the set of outcomes of canonical Nash equilibria in which players avoid codominated actions.

We also show that the stronger *communication revelation principle* holds for conditional probability perfect Bayesian equilibrium, but not for sequential equilibrium. In particular, while the set of sequential equilibrium-implementable outcomes equals the set of outcomes of canonical Nash equilibria in which players avoid codominated actions, it may be necessary to allow one extra message to implement some of these outcomes as sequential equilibria.

There are however some important settings where the communication revelation principle does hold for sequential equilibrium. These include games where no player can perfectly detect another's deviation, games with a single agent, social learning games, and games of pure adverse selection or pure moral hazard.

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 a 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

[31] Watson (2017) defines plain PBE, which does not require the existence of a common CPS across players. His lectures (available at https://econweb.ucsd.edu/~jwatson/#other) further define "mutual PBE," which does require a common CPS.

(2007; Example 3) showed that, in general, not all SCE outcomes are implementable in MSE. However, we have shown that Claims 1 through 3 of Proposition 4 hold for MSE, as does a "virtual-implementation" version of Claim 4. In particular, 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.[32] We conjecture that the communication RP for CPPBE can be extended to infinite games under suitable measurability conditions. This extension is not immediate, because we build on Myerson's characterization of CPS's as limits of full-support move distributions, which does not apply in infinite games. Nonetheless, 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 a period (e.g., Doval and Ely, 2020). Modeling communication equilibrium in general extensive-form games is a long-standing unresolved issue, and different approaches are possible (e.g., Forges, 1986; von Stengel and Forges, 2008). Characterizing implementable outcomes in such games is 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] Ball, I. (2020), "Dynamic Information Provision: Rewarding the Past and Guiding the Future," *working paper*.

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

[32] For a recent attempt to extend SE to infinite games, see Myerson and Reny (2020).

[5] Battaglini, M. and R. Lamba (2019), "Optimal Dynamic Contracting: The First-Order Approach and Beyond," *Theoretical Economics*, 14, 1435-1482.

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

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

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

[9] Courty, P. and H. Li (2000), "Sequential Screening," *Review of Economic Studies*, 67, 697-717.

[10] Doval, L. and J. Ely (2020), "Sequential Information Design," *Econometrica*, Forthcoming.

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

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

[13] Ely, J.C. and M. Szydlowski, "Moving the Goalposts," *Journal of Political Economy*, 128, 468-506.

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

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

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

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

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

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

[20] Halac, M., N. Kartik, and Q. Liu (2017), "Contests for Experimentation," *Journal of Political Economy*, 125, 1523-1569.

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

[22] 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.

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

[24] Kreps, D.M. and G. Ramey (1987), "Structural Consistency, Consistency and Sequential Rationality," *Econometrica*, 55, 1331-1348.

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

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

[27] Makris, M. and L. Renou (2020), "Information Design in Multi-Stage Games," *working paper*.

[28] Mertens, J.-F., S. Sorin, and S. Zamir (2015), *Repeated Games*, Cambridge University Press.

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

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

[31] Myerson, R.B. and P.J. Reny (2020), "Perfect Conditional ε-Equilibria of Multi-Stage Games with Infinite Sets of Signals and Actions," *Econometrica*, 88, 495-531.

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

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

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

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

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

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

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

Appendix: Omitted Proofs

A Recursive Definition of Codomination

Fix a direct-communication game G^* and a set of mediation plans $\mathcal{F} \subset F$. For any t, given a correspondence $\mathfrak{A}'_t = \left(\mathfrak{A}'_{i,t}\right)_{i\neq 0}$ with $\mathfrak{A}'_{i,t} : Y_i^t \rightrightarrows A_{i,t}$, say that $\sigma'_i \in \Sigma_i$ is $\mathfrak{A}'_{i,t}$-*obedient* if σ'_i is honest and obedient at every history $h_i^{t'}$ with $t' \geq t$ such that $m_{i,t} \in \mathfrak{A}'_{i,t}(y_i^t)$. Say that a correspondence $\mathfrak{A}_t = (\mathfrak{A}_{i,t})_{i\neq 0}$ with $\mathfrak{A}_{i,t} : Y_i^t \rightrightarrows A_{i,t}$ is $(\mathcal{F}, \mathfrak{A}'_t)$-*motivatable* if there exists a distribution $\pi_t \in \Delta(\mathcal{F} \times Y^t)$ such that, for all $i \neq 0$ and $\mathfrak{A}'_{i,t}$-obedient σ'_i,

$$\sum_{(f,y^t)\in F\times Y^t} \pi_t\left(f, y^t\right) \bar{u}_i\left(\sigma^*, f|\bar{h}\left(f, y^t\right)\right) \geq \sum_{(f,y^t)\in F\times Y^t} \pi_t\left(f, y^t\right) \bar{u}_i\left(\sigma'_i, \sigma^*_{-i}, f|\bar{h}\left(f, y^t\right)\right),$$

and for all y_i^t and all $a_{i,t} \in \mathfrak{A}_{i,t}(y_i^t)$ there exist $f \in \mathcal{F}$ and $y_{-i}^t \in \prod_{j\neq i} Y_j^t$ satisfying $f_{i,t}\left(y_i^t, y_{-i}^t\right) = a_{i,t}$, $\left(y_i^t, y_{-i}^t\right) \in Y^t$, and $\pi_t\left(f, y_i^t, y_{-i}^t\right) > 0$.

We now characterize $\mathfrak{D}_{i,t}$ and its complement $\mathfrak{D}^{\mathrm{c}}_{i,t}$ by backward induction. We first recursively construct a finite sequence of correspondences $\mathfrak{A}_T^{[0]}, \mathfrak{A}_T^{[1]}, ..., \mathfrak{A}_T^{[L_T]}$ satisfying $\mathfrak{A}_T^{[L_T]} = \mathfrak{D}^{\mathrm{c}}_T$. Define $\mathfrak{A}_T^{[0]}$ by $\mathfrak{A}_T^{[0]}(y_i^T) = \emptyset$ for each $i \neq 0$ and $y_i^T \in Y_i^T$. Recursively, for each $l \geq 1$, let $\mathfrak{A}_T^{[l]}$ denote the union of all $\left(F, \mathfrak{A}_T^{[l-1]}\right)$-motivatable correspondences $\mathfrak{A}_T$.[33] Let L_T be the smallest integer l such that $\mathfrak{A}_T^{[l]} = \mathfrak{A}_T^{[l+1]}$, and let $\mathfrak{D}^{\mathrm{c}}_T = \mathfrak{A}_T^{[L_T]}$. (Such L_T exists because A is finite.)

By backward induction, for each $t < T$, let $\mathcal{F}_t \subset F$ denote the set of mediation plans f such that $f_{i,t'}\left(y_i^{t'}\right) \in \mathfrak{D}^{\mathrm{c}}_{i,t'}\left(y_i^{t'}\right)$ for all $i \neq 0$, $t' > t$, and $y_i^{t'} \in Y_i^{t'}$.[34] Let $\mathfrak{A}_t^{[0]}\left(y_i^t\right) = \emptyset$ for each $i \neq 0$ and $y_i^t \in Y_i^t$. Recursively, for each $l \geq 1$, let $\mathfrak{A}_t^{[l]}$ denote the union of all $\left(\mathcal{F}_t, \mathfrak{A}_t^{[l-1]}\right)$-motivatable correspondences $\mathfrak{A}_t$. Let L_t be the smallest integer l such that $\mathfrak{A}_t^{[l]} = \mathfrak{A}_t^{[l+1]}$, and let $\mathfrak{D}^{\mathrm{c}}_t = \mathfrak{A}_t^{[L_t]}$. Finally, let $\mathfrak{D}$ denote the complement of $\mathfrak{D}^{\mathrm{c}}$: that is, $\mathfrak{D}_{i,t}\left(y_i^t\right) = A_{i,t}\backslash\mathfrak{D}^{\mathrm{c}}_{i,t}\left(y_i^t\right)$ for each $i \neq 0$, t, and y_i^t. For a player $i \neq 0$, period t, and payoff-relevant history y_i^t, an action $a_{i,t} \in A_{i,t}$ is *codominated* if $a_{i,t} \in \mathfrak{D}_{i,t}\left(y_i^t\right)$.

The equivalence of this recursive definition and the fixed-point definition given in the text follows from the fact that, for any correspondence $\mathfrak{B}_t = (\mathfrak{B}_{i,t})_{i\neq 0}$ with $\mathfrak{B}_{i,t} : Y_i^t \rightrightarrows A_{i,t}$ that is *not* a codomination correspondence, there exists a correspondence $\mathfrak{B}'_t \subset \mathfrak{B}_t$ such that every action in $\mathfrak{B}_t\backslash\mathfrak{B}'_t$ is $(\mathcal{F}_t, \mathfrak{B}^{\mathrm{c}}_t)$-motivatable, where $\mathcal{F}_t$ is the set of mediation plans that never recommend codominated actions after period t.[35] Hence, if some action outside $\mathfrak{A}_t^{[l]}$ is not codominated—so $\left(\mathfrak{A}_t^{[l]}\right)^{\mathrm{c}}$ is not a codominated correspondence—then there exists a

[33] It may be helpful to note that $\mathcal{A}_{i,T}^{[1]}\left(y_i^T\right)$ is the set of actions that are played with positive probability by type y_i^T in the one-shot game where each player j's type space is Y_j^T, for some correlated equilibrium and some prior on $\prod_j Y_j^T$.

[34] For $r_i^{t'+1} \in R_i^{t'+1} \setminus Y_i^{t'}$, $f_{i,t'}(r_i^{t'+1})$ is not restricted.

[35] This fact is established by the first paragraph of the proof of Lemma 3 of Myerson (1986), setting $k = t$ and $B^1 = \mathfrak{B}_t \times \prod_{\tau\geq t+1} \mathfrak{D}_\tau$.

correspondence $\left(\mathfrak{A}_t^{[l+1]}\right)^{\mathrm{c}} \subset \left(\mathfrak{A}_t^{[l]}\right)^{\mathrm{c}}$ such that every action in $\left(\mathfrak{A}_t^{[l]}\right)^{\mathrm{c}} \setminus \left(\mathfrak{A}_t^{[l+1]}\right)^{\mathrm{c}} = \mathfrak{A}_t^{[l+1]} \setminus \mathfrak{A}_t^{[l]}$ is $\left(\mathcal{F}_t, \mathfrak{A}_t^{[l]}\right)$-motivatable.

B Proof of Proposition 2

Fix a game $G = (\Gamma, \mathfrak{C})$ and a NE (σ, ϕ). We construct a canonical NE $\left(\tilde{\sigma}, \tilde{\phi}\right)$ in $G^* = (\Gamma, \mathfrak{C}^*)$ with $\rho^{\tilde{\sigma},\tilde{\phi}} = \rho^{\sigma,\phi}$. We take $\tilde{\sigma} = \sigma^*$: players are honest and obedient at every history. The mediator's strategy $\tilde{\phi}$ 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 $\sigma_{i,1}^R(\tilde{s}_{i,1})$ (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 $\phi_1(m_1|r_1)$. 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 $\sigma_{i,1}^A(\tilde{m}_{i,1}|\tilde{s}_{i,1}, r_{i,1}, m_{i,1})$, 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 $\sigma_{i,t}^R(\tilde{s}_i^t, r_i^t, m_i^t, \tilde{a}_i^t, \tilde{s}_{i,t})$, independently across players.[36] Given the resulting vector $r_t = (r_{i,t})_i$, the mediator draws $m_t \in M_t$ according to $\phi_t(m_t|r^t, m^t, r_t)$. 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 $\sigma_{i,t}^A(\tilde{s}_i^t, r_i^t, m_i^t, \tilde{a}_i^t, \tilde{s}_{i,t}, r_{i,t}, m_{i,t})$, independently across players.[37] Finally, the mediator sends message $\tilde{m}_{i,t}$ to player i.

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

It remains to prove that $\left(\tilde{\sigma}, \tilde{\phi}\right)$ is a NE in G^*. We first show that, for any deviant strategy $\tilde{\sigma}_i'$ that player i can play against $\left(\tilde{\sigma}_{-i}, \tilde{\phi}\right)$ in game G^*, there exists a strategy σ_i' that yields the same outcome when played against (σ_{-i}, ϕ) in game G.

Lemma 1 *For each i and strategy $\tilde{\sigma}_i' \in \Sigma_i^*$, there exists a strategy $\sigma_i' \in \Sigma_i$ such that $\rho^{\sigma_i',\sigma_{-i},\phi} = \rho^{\tilde{\sigma}_i',\tilde{\sigma}_{-i},\tilde{\phi}}$.*

Proof. Fix i and $\tilde{\sigma}_i' \in \Sigma_i^*$. We construct $\sigma_i' \in \Sigma_i$ 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 $\tilde{\sigma}_{i,1}'^R(s_{i,1})$. Player i then sends report $r_{i,1} \in R_{i,1}$ according to $\sigma_{i,1}^R(\tilde{r}_{i,1})$. 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 $\sigma_{i,1}^A(\tilde{r}_{i,1}, r_{i,1}, m_{i,1})$. Finally, player i takes action $a_{i,1} \in A_{i,1}$ according to $\tilde{\sigma}_{i,1}'^A(s_{i,1}, \tilde{r}_{i,1}, \tilde{m}_{i,1})$.

[36] If $(\tilde{s}_i^t, \tilde{a}_i^t, \tilde{s}_{i,t}) \notin Y_i^t$, the mediator can draw $r_{i,t} \in R_{i,t}$ arbitrarily (e.g., uniformly at random).

[37] Again, if $(\tilde{s}_i^t, \tilde{a}_i^t, \tilde{s}_{i,t}) \notin Y_i^t$, the mediator can draw $\tilde{m}_{i,t}$ arbitrarily.

Recursively, for $t = 2, \ldots, T$, given her past signals and actions $(s_{i,\tau}, a_{i,\tau})_{\tau=1}^{t-1}$, her past fictitious type reports and action recommendations $(\tilde{r}_{i,\tau}, \tilde{m}_{i,\tau})_{\tau=1}^{t-1}$, her past reports and messages $(r_{i,\tau}, m_{i,\tau})_{\tau=1}^{t-1}$, and her current 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 $\tilde{\sigma}_{i,t}^{\prime R}(s_i^t, \tilde{r}_i^t, \tilde{m}_i^t, a_i^t, s_{i,t})$. Player i then sends $r_{i,t} \in R_{i,t}$ according to $\sigma_{i,t}^R(\tilde{r}_i^t, r_i^t, m_i^t, \tilde{r}_{i,t})$.[38] 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 $\sigma_{i,t}^A(\tilde{r}_i^t, r_i^t, m_i^t, \tilde{r}_{i,t}, r_{i,t}, m_{i,t})$. Finally, player i takes action $a_{i,t} \in A_{i,t}$ according to $\tilde{\sigma}_{i,t}^{\prime A}(s_i^t, \tilde{r}_i^t, \tilde{m}_i^t, a_i^t, s_{i,t}, \tilde{r}_{i,t}, \tilde{m}_{i,t})$.

Given this construction, $\rho^{\sigma_i', \sigma_{-i}, \phi} = \rho^{\tilde{\sigma}_i', \tilde{\sigma}_{-i}, \tilde{\phi}}$ by the same argument as for $\rho^{\tilde{\sigma}, \tilde{\phi}} = \rho^{\sigma, \phi}$. ■

Now suppose towards a contradiction that there exist $i \neq 0$ and $\tilde{\sigma}_i' \in \Sigma_i^*$ such that $\bar{u}_i\left(\tilde{\sigma}_i', \tilde{\sigma}_{-i}, \tilde{\phi}\right) > \bar{u}_i\left(\tilde{\sigma}_i, \tilde{\sigma}_{-i}, \tilde{\phi}\right)$. By Lemma 1, there exists $\sigma_i' \in \Sigma_i$ such that

$$\bar{u}_i\left(\sigma_i', \sigma_{-i}, \phi\right) = \bar{u}_i\left(\tilde{\sigma}_i', \tilde{\sigma}_{-i}, \tilde{\phi}\right) > \bar{u}_i\left(\tilde{\sigma}_i, \tilde{\sigma}_{-i}, \tilde{\phi}\right) = \bar{u}_i\left(\sigma_i, \sigma_{-i}, \phi\right).$$

This contradicts the hypothesis that (σ, ϕ) is a NE in G.

C Proof of Proposition 3

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. In the online appendix (Appendix E), we extend the example to prove that restricting to direct communication is with loss of generality.

Implementability in Non-Canonical SE Propositions 1 and 5 show that any outcome that is implementable in a canonical NE in which codominated actions are never played is SE-implementable. It thus suffices to construct a canonical NE that implements $\frac{1}{2}(A, A, N) + \frac{1}{2}(B, B, N)$ in which players avoid codominated actions. Such a NE is: the mediator recommends $m_1 = A$ 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 $s = 1$. Note that each $a_1 \in \{A, B\}$ and $a_2 = N$ are never codominated, and $a_2 = P$ is not codominated after $s = 1$ as P is optimal if $(a_1, \theta) = (C, p)$.

Non-Implementability in 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 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 of the mediator and player 2, and for 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), since neither the mediator nor player 2 observes (a_1, θ). 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, with probability 1 conditional on $\left(m_1, a_1, \theta, \hat{a}_1, \hat{\theta}, s\right)$, $a_2 = P$ iff

[38] If $\tilde{r}_i^{t+1} \notin Y_i^t$, player i can draw $r_{i,t}$ arbitrarily, and similarly for $\tilde{m}_{i,t}$ in what follows.

$m_2 = P$). 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

$$\Pr\left(m_2 = P | m_1, \hat{a}_1 = C, s = 1\right) > 0 \text{ for each } m_1 \in \{A, B\}.$$

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, after each $m_1 \in \{A, B\}$, 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,

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

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

$$\lim_{k\to\infty} \Pr^k\left(m_2 = P | m_1, a_1 = C, \theta = n, s = 1\right) = \lim_{k\to\infty} \Pr^k\left(m_2 = P | m_1, a_1 = C, \theta = p, s = 1\right) > 0. \tag{6}$$

Therefore,

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

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 (6) 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 playing A and B with equal probability after each m_1, so $\frac{1}{2}(A, A, N) + \frac{1}{2}(B, B, N)$ cannot be implemented.

D Main Results for Sequential Equilibrium

This section contains our analysis of SE, culminating in the proofs of Propositions 4 and 5.

D.1 Quasi-Strategies and Quasi-SE

We begin by introducing notions of "quasi-strategy," which is simply a partially defined strategy, and "quasi-equilibrium," which is a profile of quasi-strategies where incentive constraints are satisfied wherever strategies are defined. We use these concepts to show that defining strategies and assessing sequential rationality only after a subset of histories (which necessarily includes all on-path histories) suffices to establish the existence of a SE with the specified on-path behavior. The basic idea is that strategies outside the specified subset can be defined implicitly without affecting incentives at histories within the subset.

Fix a game $G = (\Gamma, \mathfrak{C})$.

Intuitively, a quasi-strategy for player i consists of a subset of histories $J_i \subset H_i$ and a strategy χ_i that is defined only on J_i. Formally, for each player i, a *quasi-strategy* (χ_i, J_i) consists of

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

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

The key requirement in the definition of a quasi-strategy (χ_i, J_i) is thus that for each history $h_i^{R,t} \in J_i^{R,t}$ there is some continuation path of play that terminates at a history $h_i^{T+1} \in J_i^{A,T+}$, and conversely any history $h_i^{R,t}$ reached along the path of play leading to any terminal history $h_i^{T+1} \in J_i^{A,T+}$ is contained in $J_i^{R,t}$ (and similarly for $h_i^{A,t} \in J_i^{A,t}$). We also let $J = \{h \in H : h_i \in J_i \ \forall i = 0, ..., N\}$. Note that $h^{R,t} \in J^{R,t}$ if and only if $h_i^{R,t} \in J_i^{R,t}$ for all i, and similarly for $\left(h^{R,t}, r_t\right) \in J^{R,t+}$, $h^{A,t} \in J^{A,t}$, and $\left(h^{A,t}, a_t\right) \in J^{A,t+}$.

Similarly, a *quasi-strategy* (ψ, K) for the mediator consists of

1. A set of histories $K = \bigcup_{t=1}^{T} (K^t \cup K^{t+})$ 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 each $(r^{T+1}, m^T) \in K^T$ 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+}, with (r^{t+1}, m^t) replaced by (r^{t+1}, m^{t+1}).

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

A strategy profile (σ, ϕ) has *support within* (J, K) if (i) for each $(r^{t+1}, m^t) \in K^t$, $\phi_t(m_t | r^{t+1}, m^t) > 0$ only if $(r^{t+1}, m^{t+1}) \in K^{t+}$, (ii) for each $h_i^{R,t} \in J_i^{R,t}$, $\sigma_{i,t}^R(r_{i,t} | h_i^{R,t}) > 0$ only if $(h_i^{R,t}, r_{i,t}) \in J_i^{R,t+}$, and (iii) for each $h_i^{A,t} \in J_i^{A,t}$, $\sigma_{i,t}^A(a_{i,t} | h_i^{A,t}) > 0$ only if $(h_i^{A,t}, a_{i,t}) \in J_i^{A,t+}$.

Recall that $\dot{h}^{R,t}$ denotes the payoff-*irrelevant* component of $h^{R,t}$. Define $h^{R,t} \in H^{R,t}|_{J,K}$ if $h_i^{R,t} \in J_i^{R,t}$ for each i and $\dot{h}^{R,t} \in K^{t-1,+}$, with the convention that $\dot{h}^{R,1} \in K^{0,+}$ vacuously holds. For each i and $h_i^{R,t} \in J_i^{R,t}$, define $h^{R,t} \in H^{R,t}[h_i^{R,t}]|_{J,K}$ if $h^{R,t} \in H^{R,t}|_{J,K}$ and $h^{R,t} \in H^{R,t}[h_i^{R,t}]$. Define $h^{A,t} \in H^{A,t}|_{J,K}$, and $h^{A,t} \in H^{A,t}[h_i^{A,t}]|_{J,K}$ analogously.

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

1. $J^{R,1} = S_1$. For each $t \geq 1$, $h^{R,t} \in H^{R,t}|_{J,K}$, $i \neq 0$, σ_i, $\tau \geq t$, and $h^{R,\tau}$ with $\Pr^{\sigma_i, \chi_{-i}, \psi}\left(h^{R,\tau} | h^{R,t}\right) > 0$, we have $h_j^{R,\tau} \in J_j^{R,\tau}$ for each $j \neq i$ and $\dot{h}^{R,\tau} \in K^{\tau-1,+}$.[39] Similarly, for each r_τ with $\Pr^{\sigma_i, \chi_{-i}, \psi}\left(h^{R,\tau}, r_\tau | h^{R,t}\right) > 0$, we have $\left(\dot{h}^{R,\tau}, r_\tau\right) \in K^\tau$; and for each m_τ with $\Pr^{\sigma_i, \chi_{-i}, \psi}\left(h^{R,\tau}, r_\tau, m_\tau | h^{R,t}\right) > 0$, we have $\left(h_j^{R,\tau}, r_{j,\tau}, m_{j,\tau}\right) \in J_j^{A,\tau}$ for each $j \neq i$. The same condition holds when we replace $h^{R,t} \in H^{R,t}|_{J,K}$ by $h^{A,t} \in H^{A,t}|_{J,K}$. That is, no unilateral player-deviation leads to a history where either the mediator's or another player's quasi-strategy is undefined.

2. For each (σ, ϕ) with support within (J, K), we have $\Pr^{\sigma,\phi}(h^{T+1} \in H^{T+1}|_{J,K}) = 1$.[40]

The first requirement implies that, for every valid quasi-strategy profile (χ, ψ, J, K), every history $h^{R,t}$ (respectively, $h^{A,t}$) with $\Pr^{\chi,\psi}\left(h^{R,t}\right) > 0$ ($\Pr^{\chi,\psi}\left(h^{A,t}\right) > 0$) lies in $H^{R,t}|_{J,K}$ ($H^{A,t}|_{J,K}$). That is, $H^{T+1}|_{J,K}$ includes all on-path histories under (χ, ψ). This implies in particular that (χ, ψ) induces a well-defined outcome $\rho^{\chi,\psi} \in \Delta(X)$. The second requirement implies that the same conclusion is true for each strategy profile with support within (J, K).

Finally, a *quasi-SE* $(\chi, \psi, J, K, \beta)$ is a valid quasi-strategy profile (χ, ψ, J, K) together with 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

$$\sum_{h^{R,t} \in H^{R,t}[h_i^{R,t}]|_{J,K}} \beta_i\left(h^{R,t} | h_i^{R,t}\right) \bar{u}_i\left(\chi, \psi | h^{R,t}\right) \geq \sum_{h^{R,t} \in H^{R,t}[h_i^{R,t}]|_{J,K}} \beta_i\left(h^{R,t} | h_i^{R,t}\right) \bar{u}_i\left(\sigma_i', \chi_{-i}, \psi | h^{R,t}\right). \tag{7}$$

[39] If there exists $j \neq i$ with $h_j^{A,\tau-1} \notin J_j^{A,\tau-1}$, or if $\dot{h}^{R,\tau} \notin K^{\tau-1}$, then $\Pr^{\sigma_i, \chi_{-i}, \psi}\left(h^{R,\tau} | h^{R,t}\right)$ is not well-defined. In this case, the above condition vacuously holds. The same caution applies to the following conditions.

[40] Note that, since (σ, ϕ) is a fully-specified strategy profile, $\Pr^{\sigma,\phi}$ is well-defined.

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

$$\sum_{h^{A,t} \in H^{A,t}[h_i^{A,t}]|_{J,K}} \beta_i\left(h^{A,t}|h_i^{A,t}\right) \bar{u}_i\left(\chi, \psi | h^{A,t}\right) \geq \sum_{h^{A,t} \in H^{A,t}[h_i^{A,t}]|_{J,K}} \beta_i\left(h^{A,t}|h_i^{A,t}\right) \bar{u}_i\left(\sigma_i', \chi_{-i}, \psi | h^{A,t}\right). \tag{8}$$

3. *[Kreps-Wilson consistency]* There exists a sequence of strategy profiles $\left(\sigma^k, \phi^k\right)_{k=1}^{\infty}$ such that

   (a) $\left(\sigma^k, \phi^k\right)$ has support within (J, K) for each k.

   (b) $\Pr^{\sigma^k, \phi^k}\left(h_i^{R,t}\right) > 0$ and $\Pr^{\sigma^k, \phi^k}\left(h_i^{A,t}\right) > 0$ for all i, $h_i^{R,t} \in J_i^{R,t}$, and $h_i^{A,t} \in J_i^{A,t}$.

   (c) $\lim_{k\to\infty} \sigma_{i,t}^{R,k}\left(h_i^{R,t}\right) = \chi_{i,t}\left(h_i^{R,t}\right)$ for each i and $h_i^{R,t} \in J_i^{R,t}$, $\lim_{k\to\infty} \sigma_{i,t}^{A,k}\left(h_i^{A,t}\right) = \chi_{i,t}\left(h_i^{A,t}\right)$ for each i and $h_i^{A,t} \in J_i^{A,t}$, and $\lim_{k\to\infty} \phi_t^k\left(r^{t+1}, m^t\right) = \psi_t\left(r^{t+1}, m^t\right)$ for each $\left(r^{t+1}, m^t\right) \in K^t$.

   (d)

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

   for each i, $h_i^{R,t} \in J_i^{R,t}$, $h_i^{A,t} \in J_i^{A,t}$, $h^{R,t} \in H^{R,t}[h_i^{R,t}]|_{J,K}$, and $h^{A,t} \in H^{A,t}[h_i^{A,t}]|_{J,K}$.

The following lemma shows that it is without loss to consider quasi-SE rather than fully specified SE.

Lemma 2 *For any game G and outcome $\rho \in \Delta(X)$, ρ is a SE outcome in G if and only if $\rho = \rho^{\chi,\psi}$ for some quasi-SE $(\chi, \psi, J, K, \beta)$ in G. Moreover, for any quasi-SE $(\chi, \psi, J, K, \beta)$, there exists a SE (σ, ϕ) such that (σ, ϕ) and (χ, ψ) coincide on (J, K).*

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

For the converse, fix a quasi-SE $(\chi, \psi, J, K, \beta)$ and a corresponding sequence of strategy profiles $\left(\tilde{\sigma}^k, \tilde{\phi}^k\right)_k$ satisfying the conditions of Kreps-Wilson consistency on (J, K). For each k, let

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

Let $\mathfrak{R}_{i,t}^k(h_i^{R,t}) = \operatorname{supp} \tilde{\sigma}_{i,t}^{R,k}(\cdot|h_i^{R,t})$, $\mathfrak{A}_{i,t}^k(h_i^{A,t}) = \operatorname{supp} \tilde{\sigma}_{i,t}^{A,k}(\cdot|h_i^{A,t})$, and $\mathfrak{M}_t^k(r^{t+1}, m^t) = \operatorname{supp} \tilde{\phi}_t^k(\cdot|r^{t+1}, m^t)$. Since $\left(\tilde{\sigma}^k, \tilde{\phi}^k\right)$ has support within (J, K), $\left(h_i^{R,t}, r_{i,t}\right) \in J_i^{R,t+}$ for each $h_i^{R,t} \in J_i^{R,t}$ and $r_{i,t} \in$

$\mathfrak{R}_{i,t}^k(h_i^{R,t})$, $\left(h_i^{A,t}, a_{i,t}\right) \in J_i^{A,t+}$ for each $h_i^{A,t} \in J_i^{A,t}$ and $a_{i,t} \in \mathfrak{A}_{i,t}^k(h_i^{A,t})$, and $(r^{t+1}, m^{t+1}) \in K^{t+}$ for each $(r^{t+1}, m^t) \in K^t$ and $m_t \in \mathfrak{M}_t^k(r^{t+1}, m^t)$.

We now define an auxiliary game $\left(\Gamma^k, \mathfrak{C}\right)$ indexed by k. In this game, each player i chooses a strategy $\sigma_i^k \in \Sigma_i$ and the mediator chooses a behavioral mediation plan ϕ^k, subject to the requirement that their choices coincide with $\left(\tilde{\sigma}^k, \tilde{\phi}^k\right)$ at histories in $H|_{J,K}$. These strategies are then perturbed so that every history h^t occurs with positive probability, but when k is large all histories outside $H|_{J,K}$ occur with much smaller probability than any history within $H|_{J,K}$. Formally, the game $\left(\Gamma^k, \mathfrak{C}\right)$ is defined as follows:

1. The mediator chooses probability distributions $\phi_t^k(\cdot|r^{t+1}, m^t) \in \Delta(M_t)$ 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 is required to choose $\phi_t^k(\cdot|r^{t+1}, m^t) = \tilde{\phi}_t^k(\cdot|r^{t+1}, m^t)$.

2. Each player i chooses probability distributions $\sigma_{i,t}^{R,k}(\cdot|h_i^{R,t}) \in \Delta(R_{i,t})$ and $\sigma_{i,t}^{A,k}(\cdot|h_i^{A,t}) \in \Delta(A_{i,t})$ for each t, $h_i^{R,t} \in H_i^{R,t} \setminus J_i^{R,t}$, and $h_i^{A,t} \in H_i^{A,t} \setminus J_i^{A,t}$. At histories $h_i^{R,t} \in J_i^{R,t}$ and $h_i^{A,t} \in J_i^{A,t}$, player i is required to choose $\sigma_{i,t}^{R,k}(\cdot|h_i^{R,t}) = \tilde{\sigma}_{i,t}^{R,k}\left(\cdot|h_i^{R,t}\right)$ and $\sigma_{i,t}^{A,k}(\cdot|h_i^{A,t}) = \tilde{\sigma}_{i,t}^{A,k}\left(\cdot|h_i^{A,t}\right)$.

3. Given $\left(\sigma^k, \phi^k\right)$, the distribution of terminal histories 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

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

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

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

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

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

   Given $h^{A,t} \in H_t^{A,t}$ and $a_t \in A_t$, 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. Player i's payoff at terminal history $h^{t+1} \in H^{T+1}$ is $u_i\left(\mathring{h}^{t+1}\right)$.

The interpretation of the distribution of $r_{i,t}$ in part 3 is as follows (the interpretation of the distributions of m_t and $a_{i,t}$ are similar): If the current history $h_i^{R,t}$ lies in $J_i^{R,t}$, then player i reports each $r_{i,t} \in \mathfrak{R}_{i,t}^k(h_i^{R,t})$ with probability $\tilde{\sigma}_{i,t}^{R,k}(r_{i,t}|h_i^{R,t})$ (in which case the resulting pair $\left(h_i^{R,t}, r_{i,t}\right)$ lies in $J_i^{R,t+}$), barring a low-probability tremble to a report outside $\mathfrak{R}_{i,t}^k(h_i^{R,t})$. Such low-probability trembles occur with uniform probability $\frac{\varepsilon_k}{k}$, which is much smaller than the probability of any report in the support of $\tilde{\sigma}_i^{R,k}$ when k is large. Finally, if the current history $h_i^{R,t}$ is already outside $J_i^{R,t}$, then player i follows her chosen strategy $\sigma_i^{R,k}$, while trembling uniformly with probability $\frac{\varepsilon_k}{k}$.

Note that the strategy set in the game $\left(\Gamma^k, \mathfrak{C}\right)$ 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 $\left(\Gamma^k, \mathfrak{C}\right)$. Moreover, since $\left(\sigma^k, \phi^k\right)$ has full support on Z for any strategy profile σ^k in $\left(\Gamma^k, \mathfrak{C}\right)$, Bayes' rule defines a belief system β^k by

$$\beta_i^k\left(h^t|h_i^t\right) = \frac{\Pr_{\Gamma^k,\mathfrak{C}}^{\sigma^k,\phi^k}(h^t)}{\Pr_{\Gamma^k,\mathfrak{C}}^{\sigma^k,\phi^k}(h_i^t)}$$

for all $i \neq 0$, all h_i^t, and all $h^t \in H^t[h_i^t]$, where $\Pr_{\Gamma^k,\mathfrak{C}}$ denotes probability in game $\left(\Gamma^k, \mathfrak{C}\right)$.

So let $(\bar{\sigma}^k, \bar{\phi}^k, \bar{\beta}^k)_k$ denote a sequence of NE $\left(\bar{\sigma}^k, \bar{\phi}^k\right)$ in $\left(\Gamma^k, \mathfrak{C}\right)$ with corresponding beliefs $\bar{\beta}^k$. Taking a subsequence if necessary to guarantee convergence, let $(\bar{\sigma}, \bar{\phi}, \bar{\beta}) = \lim_{k\to\infty}(\bar{\sigma}^k, \bar{\phi}^k, \bar{\beta}^k)$. Note that $\left(\bar{\sigma}, \bar{\phi}\right)$ and (χ, ψ) coincide on $H|_{J,K}$. We claim that $(\bar{\sigma}, \bar{\phi}, \bar{\beta})$ is a SE in $(\Gamma, \mathfrak{C})$. Since $\bar{\beta}$ satisfies Kreps-Wilson consistency by construction, 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 $h_i^{T+1} \notin J_i^{A,T+}$ for all h_i^{T+1} that follow $h_i^{R,t}$, so by inspection the outcome distribution (and hence player i's expected payoff) conditional on $h_i^{R,t}$ is continuous in σ^k, ϕ^k, ε_k, and k. Since $\bar{\sigma}_{i,t}^{R,k}\left(\cdot|h_i^{R,t}\right)$ is sequentially rational in $\left(\Gamma^k, \mathfrak{C}\right)$ (as $\left(\bar{\sigma}^k, \bar{\phi}^k\right)$ is a NE in $\left(\Gamma^k, \mathfrak{C}\right)$, where the distribution over h^{T+1} has full support), it follows that $\bar{\sigma}_{i,t}^{R}\left(\cdot|h_i^{R,t}\right)$ is sequentially rational in $(\Gamma, \mathfrak{C})$.

Now consider the case where $h_i^{R,t} \in J_i^{R,t}$. We show that player i believes that $h^{R,t} \in H^{R,t}[h_i^{R,t}]|_{J,K}$ with probability 1. Note that, for each $h_i^{T+1} \in J_i^{A,T+}$ and h_{-i}^{T+1} with $\left(h_i^{T+1}, h_{-i}^{T+1}\right) \notin H^{T+1}[h_i^{T+1}]|_{J,K}$, there exists $h_{-i}'^{T+1}$ such that $(h_i^{T+1}, h_{-i}'^{T+1}) \in H^{T+1}[h_i^{T+1}]|_{J,K}$ and

$$\lim_{k\to\infty} \frac{\Pr^{\sigma^k,\phi^k}(h_i^{T+1}, h_{-i}^{T+1})}{\Pr^{\sigma^k,\phi^k}(h_i^{T+1}, h_{-i}'^{T+1})} = 0.$$

This follows because in $\left(\Gamma^k, \mathfrak{C}\right)$ each "tremble" leading to a history outside J occurs with probability at most ε_k/k (this is an implication of Condition 3(a) of Kreps-Wilson consistency for quasi-SE and the third condition in the definition of a valid quasi strategy profile),

while every history $h_i^{T+1} \in J_i^{A,T+}$ occurs with positive probability given $\left(\sigma^k, \phi^k\right)$ (this is an implication of Condition 3(b) of Kreps-Wilson consistency for quasi-SE).

Therefore, for each $h_i^{R,t} \in J_i^{R,t}$, we have $\bar{\beta}_i\left(h^{R,t}|h_i^{R,t}\right) = \beta_i(h^{R,t}|h_i^{R,t})$ for all $h^{R,t} \in H^{R,t}[h_i^{R,t}]|_{J,K}$. Moreover, by the second condition in the definition of a valid quasi strategy profile, for any σ_i', player i believes that players $-i$ follow χ_{-i}. Hence, the fact that (7) holds for belief β_i implies that $\sigma_{i,t}^R(\cdot|h_i^{R,t}) = \chi_{i,t}\left(\cdot|h_i^{R,t}\right)$ is sequentially rational in $(\Gamma, \mathfrak{C})$. ■

In the proofs of Propositions 4 and 5, it will be convenient to describe the mediator's strategy as 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.

D.2 Proof of Proposition 5

Here we prove that every SCE outcome is SE-implementable with pseudo-direct communication, and hence $SCE \subset SE$. As discussed in Section 4, the reverse inclusion follows from Propositions 6, 7, and 8.

By Proposition 1, it suffices to show that every outcome that arises in a canonical NE in which codominated actions are never recommended at any history is SE-implementable with pseudo-direct communication.

Under pseudo-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 obeyed all action recommendations. Note that faithfulness places no restriction on player i's action in periods τ in which she received message $\star$. Faithfulness at histories $h_i^{R,t}$ and $h_i^{A,t}$ are similarly defined.

Trembling-Hand Perfect Equilibrium As previewed in Section 3.4, our construction begins by defining an arbitrary trembling-hand perfect equilibrium (PE).

Fix $(\varepsilon_k)_{k\in\mathbb{N}}$ satisfying $\varepsilon_k \to 0$ and $k\left(\varepsilon_k\right)^{NT} \to \infty$. For each k, let $\hat{\sigma}^k$ be a NE in the unmediated, ε_k-constrained game where each player is required to play each action with probability at least ε_k at each information set. Taking a subsequence if necessary, $\left(\hat{\sigma}^k\right)_{k\in\mathbb{N}}$ converges to a PE $\hat{\sigma}$ in the unconstrained game $(\Gamma, \emptyset)$. Thus, for each i, t, y_i^t, and strategy σ_i', we have

$$\sum_{y^t \in Y^t[y_i^t]} \hat{\beta}_{i,t}\left(y^t|y_i^t\right) \bar{u}_i\left(\hat{\sigma}|y^t\right) \geq \sum_{y^t \in Y^t[y_i^t]} \hat{\beta}_{i,t}\left(y^t|y_i^t\right) \bar{u}_i\left(\sigma_i', \hat{\sigma}_{-i}|y^t\right), \tag{9}$$

where $Y^t[y_i^t]$ is the set of $y^t \in Y^t$ with i component equal to y_i^t and

$$\hat{\beta}_{i,t}\left(y^t|y_i^t\right) = \lim_{k\to\infty} \frac{\Pr^{\hat{\sigma}^k}\left(y^t\right)}{\Pr^{\hat{\sigma}^k}\left(y_i^t\right)}.$$

For future reference, for each y_i^t, let $\hat{B}_{i,t}\left(y_i^t\right) = A_{i,t} \backslash \operatorname{supp} \hat{\sigma}_{i,t}\left(y_i^t\right)$ denote the set of actions that

are taken at y_i^t only when player i trembles. For $r_i^{t+1} \in R_i^{*t+1} \setminus Y_i^t$, we define $\hat{B}_{i,t}\left(r_i^{t+1}\right) = \emptyset$.

Consider now the mediated, unrestricted, direct-communication game $(\Gamma, \mathfrak{C}^*)$. Suppose the mediator performs all randomizations in the PE $\hat{\sigma}$ on behalf of the players, so that the outcome of $\hat{\sigma}$ results if players are honest and obedient: that is, the mediator follows the behavioral mediation plan $\tilde{\phi}$ constructed from $\hat{\sigma}$ as in the proof of Proposition 2. Note that, since $\hat{\sigma}$ is a strategy profile in the unmediated game, $\tilde{\phi}_{i,t}\left(r^{t+1}, m^t\right)$ depends only on r_i^{t+1}, for all i, t, and (r^{t+1}, m^t). In particular, the correspondence defined by $\tilde{Q}_{i,t}\left(r_i^{t+1}\right) = \operatorname{supp} \tilde{\phi}_{i,t}\left(r_i^{t+1}\right)$ for all i and t is a mediation range.[41] Moreover, since player i's recommendations and player j's recommendations are drawn independently, we can write $\tilde{\phi}_t(m_t | r^{t+1}, m^t) = \prod_i \tilde{\phi}_{i,t}\left(m_{i,t} | r_i^{t+1}\right)$ for all t, m_t, and (r^{t+1}, m^t).

Now let $\left(\tilde{\sigma}^k, \tilde{\phi}\right)$ denote the profile in the mediated game $(\Gamma, \mathfrak{C}^*)$ where players are honest and obedient, while trembling uniformly over actions with probability ε_k. Note that $\tilde{\sigma}^k$ converges to the fully canonical strategy σ^*. Define $\chi = \sigma^*$, $\psi = \tilde{\phi}$, $J_i^{R,t} = \{(s_i^{t+1}, r_i^t, m_i^t, a_i^t) : r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau}) \wedge m_{i,\tau} \in \tilde{Q}_{i,t}\left(r_i^{\tau+1}\right) \ \forall \tau \leq t-1\}$ for all i and t (and similarly for $J_i^{R,t+}$, $J_i^{A,t}$, and $J_i^{A,t+}$), $K^t = \{(r^t, m^t) : m_{i,\tau} \in \tilde{Q}_{i,\tau}\left(r_i^{\tau+1}\right) \ \forall i, \tau \leq t \ \}$ for all t (and similarly for K^+), and

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

for all i, t, $h_i^{R,t}$, and $h^{R,t} \in H^{R,t}[h_i^{R,t}]|_{J,K}$ (and similarly for $\tilde{\beta}_{i,t}\left(h^{A,t} | h_i^{A,t}\right)$). For $y^t \in Y^t[\mathring{h}_i^{R,t}]$, we write

$$\tilde{\beta}_{i,t}\left(y^t | h_i^{R,t}\right) = \sum_{h^{R,t} \in H^{R,t}[h_i^{R,t}]|_{J,K} \text{ with } \mathring{h}^{R,t} = y^t} \tilde{\beta}_{i,t}\left(h^{R,t} | h_i^{R,t}\right),$$

and we write $\bar{u}_i\left(\sigma_i, \sigma_{-i}^*, \tilde{\phi} | \dot{h}_i^{R,t}, y^t\right)$ for player i's continuation payoff at the history $h^{R,t}$ where $h_i^{R,t} = \left(y_i^t, \dot{h}_i^{R,t}\right)$ (recall that $\dot{h}_i^{R,t}$ is the payoff-irrelevant components of $h_i^{R,t}$) and $h_j^{R,t} = \bar{h}_j(y_j^t)$ for each $j \neq i$.[42]

Lemma 3 *$\left(\chi, \psi, J, K, \tilde{\beta}\right)$ is a quasi-SE. Moreover, for each i, t, σ_i', $h_i^{R,t} \in J_i^{R,t}$, and $h_i'^{R,t} \in J_i^{R,t}$ with $\mathring{h}_i^{R,t} = \mathring{h}_i'^{R,t}$, we have*

$$\sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \tilde{\beta}_{i,t}\left(y^t | h_i^{R,t}\right) \bar{u}_i\left(\sigma^*, \tilde{\phi} | \dot{h}_i'^{R,t}, y^t\right) \geq \sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \tilde{\beta}_{i,t}\left(y^t | h_i^{R,t}\right) \bar{u}_i\left(\sigma_i', \sigma_{-i}^*, \tilde{\phi} | \dot{h}_i'^{R,t}, y^t\right). \tag{10}$$

Proof. For each i, let J_i be the set of histories where players are honest and all past messages

[41] Since we construct $\tilde{\phi}$ in Proposition 2 such that, after $r_i^{t+1} \in R_i^{*t+1} \setminus Y^t$, the mediator sends all action recommendations with positive probability. Hence, $\tilde{Q}_{i,t}\left(r_i^{t+1}\right) = A_{i,t}$ for $r_i^{t+1} \in R_i^{*t+1} \setminus Y^t$.

[42] Here, $\bar{h}_j(y_j^t)$ is the history for player j that obtains under honesty and obedience given payoff-relevant history y_j^t.

lie in the mediation range: for each t,

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

Similarly, for each t, let

$$\begin{aligned}
K^t &= \left\{\left(r^{t+1}, m^t\right) \in R^{t+1} \times M^t : m_{i,\tau} \in Q_{i,\tau}\left(r_i^\tau, m_i^\tau, r_{i,\tau}\right) \quad \forall i, \tau < t\right\}, \\
K^{t+} &= \left\{\left(r^{t+1}, m^{t+1}\right) \in R^{t+1} \times M^{t+1} : m_{i,\tau} \in Q_{i,\tau}\left(r_i^\tau, m_i^\tau, r_{i,\tau}\right) \quad \forall i, \tau \leq t\right\}.
\end{aligned}$$

The quasi-strategy profile (χ, ψ, J, K) satisfies the two defining conditions for validity in G, since (i) $J^{R,1} = S^1$ by definition, (ii) histories outside J_i cannot arise as long as player i is honest and the mediator follows ψ, and (iii) for each i, a terminal history h_i^{T+1} arises with positive probability when players are honest and take all actions with positive probability and the mediator sends all messages within the mediation range with positive probability if and only if all reports in h_i^{T+1} are honest and all messages in h_i^{T+1} lie in the mediation range.

To establish that $\left(\chi, \psi, J, K, \tilde{\beta}\right)$ is a quasi-SE, it remains to show sequential rationality:

$$\sum_{h^{R,t} \in H^{R,t}[h_i^{R,t}]|_{J,K}} \tilde{\beta}_{i,t}\left(h^{R,t}|h_i^{R,t}\right) \bar{u}_i\left(\sigma^*, \tilde{\phi}|h^{R,t}\right) \geq \sum_{h^{R,t} \in H^{R,t}[h_i^{R,t}]|_{J,K}} \tilde{\beta}_{i,t}\left(h^{R,t}|h_i^{R,t}\right) \bar{u}_i\left(\sigma_i', \sigma_{-i}^*, \tilde{\phi}|h^{R,t}\right) \tag{11}$$

for each i, t, $h_i^{R,t} \in J_i^{R,t}$, and strategy σ_i' in game $(\Gamma, \mathfrak{C}^*)$ (and similarly for $h_i^{A,t}$). Since (i) $h_i^{R,t} \in J_i^{R,t}$ implies that player i has been honest, (ii) $\tilde{\phi}_t(r^{t+1}, m^t) = \prod_i \tilde{\phi}_{i,t}\left(r_i^{t+1}\right)$, and (iii) player i always believes that her opponents are honest, for each i, t, $h_i^{R,t} \in J_i^{R,t}$, and σ_i', we can write (11) as

$$\sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \tilde{\beta}_{i,t}\left(y^t|h_i^{R,t}\right) \bar{u}_i\left(\sigma^*, \tilde{\phi}|\dot{h}_i^{R,t}, y^t\right) \geq \sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \tilde{\beta}_{i,t}\left(y^t|h_i^{R,t}\right) \bar{u}_i\left(\sigma_i', \sigma_{-i}^*, \tilde{\phi}|\dot{h}_i^{R,t}, y^t\right). \tag{12}$$

Moreover, for each $y^t \in Y^t[\mathring{h}_i^{R,t}]$, we have

$$\begin{aligned}
\tilde{\beta}_{i,t}\left(y^t|h_i^{R,t}\right) &= \lim_{k\to\infty} \frac{\Pr^{\tilde{\sigma}^k,\tilde{\phi}}(y^t) \Pr^{\tilde{\sigma}^k,\tilde{\phi}}(m_i^t|y^t)}{\Pr^{\tilde{\sigma}^k,\tilde{\phi}}(y_i^t) \Pr^{\tilde{\sigma}^k,\tilde{\phi}}(m_i^t|y_i^t)} = \lim_{k\to\infty} \frac{\Pr^{\tilde{\sigma}^k,\tilde{\phi}}(y^t) \Pr^{\tilde{\sigma}^k,\tilde{\phi}}(m_i^t|y_i^t)}{\Pr^{\tilde{\sigma}^k,\tilde{\phi}}(y_i^t) \Pr^{\tilde{\sigma}^k,\tilde{\phi}}(m_i^t|y_i^t)} \\
&= \lim_{k\to\infty} \frac{\Pr^{\tilde{\sigma}^k,\tilde{\phi}}(y^t)}{\Pr^{\tilde{\sigma}^k,\tilde{\phi}}(y_i^t)} = \lim_{k\to\infty} \frac{\Pr^{\hat{\sigma}^k}(y^t)}{\Pr^{\hat{\sigma}^k}(y_i^t)} = \hat{\beta}_{i,t}\left(y^t|y_i^t\right),
\end{aligned}$$

where the second equality uses $\Pr^{\tilde{\sigma}^k,\tilde{\phi}}(m_i^t|y^t) = \Pr^{\tilde{\sigma}^k,\tilde{\phi}}(m_i^t|y_i^t)$, which holds since $\tilde{\phi}_t(m_t|r^{t+1}, m^t) =$

$\prod_i \tilde{\phi}_{i,t}\left(m_{i,t}|r_i^{t+1}\right)$. Hence, (12) is equivalent to

$$\sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \hat{\beta}_{i,t}\left(y^t|\mathring{h}_i^{R,t}\right) \bar{u}_i\left(\sigma^*, \tilde{\phi}|\dot{h}_i^{R,t}, y^t\right) \geq \sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \hat{\beta}_{i,t}\left(y^t|\mathring{h}_i^{R,t}\right) \bar{u}_i\left(\sigma_i', \sigma_{-i}^*, \tilde{\phi}|\dot{h}_i^{R,t}, y^t\right).$$

By the same argument as in the proof of Proposition 2, $\left(\sigma^*, \tilde{\phi}\right)$ induces the same outcome distribution in $(\Gamma, \mathfrak{C}^*)$ as $\hat{\sigma}$ does in $(\Gamma, \emptyset)$. Hence,

$$\sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \hat{\beta}_{i,t}\left(y^t|\mathring{h}_i^{R,t}\right) \bar{u}_i\left(\sigma^*, \tilde{\phi}|\dot{h}_i^{R,t}, y^t\right) = \sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \hat{\beta}_{i,t}\left(y^t|\mathring{h}_i^{R,t}\right) \bar{u}_i\left(\hat{\sigma}|y^t\right).$$

In addition, by the same construction as in the proof of Lemma 1, for each $h_i^{R,t} \in J_i^{R,t}$ and every strategy σ_i' in game $(\Gamma, \mathfrak{C}^*)$, there exists a strategy $\hat{\sigma}_i'$ in game $(\Gamma, \emptyset)$ such that

$$\sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \hat{\beta}_{i,t}\left(y^t|\mathring{h}_i^{R,t}\right) \bar{u}_i\left(\sigma_i', \sigma_{-i}^*, \tilde{\phi}|\dot{h}_i^{R,t}, y^t\right) = \sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \hat{\beta}_{i,t}\left(y^t|\mathring{h}_i^{R,t}\right) \bar{u}_i\left(\hat{\sigma}_i', \hat{\sigma}_{-i}|y^t\right).$$

Hence, (9) implies (12).

To complete the proof, we show that (12) can be strengthened to (10). To see this, note that, under strategy $\sigma_i = \sigma_i^*$, since $\tilde{\phi}$ implements the play of a PE in the unmediated game, player i's continuation play in periods $\tau \geq t$ does not depend on m_i^t. Hence, for each $h_i^{R,t} \in J_i^{R,t}$ and $h_i'^{R,t} \in J_i^{R,t}$ with $\mathring{h}_i^{R,t} = \mathring{h}_i'^{R,t}$, we have

$$\sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \tilde{\beta}_{i,t}\left(y^t|h_i^{R,t}\right) \bar{u}_i\left(\sigma^*, \tilde{\phi}|\dot{h}_i^{R,t}, y^t\right) = \sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \tilde{\beta}_{i,t}\left(y^t|h_i^{R,t}\right) \bar{u}_i\left(\sigma^*, \tilde{\phi}|\dot{h}_i'^{R,t}, y^t\right).$$

Moreover, again because $\tilde{\phi}$ pertains to the unmediated game, players' recommendations are independent. Hence, for any σ_i' that depends on m_i^t, there is another strategy σ_i'' that does not depend on m_i^t and achieves the same payoff. Hence, for each $h_i^{R,t} \in J_i^{R,t}$ and $h_i'^{R,t} \in J_i^{R,t}$ with $\mathring{h}_i^{R,t} = \mathring{h}_i'^{R,t}$, we have

$$\max_{\sigma_i'} \sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \tilde{\beta}_{i,t}\left(y^t|h_i^{R,t}\right) \bar{u}_i\left(\sigma_i', \sigma_{-i}^*, \tilde{\phi}|\dot{h}_i^{R,t}, y^t\right) = \max_{\sigma_i'} \sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \tilde{\beta}_{i,t}\left(y^t|h_i^{R,t}\right) \bar{u}_i\left(\sigma_i', \sigma_{-i}^*, \tilde{\phi}|\dot{h}_i'^{R,t}, y^t\right).$$

Hence, (12) implies (10). ■

By Kuhn's theorem, there exists a mixed mediation plan $\tilde{\mu} \in \Delta(F)$ such that $(\sigma, \tilde{\mu})$ and $\left(\sigma, \tilde{\phi}\right)$ induce the same distribution on terminal histories Z in $(\Gamma, \mathfrak{C}^*)$ for all strategies σ. Since $\tilde{\phi}_t(m_t|r^{t+1}, m^t) = \prod_i \tilde{\phi}_{i,t}\left(m_{i,t}|r_i^{t+1}\right)$ for all i, t, m_t, and (r^{t+1}, m^t), we have

$$\tilde{\mu}(f) = \textstyle\prod_{i,t} \tilde{\mu}_{i,t}(f_{i,t}) \text{ for all } f. \tag{13}$$

For each t, let F_t denote the set of functions $f_t : R^{t+1} \to M_t$, and define $F^{<t} = \times_{\tau<t} F_\tau$ and $F^{\geq t} = \times_{\tau \geq t} F_\tau$. Thus, $\left(f^{<t}, f^{\geq t}\right) \in F$ for each t. For $f^{<t} \in F^{<t}$, we write $f^{<t} \in \operatorname{supp} \tilde{\mu}^{<t}$ if

there exists $f^{\geq t} \in F^{\geq t}$ such that $\left(f^{<t}, f^{\geq t}\right) \in \operatorname{supp} \tilde{\mu}$. Define $f_i \in \operatorname{supp} \tilde{\mu}_i$ and $f_i^{<t} \in \operatorname{supp} \tilde{\mu}_i^{<t}$ analogously, focusing on the mediation plan for player i.

Rationalizing Non-Codominated Actions We now construct a sequence of mixed mediation plans that induce all non-codominated actions with positive probability.

Let

$$F^* = \left\{f \in F : f_{i,t}(r^{t+1}) \in A_{i,t} \setminus \mathfrak{D}_{i,t}(r_i^{t+1}) \ \forall i, t, r^{t+1}\right\}$$

be the set of mediation plans that never recommend codominated actions, and let

$$F^{*\geq t} = \left\{f^{\geq t} \in F^{\geq t} : f_{i,\tau}(r^{\tau+1}) \in A_{i,\tau} \setminus \mathfrak{D}_{i,t}(r_i^{\tau+1}) \ \forall i, \tau \geq t, r^{\tau+1}\right\}$$

be the projection of F^* on $F^{\geq t}$.[43] As in Myerson's Lemma 3, the definition of codomination requires that there exist $L \geq 1$ and distributions $\left((\pi_t^{[l]})_{l=1}^L\right)_{t=1}^T$ with $\pi_t^{[l]} \in \Delta\left(F^{*\geq t} \times X^t \times S_t\right)$ for each t and l satisfying the following conditions:

First, define $\pi_t^k = \sum_{l=1}^L \left(\frac{1}{k}\right)^{l-1} \pi_t^{[l]}$. Then, for each i, t, and y_i^t, denote the support of $f_{i,t}^{\geq t}$ at y_i^t under π_t^k by

$$\operatorname{supp}_{i,t}\left(y_i^t\right) = \left\{m_{i,t} \in A_{i,t} : \begin{array}{c}\text{there exist } f^{\geq t} \text{ and } y^t \text{ with } i\text{-component equal to } y_i^t \\ \text{such that } \pi_t^k\left(f^{\geq t}, y^t\right) > 0 \wedge f_{i,t}^{\geq t}(y^t) = m_{i,t}\end{array}\right\}.$$

Note that this set is the same for all $k \in \mathbb{N}$. Finally, for each t, let

$$\Pr^{\sigma^*,\pi_t^k}(f^{\geq t}, \mathring{h}^{T+1}, m^{t:T}) = \pi_t^k\left(f^{\geq t}, \mathring{h}^{R,t}\right) \Pr^{\sigma^*, f^{\geq t}}\left(\mathring{h}^{T+1}, m^{t:T} | \bar{h}\left(\mathring{h}^{R,t}\right)\right),$$

where $m^{t:T+1} = (m_t, \ldots, m_T)$. Note that $\operatorname{supp} \Pr^{\sigma^*,\pi_t^k}$ is the same for all $k \in \mathbb{N}$.

With these definitions, the required conditions on $\left((\pi_t^{[l]})_{l=1}^L\right)_{t=1}^T$ are

1. Every non-codominated action is recommended with positive probability:

$$\operatorname{supp}_{i,t}(y_i^t) = A_{i,t} \setminus \mathfrak{D}_{i,t}(y_i^t).$$

2. At every history reached with positive probability under profile $\left(\sigma^*, \pi_t^k\right)$, honesty and obedience is optimal under the beliefs β^{σ^*,π_t} derived from $\left(\sigma^*, \pi_t^k\right)$ as $k \to \infty$: for each i, t, $\tau \geq t$, $(\mathring{h}_i^{R,\tau}, m_i^{t:\tau})$ satisfying $\Pr^{\sigma^*,\pi_t^k}(\mathring{h}_i^{R,\tau}, m_i^{t:\tau}) > 0$ (for any $k \in \mathbb{N}$), and strategy σ_i', we have

$$\begin{aligned} &\sum_{f^{\geq t} \in F^{*\geq t}, y^\tau \in Y^\tau[\mathring{h}_i^{R,\tau}]} \beta_{i,\tau}^{\sigma^*,\pi_t}\left(f^{\geq t}, y^\tau | \mathring{h}_i^{R,\tau}, m_i^{t:\tau}\right) \bar{u}_i\left(\sigma^* | f^{\geq t}, y^\tau\right) \\ \geq &\sum_{f^{\geq t} \in F^{*\geq t}, y^\tau \in Y^\tau[\mathring{h}_i^{R,\tau}]} \beta_{i,\tau}^{\sigma^*,\pi_t}\left(f^{\geq t}, y^\tau | \mathring{h}_i^{R,\tau}, m_i^{t:\tau}\right) \bar{u}_i\left(\sigma_i', \sigma_{-i}^* | f^{\geq t}, y^\tau\right), \end{aligned} \tag{14}$$

[43] Note that in the definition of $F^{*\geq t}$, there is no restriction for $f_{i,\tau}(r^{\tau+1})$ if $r^{\tau+1} \in R_i^{*\tau+1} \setminus Y_i^\tau$.

where

$$\beta_{i,\tau}^{\sigma^*,\pi_t}\left(f^{\geq t},y^\tau|\mathring{h}_i^{R,\tau},m_i^{t:\tau}\right)=\lim_{k\to\infty}\frac{\Pr^{\sigma^*,\pi_t^k}(f^{\geq t},y^\tau,m_i^{t:\tau})}{\Pr^{\sigma^*,\pi_t^k}(y_i^\tau,m_i^{t:\tau})}.$$

Since every reachable history $h^{R,\tau}$ under $\left(\sigma^*,\pi_t^k\right)$ is faithful and $f^{\geq t}$ sends only action recommendations (that is, $\star$ is not recommended), $\left(f^{\geq t},y^\tau\right)$ uniquely pins down $h^{R,\tau}=\bar{h}\left(y^\tau\right)$ given that we have $h^{R,t}=\bar{h}\left(y^t\right)$. Hence, we write $\bar{u}_i\left(\sigma_i',\sigma_{-i}^*|f^{\geq t},y^\tau\right)$ for $\bar{u}_i\left(\sigma_i',\sigma_{-i}^*|f^{\geq t},\bar{h}\left(y^\tau\right)\right)$.

"Motivating Equilibrium" Construction We next define a sequence of quasi-strategy profiles $\left(\sigma^k,\phi^k,J,K\right)_k$ in the game with pseudo-direct communication $(\Gamma,\mathfrak{C}^{**})$, where the quasi-SE profile for the desired "motivating equilibrium" is given by $(\sigma,\phi,J,K)=\lim_{k\to\infty}\left(\sigma^k,\phi^k,J,K\right)$.

Players' strategies σ^k: Each player i is faithful: with probability 1, she plays $a_{i,t}=m_{i,t}$ after each $m_{i,t}\in A_{i,t}$ and reports $r_{i,t}=(a_{i,t-1},s_{i,t})$ after each $(a_{i,t-1},s_{i,t})\in A_{i,t-1}\times S_{i,t}$. After receiving message $m_{i,t}=\star$, with probability $1-\sqrt{\varepsilon_k}$ player i takes $a_{i,t}$ according to her PE strategy $\hat{\sigma}_{i,t}(y_i^t)$, 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 $\Pr\left(\zeta_{i,t}=0\right)=1-\left(\frac{1}{k}\right)^{2(L+1)T}$. Third, independently for each i, he draws $\tilde{f}_i$ from $\tilde{\mu}_i$. 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*

$$\omega_t\in\bigcup\nolimits_{f\in\operatorname{supp}\tilde{\mu}}(0,f)\cup\left(\bigcup\nolimits_{1\leq t^*\leq T,f\in F^{*,\geq t^*}}(t^*,f)\right),$$

with initial state $\omega_0=\left(0,\tilde{f}\right)$. Let $\omega=\omega^{T+1}$. Given $\theta=\theta^{T+1}$ and $\zeta=\zeta^{T+1}$, 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:** For each i, t, and y_i^{t-1}, denote the number of tuples $\left(f_i^{<t},\theta_i^t,m_i^t\right)$ with $f_i^{<t}\in\operatorname{supp}\tilde{\mu}_i^{<t}$ such that, for all $\tau\leq t-1$, the mediator sends $m_{i,\tau}=f_{i,\tau}^{<t}(y_i^\tau)$ if $\theta_{i,\tau}=0$ and sends $m_{i,\tau}=\star$ if $\theta_{i,\tau}=1$ by

$$\#M_i(y_i^{t-1})=\left|\left\{\begin{array}{c}\left(f_i^{<t},\theta_i^t,m_i^t\right)\in\operatorname{supp}\tilde{\mu}_i^{<t}\times\{0,1\}^{t-1}\times\prod_{\tau=1}^{t-1}\left(A_{i,\tau}\cup\{\star\}\right)\\ :\left\{\begin{array}{l}m_{i,\tau}=f_{i,\tau}^{<t}(y_i^\tau)\ \forall\tau\leq t-1\text{ s.t. }\theta_{i,\tau}=0,\\ m_{i,\tau}=\star\ \forall\tau\leq t-1\text{ s.t. }\theta_{i,\tau}=1\end{array}\right.\end{array}\right\}\right|. \quad (15)$$

Let $\#M(y^{t-1})=\prod_{i=0}^N\#M_i(y_i^{t-1})$. In addition, for $\tilde{f}^{<t}\in\operatorname{supp}\tilde{\mu}^{<t}$, denote the number of recommendation strategies $\tilde{f}'\in\operatorname{supp}\tilde{\mu}$ which coincide with $\tilde{f}^{<t}$ for the first $t-1$ periods by $\#(\tilde{f}^{<t})=|\{\tilde{f}'\in\operatorname{supp}\tilde{\mu}:\tilde{f}'^{<t}=\tilde{f}^{<t}\}|$.

- **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 $\omega_t=\omega_{t-1}$ with probability 1. If $\omega_{t-1}=\omega_0$ then the mediator calculates the probability of $(\theta,\zeta,\omega_{t-1}=\omega_0,r^{t+1},m^t)$ given σ^k and the construction of ϕ^k up to period t. Denote this probability by $p^k\left(\theta,\zeta,\omega_0,r^{t+1},m^t\right)$.

If $p^k(\theta,\zeta,\omega_0,r^{t+1},m^t)=0$ then $\omega_t=\omega_{t-1}$ with probability 1. If $p^k(\theta,\zeta,\omega_0,r^{t+1},m^t)>0$ then, for each $f^{\geq t}\in F^{*\geq t}$, the mediator draws $\omega_t=\left(t,f^{\geq t}\right)$ with probability

$$q^k\left(\omega_t|\theta,\zeta,\omega_0,r^{t+1},m^t\right)=\left(\frac{1}{k}\right)^{(L+1)t+2(L+1)T|\zeta|}\times\frac{1}{p^k(\theta,\zeta,\omega_0,r^{t+1},m^t)}\times\frac{\pi_t^k(f^{\geq t},r^{t+1})}{\#(\tilde{f}^{<t})\#M(r^t)},\tag{16}$$

and draws $\omega_t=\omega_0$ with the remaining probability. Note that $p^k\left(\theta,\zeta,\omega_0,r^{t+1},m^t\right)>0$ implies that r^{t+1} corresponds to some $y^t\in Y^t$, so $\pi_t^k(f^{\geq t},r^{t+1})$ and $\#M(r^t)$ are well-defined.

- **Calculation of m_t:** If $\omega_t=\left(0,\tilde{f}\right)$, then the mediator recommends $m_{i,t}=\tilde{f}_{i,t}(r_i^{t+1})$ 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 $A_{i,t}\setminus\mathfrak{D}_{i,t}(r_i^{t+1})$ with equal probability if $\zeta_{i,t}=1$. If $\omega_t=\left(t,f^{\geq t}\right)$ with $t\geq 1$, then the mediator recommends $m_{i,t}=f_{i,t}^{\geq t}(r^{t+1})$.

Definition of K and J: Let $K^{T+}=\left\{\left(r^{T+1},m^{T+1}\right):m_t\in\prod_i A_{i,t}\cup\{\star\}\setminus\mathfrak{D}_{i,t}(r_i^{t+1})\ \forall t\right\}$; for each t, K^t and K^{t+} consist of all truncations of histories in K^{T+}. Let $J_i^{A,T+}$ be the set of player i's histories h_i^{T+1} such that (i) $m_{i,t}\in A_{i,t}\cup\{\star\}\setminus\mathfrak{D}_{i,t}(r_i^{t+1})\ \forall t$, (ii) $r_{i,t}=(a_{i,t-1},s_{i,t})$ $\forall t$, and (iii) $a_{i,t}=m_{i,t}\ \forall t$ with $m_{i,t}\in A_{i,t}\setminus\mathfrak{D}_{i,t}(r_i^{t+1})$; the other elements of J_i consist of all truncations of histories in $J_i^{A,T+}$.

Let us give some interpretation of the mediator's strategy. The mediator's state ω_t indicates whether the mediator currently intends to implement the PE $\hat{\sigma}$ (in which case $\omega_t=\left(0,\tilde{f}\right)$) or has switched to implementing some other mediation plan $f^{\geq\tau}$ for some $\tau\leq t$ (in which case $\omega_t=\left(\tau,f^{\geq\tau}\right)$, where τ is the period when the mediator switched). The mediator's state switches at most once in the course of the game: that is, every state ω_t except $\left(0,\tilde{f}\right)$ is absorbing. Moreover, the probability that the mediator's state ever switches converges to 0 as $k\to\infty$. A crucial feature of the construction is that the mediator's state transition probability, q^k, is determined so that, conditional on the event that the mediator's state switches in period t, the likelihood ratio between any two mediation plans and payoff relevant histories $\left(f^{\geq t},\mathring{h}^t\right)$ and $\left(f'^{\geq t},\mathring{h}'^t\right)$ is the same as the likelihood specified by π_t^k. This feature will guarantee that the recommendation to play any non-codominated action is incentive compatible in the limit as $k\to\infty$.

In addition to possibly "trembling" from the initial state $\left(0,\tilde{f}\right)$ to another state $\left(t,f^{\geq t}\right)$, the mediator can also tremble in his recommendations while remaining in state $\left(0,\tilde{f}\right)$. Specifically, when $\omega_t=\left(0,\tilde{f}\right)$, $\theta_{i,t}=1$ indicates a mediator tremble that sends message $\star$ to player i. Player i then plays her PE strategy $\hat{\sigma}_i$ in period t but trembles with probability $\sqrt{\varepsilon_k}$. Since $\Pr\left(\theta_{i,t}=1\right)=\sqrt{\varepsilon_k}$ from the perspective of each of i's opponents, they assess that player i trembles with probability $\sqrt{\varepsilon_k}\times\sqrt{\varepsilon_k}=\varepsilon_k$, exactly as in strategy profile $\hat{\sigma}^k$. This argument is formalized in Lemma 5.

Also, when $\omega_t=\left(0,\tilde{f}\right)$, $\zeta_{i,t}=1$ indicates a mediator tremble that recommends all non-codominated actions with positive probabilities. This event is very rare, so that when player

i is recommended a non-codominated action outside $\operatorname{supp}\hat{\sigma}_{i,t}\left(y_i^t\right)$ in period t, she believes that $\zeta_{i,t}=0$ and this surprising recommendation is instead due to a switch in the mediator's state in period t. However, if she later reaches a history inconsistent with this explanation, she updates her belief to $\zeta_{i,t}=1$. This is formalized in Lemma 4.

We note that the quasi-strategy profile (σ,ϕ,J,K) is valid. To see this, observe that a profile (σ',ϕ') has full support on (J,K) if and only if each player i is faithful and the support of the mediator's recommendation equals $\prod_i A_{i,t}\cup\{\star\}\setminus\mathfrak{D}_{i,t}(r_i^{t+1})$ for each t and (r^{t+1},m^t). Given this observation, the three conditions of the definition of validity are immediate.

Joint Distribution of Histories and Mediator States Given the quasi-strategy profile $\left(\sigma^k,\phi^k\right)$ just defined, we calculate the joint distribution of $\left(\theta,\zeta,\omega,h^{T+1}\right)$, which we denote by δ^k.

For each $(\theta,\zeta,\omega_0,r^{t+1},m^t)$ such that $p^k(\theta,\zeta,\omega_0,r^{t+1},m^t)>0$, we have

$$p^k(\theta,\zeta,\omega_0,r^{t+1},m^t)\geq\tilde{\mu}\left(\tilde{f}\right)\times\frac{\left(\varepsilon_k\right)^{NT}}{\left|A^T\right|}\times\left(\frac{1}{k}\right)^{2(L+1)T|\zeta|}.$$

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

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

Since $k\left(\varepsilon_k\right)^{NT}\to\infty$ as $k\to\infty$, this implies

$$\lim_{k\to\infty}q^k\left(\omega_t|\theta,\zeta,\omega_0,r^{t+1},m^t\right)=0. \tag{17}$$

Given y^t, $f^{<t}\in\operatorname{supp}\tilde{\mu}^{<t}$, θ^t, and ζ^t, let $M^t(f^{<t},\theta^t,\zeta^t,y^t)$ denote the set of m^t such that, for each i and $\tau=1,...,t-1$, (i) $m_{i,\tau}=f_{i,\tau}^{<t}(y_i^\tau)$ 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) $m_{i,\tau}\in A_{i,\tau}\setminus\mathfrak{D}_{i,\tau}(y_i^\tau)$ if $\zeta_{i,\tau}=1$.

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

and $f(\omega) = \left(\tilde{f}^{<t}, \hat{f}^{\geq t}\right)$. We have[44]

$$\begin{aligned}
&\delta^k(\theta, \zeta, \omega, h^{T+1}) \\
= \;& 1_{\{\omega_t = \omega_0 \forall t\}} \times \Pr^{\phi^k}(\omega_t = \omega_0 \forall t) \times \Pr^{\phi^k}(\theta, \zeta) \times \tilde{\mu}(f(\omega)) \times \Pr^{\sigma^k}\left(h^{T+1} | f(\omega), \theta, \zeta\right) \\
& + \sum_{t=1}^{T} 1_{\{t^*(\omega)=t\}} \sum_{\tilde{f}: \tilde{f}^{<t} = f^{<t}(\omega)} \left(\begin{array}{l} p^k(\theta, \zeta, \omega_0, \mathring{h}^{R,t}, m^t) \left(\frac{1}{k}\right)^{(L+1)t + 2(L+1)T|\zeta|} \\ \times \frac{1}{p^k(\theta, \zeta, \omega_0, \mathring{h}^{R,t}, m^t)} \frac{\pi_t^k(f^{\geq t}(\omega), \mathring{h}^{R,t})}{\#(f^{<t}(\omega)) \# M(\mathring{h}^{R,t})} \\ \times \Pr^{\sigma^*}\left(\mathring{h}^{T+1} | \mathring{h}^{R,t}, f^{\geq t}(\omega)\right) \\ \times 1_{\left\{m^t \in M^t(f^{<t}(\omega), \theta^t, \zeta^t, \mathring{h}^t) \text{ and } m_\tau = f^{\geq t}(\omega)(\mathring{h}^{R,\tau}) \; \forall \tau \geq t\right\}} \end{array} \right).
\end{aligned}$$

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

$$\begin{aligned}
&\delta^k(\theta, \zeta, \omega, h^{T+1}) \\
= \;& 1_{\{\omega_t = \omega_0 \forall t\}} \times \Pr^{\phi^k}(\omega_t = \omega_0 \forall t) \times \Pr^{\phi^k}(\theta, \zeta) \times \tilde{\mu}(f(\omega)) \times \Pr^{\sigma^k}\left(h^{T+1} | f(\omega), \theta, \zeta\right) \\
& + \sum_{t=1}^{T} 1_{\{t^*(\omega)=t\}} \left(\begin{array}{l} \left(\frac{1}{k}\right)^{(L+1)t + 2(L+1)T|\zeta|} \frac{\pi_t^k(f^{\geq t}(\omega), \mathring{h}^{R,t})}{\#M(\mathring{h}^{R,t})} \Pr^{\sigma^*}\left(\mathring{h}^{T+1} | \mathring{h}^{R,t}, f^{\geq t}(\omega)\right) \\ \times 1_{\left\{m^t \in M^t(f^{<t}(\omega), \theta^t, \zeta^t, \mathring{h}^t) \text{ and } m_\tau = f^{\geq t}(\omega)(\mathring{h}^{R,\tau}) \; \forall \tau \geq t\right\}} \end{array} \right).
\end{aligned}$$

Conditions for Quasi-Sequential Equilibrium The main remaining step in the proof is showing that (σ, ϕ, J, K), together with a belief system β, is a quasi-SE. (The proof will then be completed by mixing this quasi-SE with the target SCE outcome.) We first show that, for each k, i and t, $\delta^k\left(h_i^{R,t}\right) > 0$ for all $h_i^{R,t} \in J_i^{R,t}$ and $\delta^k\left(h_i^{A,t}\right) > 0$ for all $h_i^{A,t} \in J_i^{A,t}$. We then compute β as the limit of conditional probabilities under δ^k. We will then be ready to verify sequential rationality given beliefs β.

We show that $\delta^k\left(h_i^{A,T+}\right) > 0$ for each k, i, and $h_i^{A,T+} = \left(s_i^{T+1}, r_i^{T+1}, m_i^{T+1}, a_i^{T+1}\right) \in J_i^{A,T+}$. Fix any $\left(s_{-i}^{T+1}, a_{-i}^{T+1}\right)$ such that $\left(s_i^{T+1}, s_{-i}^{T+1}, a_i^{T+1}, a_{-i}^{T+1}\right) \in X^{T+1}$. For each t, define $\theta_{j,t} = 1$ for each j and t; for each t, define $m_{j,t} = \star$ and $\zeta_{j,t} = 0$ for each $j \neq i$ and t; and for each t, define $\zeta_{i,t} = 1$ if and only if $m_{i,t} \neq \star$. It suffices to show that

$$\left(s_i^{T+1}, s_{-i}^{T+1}, r^{T+1}, m^{T+1}, a_i^{T+1}, a_{-i}^{T+1}\right)$$

happens with a positive probability given δ^k if $r_t = (a_{t-1}, s_t)$ for each t and $a_{i,t} = m_{i,t}$ for each t with $m_{i,t} \in A_{i,t}$. Since any $m_{i,t} \in A_{i,t} \setminus \mathfrak{D}_{i,t}(r_i^{t+1})$ is recommended with a positive probability given $\zeta_{i,t} = 1$, $m_{j,t} = \star$ is recommended given $\zeta_{j,t} = 0$ and $\theta_{j,t} = 1$ for each j and t, and each player j takes all actions with positive probability given $m_{j,t} = \star$, we have $\delta^k\left(h_i^{A,T+}\right) > 0$.

[44] Here $\mathring{h}^t$ is the projection of h^{T+1} on X^t. Since players are faithful, $\mathring{h}^t$ fully determines the reports, and hence r^{T+1} does not appear in this calculation. We also write $\Pr^{\sigma^*}\left(\mathring{h}^{T+1} | \mathring{h}^{R,t}, f^{\geq t}(\omega)\right)$ instead of $\Pr^{\sigma^k}\left(\mathring{h}^{T+1} | \mathring{h}^{R,t}, f^{\geq t}(\omega)\right)$ since (i) if $\omega_t = \left(\tau, \hat{f}^{\geq \tau}\right)$ for some $\tau \neq 0$, then $m_t \in A_t$ and (ii) players follow all non-$\star$ recommendations given σ^k.

We now define the belief system β by

$$\beta_{i,t}\left(h^{R,t}|h_i^{R,t}\right)=\sum_{\theta,\zeta,\omega_t}\beta_{i,t}\left(\theta,\zeta,\omega_t,h^{R,t}|h_i^{R,t}\right),$$

where

$$\beta_{i,t}\left(\theta,\zeta,\omega_t,h^{R,t}|h_i^{R,t}\right)=\lim_{k\to\infty}\frac{\delta^k(\theta,\zeta,\omega,h^{R,t})}{\delta^k(h_i^{R,t})}$$

for i-component of $h^{R,t}$ being equal to $h_i^{R,t}$. By construction, this belief system satisfies Kreps-Wilson consistency given quasi-strategy profile (σ,ϕ,J,K). Thus, to establish that (σ,ϕ,J,K) is a quasi-SE, it remains to verify

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

$$\sum_{h^{R,t}\in H^{R,t}[h_i^{R,t}]|_{J,K}}\beta_{i,t}\left(h^{R,t}|h_i^{R,t}\right)\bar{u}_i\left(\sigma,\phi|h^{R,t}\right)\geq\sum_{h^{R,t}\in H^{R,t}[h_i^{R,t}]|_{J,K}}\beta_{i,t}\left(h^{R,t}|h_i^{R,t}\right)\bar{u}_i\left(\sigma_i',\sigma_{-i},\phi|h^{R,t}\right). \tag{18}$$

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

$$\sum_{h^{A,t}\in H^{A,t}[h_i^{A,t}]|_{J,K}}\beta_{i,t}\left(h^{A,t}|h_i^{A,t}\right)\bar{u}_i\left(\sigma,\phi|h^{A,t}\right)\geq\sum_{h^{A,t}\in H^{A,t}[h_i^{A,t}]|_{J,K}}\beta_{i,t}\left(h^{A,t}|h_i^{A,t}\right)\bar{u}_i\left(\sigma_i',\sigma_{-i},\phi|h^{A,t}\right). \tag{19}$$

Mediator States that Explain a Faithful History We now present Lemma 4, which was previewed above following the definition of $\left(\sigma^k,\phi^k,J,K\right)$. We first define the notion of a mediator state "explaining" a given faithful history.

Given a faithful history $h_i^{R,t}$ for some i and t, we say $(0,\zeta)$ *explains* $h_i^{R,t}$ if there exist $\tilde{f}\in\operatorname{supp}\tilde{\mu}$, θ, and $h_{-i}^{R,t}$ such that, for each j and $\tau=1,...,t-1$, (i) $m_{j,\tau}=a_{j,\tau}=\tilde{f}_{j,\tau}(\mathring{h}_j^{R,\tau})$ if $\zeta_{j,\tau}=\theta_{j,\tau}=0$, (ii) $m_{j,\tau}=\star$ if $\zeta_{j,\tau}=0$ and $\theta_{j,\tau}=1$, (iii) $m_{j,\tau}=a_{j,\tau}\in A_{j,\tau}\setminus\mathfrak{D}_{j,\tau}(\mathring{h}_j^{R,\tau})$ if $\zeta_{j,\tau}=1$, and (iv) $p\left(s_{\tau+1}|s^{\tau+1},a^{\tau+1}\right)>0$ (and also $p\left(s_1\right)>0$).

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

Similarly, given a faithful history $h_i^{A,t}$, we say $(0,\zeta)$ *explains* $h_i^{A,t}$ if (i)–(iii) hold for $\tau=1,...,t$ and (iv) holds for $\tau=0,...,t-1$; and (t^*,ζ) *explains* $h_i^{A,t}$ if $m_t=f_t^{\geq t^*}(\mathring{h}^{A,t})$ holds in addition to the above conditions (i)–(vi).

Let

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

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 $t^*<\tilde{t}^*$. Lemma 4 will establish that $(t^*,\zeta)<(\tilde{t}^*,\tilde{\zeta})$ if and only if a mediator trembles to $\pi_{t^*}^k$ with

$\zeta_{j,\tau} = 1$ for $|\zeta|$ values of (j,τ) is infinitely more likely than a mediator tremble to $\pi^k_{\tilde{t}^*}$ 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 $\xi(h_i^{A,t})$ denote the smallest pairs (t^*,ζ) that explain $h_i^{R,t}$ and $h_i^{A,t}$, respectively. Since Ξ is a finite set and $\delta^k\left(h_i^{R,t}\right)$ and $\delta^k\left(h_i^{A,t}\right)$ are strictly positive for all faithful histories $h_i^{R,t}$ and $h_i^{A,t}$, such pairs always exist. Note in addition that each realization (ζ,ω) defines a realization $(t^*,\zeta) \in \Xi$. Denote the corresponding random variable that takes realizations in Ξ by $(\mathbf{t}^*,\boldsymbol{\zeta})$.

For any $\zeta_i^t \in \{0,1\}^{t-1}$, define $\left(0,\zeta_i^t\right) = \left(0,\tilde{\zeta}^t\right)$ where $\tilde{\zeta}_i^t = \zeta_i^t$, $\tilde{\zeta}_{i,\tau} = 0$ for all $\tau \geq t$, and $\tilde{\zeta}_{j,\tau} = 0$ for all $j \neq i$ and all τ. Define $\left(t^*,\zeta_i^{t^*}\right)$ similarly.

Lemma 4 *For all faithful histories $h_i^{R,t}$ and $h_i^{A,t}$, the following claims hold:*

1. *We have*

$$\lim_{k\to\infty} \delta^k\left(\left\{(\mathbf{t}^*,\boldsymbol{\zeta}) = \xi(h_i^{R,t})\right\}|h_i^{R,t}\right) = 1, \tag{20}$$

$$\lim_{k\to\infty} \delta^k\left(\left\{(\mathbf{t}^*,\boldsymbol{\zeta}) = \xi(h_i^{A,t})\right\}|h_i^{A,t}\right) = 1. \tag{21}$$

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 $0 < t^* \leq t$ 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 $0 < t^* \leq t$ and $\zeta_i^{t^*}$.*

Claim 1 of the lemma says that $\xi(h_i^{R,t})$ is an infinitely more likely explanation for faithful history $h_i^{R,t}$ than any other (t^*,ζ), and $\xi(h_i^{A,t})$ is an infinitely more likely explanation for faithful history $h_i^{A,t}$ than any other (t^*,ζ).

Claim 2 says that the most likely explanation (t^*,ζ) for a faithful history has the following three properties. First, the most likely explanation never involves a player $j \neq i$ receiving a recommendation outside the support of $\hat{\sigma}_j$ while the mediator intends to implement $\hat{\sigma}$: that is, $\zeta_{j,\tau} = 0$ for all $j \neq i$ and all τ. Second, the most likely explanation never involves player i receiving a future recommendation outside the support of $\hat{\sigma}_i$ while the mediator intends to implement $\hat{\sigma}$: that is, $\zeta_{i,\tau} = 0$ for all $\tau > t$ (for $h_i^{R,t}$, we also have $\zeta_{i,t} = 0$; note that player i has not received her period-t recommendation at history $h_i^{R,t}$). Third, for an acting history $h_i^{A,t}$, the most likely explanation never involves a recommendation outside the support of $\hat{\sigma}_i$ while the mediator intends to implement $\hat{\sigma}$ in the current period: that is, for each $h_i^{A,t}$ with $m_{i,t} \in A_{i,t} \cap \hat{B}_{i,t}(r_i^{t+1})$, we have $\xi(h_i^{A,t}) = \left(t^*,\zeta_i^{t^*}\right)$ with $0 < t^* \leq t$.

Proof. We prove first Claim 2 and then Claim 1.

Claim 2: We first observe that, whenever $\xi(h_i^{R,t}) = (0,\zeta)$, we have $\zeta = \zeta_i^t$. To see this, note that whenever $(0,\zeta)$ explains $h_i^{R,t}$, so does $(0,\zeta')$ with $\zeta'_{i,\tau} = \zeta_{i,\tau}$ for all τ and $\zeta'_{j,\tau} = 0$ for all j and τ, since, for each τ and $a_{j,\tau}$, we can take $\zeta_{j,\tau} = 0$, $\theta_{j,\tau} = 1$, and $m_{j,\tau} = \star$, rather than $\zeta_{j,\tau} = 1$ and $m_{j,\tau} = a_{j,\tau}$. Moreover, whenever $(0,\zeta')$ explains $h_i^{R,t}$, so does $\left(0,\zeta_i^t\right)$. As $\left(0,\zeta_i^t\right) \leq (0,\zeta)$ with strict inequality if $\zeta_i^t \neq \zeta$, this implies the observation.

We next observe that, whenever $\xi(h_i^{R,t}) = (t^*,\zeta)$ for $t^* > 0$, we have $\zeta = \zeta_i^{t^*}$. To see this, note that whenever (t^*,ζ) explains $h_i^{R,t}$, so does (t^*,ζ') with $\zeta'_{j,\tau} = \zeta_{j,\tau}$ for all j and

$\tau < t^*$ and $\zeta'_{j,\tau} = 0$ for all j and $\tau \geq t^*$, since $m_{j,\tau}$ is independent of $\zeta_{j,\tau}$ for all j and $\tau \geq t^*$. Moreover, whenever (t^*, ζ') explains $h_i^{R,t}$, so does $\left(t^*, \zeta_i^{t^*}\right)$, since for each $\tau < t^*$ and $a_{j,\tau}$, 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 observation follows as $\left(t^*, \zeta_i^{t^*}\right) \leq (t^*, \zeta)$ with strict inequality if $\zeta_i^{t^*} \neq \zeta$. Given these two observations, Claim 2 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 $\zeta_i^{t+1} = \left(\zeta_i^t, 1\right)$. To see why this new condition holds, note that whenever $\left(0, \zeta_i^{t+1}\right)$ explains $h_i^{A,t}$, so does some $\left(t^*, \tilde{\zeta}_i^{t^*}\right)$ with $t^* = t$ and $\tilde{\zeta}_{i,\tau} = \zeta_{i,\tau}$ for each $\tau \leq t-1$. This is because, given $t^* = t$, for each r_i^{t+1}, each $m_{i,t} \in A_{i,t} \setminus \mathfrak{D}_{i,t}(r_i^{t+1})$, and each $\tilde{f}^{<t} \in \operatorname{supp} \tilde{\mu}^{<t}$, we have $m_{i,t} \in \operatorname{supp}_{i,t}(r_i^{t+1})$. Since $\left(t, \tilde{\zeta}_i^t\right) < \left(0, \zeta_i^{t+1}\right)$, we have $\xi(h_i^{A,t}) \neq \left(0, \zeta_i^{t+1}\right)$ with $\zeta_{i,t} = 1$.

Claim 1: We prove (20); the proof of (21) is analogous. Let $(t^*, \zeta) = \xi(h_i^{R,t})$.

Denote the smallest probability of any message vector m_i^{T+1} among those in the support of $\tilde{\phi}_i$ by

$$\underline{\varepsilon}_1 = \min_{i, m_i^{T+1}, \mathring{h}_i^{T+1} : \tilde{\phi}_{i,t}\left(m_{i,t} | \mathring{h}_i^t\right) > 0 \ \forall t} \prod_{t=1}^{T} \tilde{\phi}_{i,t}\left(m_{i,t} | \mathring{h}_i^t\right).$$

Denote the smallest probability of any signal vector s^{T+1} among those in the support of p by

$$\underline{\varepsilon}_2 = \min_{a^{T+1}, s^{T+1} : p(s_t | s^t, a^t) > 0 \ \forall t} \prod_{t=1}^{T} p\left(s_t | s^t, a^t\right).$$

We claim that

$$\begin{aligned} \delta^k\left(\left\{(\mathbf{t}^*, \boldsymbol{\zeta}) = (0, \zeta)\right\}, h_i^{R,t}\right) \ \geq \ & \left(\prod_{t=1}^{T}\left(1 - \left(\frac{1}{k}\right)^{(L+1)t}\right)\right) \\ & \times \left(\frac{1}{k}\right)^{2(L+1)T|\zeta|} \left(1 - \left(\frac{1}{k}\right)^{2(L+1)T}\right)^{NT - |\zeta|} \\ & \times \left(\sqrt{\varepsilon_k}\right)^{N(t-1)} \times \frac{\left(\sqrt{\varepsilon_k}\right)^{N(t-1)}}{|A^t|} \times \underline{\varepsilon}_1 \times \underline{\varepsilon}_2. \end{aligned} \tag{22}$$

The explanation is as follows: Define $\tilde{\theta}_i^t \in \{0,1\}^{t-1}$ by $\tilde{\theta}_{i,\tau} = 0$ if $m_{i,\tau} \neq \star$ and $\tilde{\theta}_{i,\tau} = 1$ if $m_{i,\tau} = \star$. First, $\prod_{t=1}^{T}\left(1 - \left(\frac{1}{k}\right)^{(L+1)t}\right)$ is a lower bound for the probability that $\omega_\tau = \omega_0$ for all $\tau \leq T$. Second,

$$\left(\frac{1}{k}\right)^{2(L+1)T|\zeta|} \left(1 - \left(\frac{1}{k}\right)^{2(L+1)T}\right)^{NT - |\zeta|}$$

is the probability that $\boldsymbol{\zeta} = \zeta$ (independently of ω). Third, conditional on any ω and ζ,

$$\left(1 - \sqrt{\varepsilon_k}\right)^{|\tau \leq t-1 : m_{i,\tau} \neq \star|} \left(\sqrt{\varepsilon_k}\right)^{|\tau \leq t-1 : m_{i,\tau} = \star|} \left(\sqrt{\varepsilon_k}\right)^{(N-1)(t-1)}$$

is a lower bound for the probability that $\theta_i^t = \tilde{\theta}_i^t$ and $\theta_{j,\tau} = 1$ for all $j \neq i$ and $\tau \leq t-1$; moreover, for sufficiently large k this product is no less than $\left(\sqrt{\varepsilon_k}\right)^{N(t-1)}$. Fourth, conditional on any θ^t such that $\theta_i^t = \tilde{\theta}_i^t$ and $\theta_{j,\tau} = 1$ for all $j \neq i$ and $\tau \leq t-1$, for any a^t, $\frac{\left(\sqrt{\varepsilon_k}\right)^{N(t-1)}}{|A^t|}$ is a lower bound for the probability that player j takes $a_{j,\tau}$ for all j (including $j = i$) and $\tau \leq t-1$ with $\theta_{j,\tau} = 1$. Fifth, conditional on $\omega_\tau = \omega_0$ for all $\tau \leq T$, any ζ, and any θ satisfying $\theta_i^t = \tilde{\theta}_i^t$, $\underline{\varepsilon}_1$ is a lower bound for the probability of $(m_{i,\tau})_{\tau \in \{1,\ldots,t-1\}: m_{i,\tau} \neq \star}$. Finally, conditional on any a^t, $\underline{\varepsilon}_2$ is a lower bound for the probability of s_i^{t+1}.

Similarly, for $t^* \geq 1$, denote the smallest probability of any tuple $\left(f^{\geq t^*}, y^{t^*}, x\right)$ in the support of $\pi_{t^*}^{[l]}$ for any t^* and l by

$$\min_{\substack{t^* \geq 1, l, f^{\geq t^*} \in F^{\geq t^*}, y^{t^*} \in Y^{t^*}, x \in X \\ : \pi_{t^*}^{[l]}\left(f^{\geq t^*}, y^{t^*}\right) \Pr^{\sigma^*}\left(x | f^{\geq t^*}, y^{t^*}\right) > 0}} \pi_{t^*}^{[l]}\left(f^{\geq t^*}, y^{t^*}\right) \Pr^{\sigma^*}\left(x | f^{\geq t^*}, y^{t^*}\right) \geq \underline{\varepsilon}_3.$$

We claim that

$$\begin{aligned}
\delta^k\left(\{(\mathbf{t}^*, \boldsymbol{\zeta}) = (t^*, \zeta)\}, h_i^{R,t}\right) \quad \geq \quad & \left(\prod_{t=1}^{t^*-1}\left(1 - \left(\frac{1}{k}\right)^{(L+1)t}\right)\right) \\
& \times \left(\frac{1}{k}\right)^{2(L+1)T|\zeta|}\left(1 - \left(\frac{1}{k}\right)^{2(L+1)T}\right)^{NT-|\zeta|} \\
& \times \left(\sqrt{\varepsilon_k}\right)^{N(t^*-1)} \times \frac{\left(\sqrt{\varepsilon_k}\right)^{N(t^*-1)}}{|A^{t^*}|} \times \underline{\varepsilon}_1 \times \underline{\varepsilon}_2 \\
& \times \left(\frac{1}{k}\right)^{(L+1)t^*} \frac{1}{\max_{t, \mathring{h}^{R,t} \in Y^t} \#M(\mathring{h}^{R,t})} \times \left(\frac{1}{k}\right)^{L} \times \underline{\varepsilon}_3.
\end{aligned}$$

The first three lines represent the same probability as (22), up to period $t^* - 1$. For the fourth line, (i) conditional on $\mathbf{t}^* \geq t^*$, $\left(\frac{1}{k}\right)^{(L+1)t^*} \frac{1}{\max_{t, \mathring{h}^{R,t} \in Y^t} \#M(\mathring{h}^{R,t})}$ is a lower bound for the probability that $\mathbf{t}^* = t^*$, (ii) conditional on $\mathbf{t}^* = t^*$, $\left(\frac{1}{k}\right)^L$ is a lower bound for the probability selecting index l for $\pi_{t^*}^{[l]}$, for each $l \in \{1, \ldots, L\}$, and (iii) conditional on $\mathbf{t}^* = t^*$ and l, $\underline{\varepsilon}_3$ is a lower bound for the probability of $\left(f^{\geq t^*}, x\right)$.

In contrast, for each $(\tilde{t}^*, \tilde{\zeta}) \neq \xi(h_i^{R,t})$, if $(\tilde{t}^*, \tilde{\zeta})$ does not explain $h_i^{R,t}$ then $\delta^k\left(\{(\mathbf{t}^*, \boldsymbol{\zeta}) = (\tilde{t}^*, \tilde{\zeta})\}, h_i^{R,t}\right) = 0$. If $(\tilde{t}^*, \tilde{\zeta})$ does explain $h_i^{R,t}$, then

$$\delta^k\left(\{(\mathbf{t}^*, \boldsymbol{\zeta}) = (\tilde{t}^*, \tilde{\zeta})\}, h_i^{R,t}\right) \leq \left(\frac{1}{k}\right)^{2(L+1)T\left|\tilde{\zeta}\right| + (L+1)\tilde{t}^*},$$

since $\left(\frac{1}{k}\right)^{2(L+1)T|\zeta|}$ is an upper bound for the probability that $\boldsymbol{\zeta} = \zeta$ and $\left(\frac{1}{k}\right)^{(L+1)\tilde{t}^*}$ is an upper bound for the probability that $\mathbf{t}^* = \tilde{t}^*$.

Ignoring terms that converge to 1 as $k \to \infty$, we have

$$\lim_{k\to\infty} \frac{\delta^k\left(\{(\mathbf{t}^*,\boldsymbol{\zeta})=(\tilde{t}^*,\tilde{\zeta})\},h_i^{R,t}\right)}{\delta^k\left(\{(\mathbf{t}^*,\boldsymbol{\zeta})=(t^*,\zeta)\},h_i^{R,t}\right)} \leq \lim_{k\to\infty} \frac{\left(\frac{1}{k}\right)^{2(L+1)T\left|\tilde{\zeta}\right|+(L+1)\tilde{t}^*}}{\left(\frac{1}{k}\right)^{2(L+1)T|\zeta|+(L+1)t^*+L}\max_{t,\mathring{h}^{R,t}\in Y^t}\#M(\mathring{h}^{R,t})\frac{(\varepsilon_k)^{NT}}{|A^T|}\underline{\varepsilon}_1\underline{\varepsilon}_2\underline{\varepsilon}_3}. \tag{23}$$

Since either $|\zeta|=\left|\tilde{\zeta}\right|$ and $t^*<\tilde{t}^*$ or $|\zeta|<\left|\tilde{\zeta}\right|$, we have

$$\begin{aligned}
&\left(2\left(L+1\right)T\left|\tilde{\zeta}\right|+\left(L+1\right)\tilde{t}^*\right)-\left(2\left(L+1\right)T\left|\zeta\right|+\left(L+1\right)t^*+L\right)\\
\geq&\begin{cases}2\left(L+1\right)T-\left(L+1\right)T-L & \text{if } |\zeta|<\left|\tilde{\zeta}\right|\\ L+1-L & \text{if } |\zeta|=\left|\tilde{\zeta}\right| \text{ and } t^*<\tilde{t}^*\end{cases}\\
\geq&\ 1.
\end{aligned}$$

Hence, the right-hand side of (23) is no more than

$$\frac{1}{k\frac{(\varepsilon_k)^{NT}}{|A^T|}\underline{\varepsilon}_1\underline{\varepsilon}_2\underline{\varepsilon}_3\max_{t,\mathring{h}^{R,t}\in Y^t}\#M(\mathring{h}^{R,t})}.$$

Since $k\left(\varepsilon_k\right)^{NT}\to\infty$ by assumption and $\underline{\varepsilon}_1$, $\underline{\varepsilon}_2$, $\underline{\varepsilon}_3$, and $\max_{t,\mathring{h}^{R,t}\in Y^t}\#M(\mathring{h}^{R,t})$ are constants independent of k, this converges to 0 as $k\to\infty$. Finally, since there are only finitely many possible values for $(\tilde{t}^*,\tilde{\zeta})$, this implies (20). ■

Sequential Rationality We now establish (18) and (19). By Lemma 4, there are two cases:

Case 1: Reporting histories satisfying $\xi(h_i^{R,t})=\left(0,\zeta_i^t\right)$ and acting histories satisfying $\xi(h_i^{A,t})=\left(0,\zeta_i^t\right)$.

Let $\Omega_0=\bigcup_{\tilde{f}\in\operatorname{supp}\tilde{\mu}}\left(0,\tilde{f}\right)$ be the set of all possible mediator states ω_0. Since $\Pr^{\phi^k}\left(\omega_t=\omega_0\forall t|\omega_0\right)\to 1$ for each $\omega_0\in\Omega_0$ by (17), $\xi(h_i^{R,t})=\left(0,\zeta_i^t\right)$ implies $\lim_{k\to\infty}\delta^k(\omega_T\in\Omega_0|h_i^{R,t})=1$, and $\xi(h_i^{A,t})=\left(0,\zeta_i^t\right)$ implies $\lim_{k\to\infty}\delta^k(\omega_T\in\Omega_0|h_i^{A,t})=1$.

For each i, t, and y_i^t, fix any action $m_{i,t}^*(y_i^t)\in A_{i,t}\setminus\hat{B}_{i,t}(y_i^t)$. With a slight abuse of notation, we write

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

where $m_{i,t}$ is the corresponding element of $h_i^{R,t+1}$. For each faithful history $h_i^{R,t}$ with $\delta^k\left(h_i^{R,t}\right)>0$, let $\lambda(h_i^{R,t})\in H_i^{R,t}$ denote the history where each message $m_{i,\tau}$ is replaced by $m_{i,\tau}^*(h_i^{R,\tau+1})\in A_{i,\tau}\setminus\hat{B}_{i,\tau}(\mathring{h}_i^{R,\tau})$ for every $\tau\leq t-1$. That is, we replace each action recommendation outside the support of $\tilde{\mu}$ with some fixed recommendation within the support. Note that $\Pr^{\tilde{\sigma}^k,\tilde{\mu}}\left(\lambda(h_i^{R,t})\right)>0$ whenever $\xi(h_i^{R,t})=\left(0,\zeta_i^t\right)$ and $\delta^k\left(h_i^{R,t}\right)>0$: this follows

because, given $\tilde{\sigma}^k$, players take each action profile a_t with probability at least $\varepsilon_k/|A_t| > 0$ at every history, and $\tilde{\mu}$ recommends all actions in $A_{i,t} \setminus \hat{B}_{i,t}(r_i^{t+1})$ in each period with positive probability. Define $\lambda(h_i^{A,t}) \in H_i^{A,t}$ analogously.

The following lemma confirms that, whenever $\xi(h_i^{R,t}) = \left(0, \zeta_i^t\right)$ or $\xi(h_i^{A,t}) = \left(0, \zeta_i^t\right)$, player i's beliefs in the constructed quasi-SE coincide with those in $\left(\tilde{\sigma}^k, \tilde{\mu}\right)$.

Lemma 5 *The following two claims hold:*

1. *For each $h_i^{R,t} \in J_i^{R,t}$ and ζ_i^t satisfying $\xi(h_i^{R,t}) = \left(0, \zeta_i^t\right)$ and each $y^t \in Y^t[\mathring{h}_i^{R,t}]$, we have*

$$\lim_{k\to\infty} \delta^k \left(y^t | h_i^{R,t}\right) = \tilde{\beta}_{i,t} \left(y^t | \lambda(h_i^{R,t})\right). \tag{24}$$

2. *For each $h_i^{A,t} \in J_i^{A,t}$ and ζ_i^t satisfying $\xi(h_i^{A,t}) = \left(0, \zeta_i^t\right)$ and each $y^t \in Y^t[\mathring{h}_i^{A,t}]$, we have*

$$\lim_{k\to\infty} \delta^k \left(y^t | h_i^{A,t}\right) = \tilde{\beta}_{i,t} \left(y^t | \lambda(h_i^{A,t})\right). \tag{25}$$

Lemma 5 follows from applying Bayes' rule inductively on t. We relegate the proof to the online appendix.

We now verify (18). $\xi(h_i^{R,t}) = \left(0, \zeta_i^t\right)$ implies $\delta^k(\omega_t \in \Omega_0 | h_i^{R,t}) = 1$. Given $\omega_t \in \Omega_0$ and $h^{R,t}$, (13) implies that the distribution of future recommendations m_τ follows $\prod_j \tilde{\phi}_{j,\tau}\left(m_{j,\tau} | r_j^{\tau+1}\right)$ for each $\tau \geq t$. Hence, Lemma 5 implies that (18) is equivalent to

$$\sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \tilde{\beta}_{i,t} \left(y^t | \lambda(h_i^{R,t})\right) \bar{u}_i \left(\sigma^*, \tilde{\phi} | \dot{h}_i^{R,t}, y^t\right) \geq \sum_{y^t \in Y^t[\mathring{h}_i^{R,t}]} \tilde{\beta}_{i,t} \left(y^t | \lambda(h_i^{R,t})\right) \bar{u}_i \left(\sigma_i', \sigma_{-i}^*, \tilde{\phi} | \dot{h}_i^{R,t}, y^t\right).$$

Finally, since the payoff-relevant component of $\lambda(h_i^{R,t})$ equals that of $h_i^{R,t}$, (10) implies (18). The proof for (19) is analogous.

Case 2: Reporting histories satisfying $\xi(h_i^{R,t}) = \left(t^*, \zeta_i^{t^*}\right)$ and acting histories satisfying $\xi(h_i^{A,t}) = \left(t^*, \zeta_i^{t^*}\right)$, for $t^* > 0$.

The next lemma confirms that, whenever $\xi(h_i^{R,t}) = \left(t^*, \zeta_i^{t^*}\right)$ or $\xi(h_i^{A,t}) = \left(t^*, \zeta_i^{t^*}\right)$, player i's beliefs in the constructed quasi-SE are given by $\beta_i^{\sigma^*, \pi_{t^*}}$.

Lemma 6 *The following two claims hold:*

1. *For each $h_i^{R,t} \in J_i^{R,t}$ and $\zeta_i^{t^*}$ satisfying $\xi(h_i^{R,t}) = \left(t^*, \zeta_i^{t^*}\right)$, each $f^{\geq t^*} \in F^{\geq t^*}$, and each $y^t \in Y^t[\mathring{h}_i^{R,t}]$, we have*

$$\lim_{k\to\infty} \delta^k \left(f^{\geq t^*}, y^t | h_i^{R,t}\right) = \beta_{i,t}^{\sigma^*, \pi_{t^*}} \left(f^{\geq t^*}, y^t | \mathring{h}_i^{R,t}, m_i^{t^*:t}\right). \tag{26}$$

2. *For each $h_i^{A,t} \in J_i^{A,t}$ and $\zeta_i^{t^*}$ satisfying $\xi(h_i^{A,t}) = \left(t^*, \zeta_i^{t^*}\right)$, each $f^{\geq t^*} \in F^{\geq t^*}$, and each $y^t \in Y^t[\mathring{h}_i^{A,t}]$, we have*

$$\lim_{k\to\infty} \delta^k \left(f^{\geq t^*}, y^t | h_i^{A,t}\right) = \beta_{i,t}^{\sigma^*, \pi_{t^*}} \left(f^{\geq t^*}, y^t | \mathring{h}_i^{A,t}, m_i^{t^*:t+1}\right). \tag{27}$$

Lemma 6 follows from another application of Bayes' rule. We again relegate the proof to the online 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 6, (18) is equivalent to (14), and therefore follows from the definition of $\pi^k_{t^*}$. The proof for (19) is analogous.

This completes the proof that $(\sigma, \phi, J, K, \beta)$ is a quasi-SE.

Final Construction Fix any canonical NE (σ^*, π^*) in which codominated actions are never recommended.[45] The proof is completed by mixing the "motivating" quasi-SE $(\sigma, \phi, J, K, \beta)$ with this NE (with almost all weight on the latter) to create a quasi-SE that implements the same outcome.

We construct a sequence of quasi-strategy profiles $\left(\bar{\sigma}^k, \bar{\phi}^k, J, K\right)$ indexed by k that limit to a quasi-SE profile $\left(\bar{\sigma}, \bar{\phi}, J, K\right)$ (with the same sets J and K as in the motivating quasi-SE) satisfying $\rho^{\bar{\sigma},\bar{\phi}} = \rho^{\sigma^*,\pi^*}$.

Players' strategies σ^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 the PE strategy $\hat{\sigma}_{i,t}(\mathring{h}_i^{R,t})$, and with probability $\sqrt{\varepsilon_k}$ she takes all actions with equal probability.

Mediator's strategy ϕ^k: At the beginning of the game, the mediator draws $f \in F^*$ according to π^* with probability $1 - \frac{1}{k}$ (and subsequently follows f), and the mediator follows quasi-strategy ϕ^k with probability $\frac{1}{k}$.

Letting $\left(\bar{\sigma}, \bar{\phi}\right) = \lim_{k\to\infty}\left(\bar{\sigma}^k, \bar{\phi}^k\right)$, we have $\rho^{\bar{\sigma},\bar{\phi}} = \rho^{\sigma^*,\pi^*}$.

Since J includes all faithful histories where no codominated actions have been recommended, $\left(\bar{\sigma}, \bar{\phi}, J, K\right)$ is valid. For each i, t, $h_i^{R,t} \in J_i^{R,t}$, and $h^{R,t}$ with i-component $h_i^{R,t}$, define

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

Define $\bar{\beta}_{i,t}(h^{A,t}|h_i^{A,t})$ analogously. Since $\Pr^{\bar{\sigma}^k,\bar{\phi}^k}(h_i^{R,t}) > 0$ for each $h_i^{R,t} \in J_i^{R,t}$ conditional on the mediator following ϕ^k, $\bar{\beta}$ is well-defined, and hence Kreps-Wilson consistent.

To prove that $\left(\bar{\sigma}, \bar{\phi}, J, K, \bar{\beta}\right)$ is a quasi-SE, it remains to verify sequential rationality. Under belief system $\bar{\beta}$, so long as a player i has been faithful and has not observed a signal or recommendation that occurs with probability 0 conditional on the mediator following π^*, she believes that with probability 1 the mediator is following π^* and other players have been faithful so far. At such a history, it is optimal for player i to be faithful, since (σ^*, π^*) is a NE. On the other hand, if player i has been faithful and does observe a signal or recommendation that occurs with probability 0 conditional on mediator strategy π^*, then she believes with probability 1 that the mediator is following ϕ^k and other players have been faithful. In this case, faithfulness is optimal by (18) and (19).

[45] Note that if (σ^{**}, π^*) is a canonical NE for some canonical (but possibly not fully canonical) player strategy profile σ^{**}, then (σ^*, π^*) is also a canonical NE, where σ^* denotes the fully canonical player strategy profile. One way of seeing this is to note that the strategy profile constructed in the proof of Proposition 2 is fully canonical.

Online Appendix

E Extending the Opening Example

In the extended example, there are four players (in addition to the mediator) and four periods. The roles of players 2 and 4 are similar to those of players 1 and 2 in the original example, respectively. The timing is as follows:

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 $\theta \in \{n, p\}$ such that $\theta = n$ with probability 3/4. The mediator takes $a_0 \in \{A_0, B_0\}$.

Period 4. The mediator and player 4 observe $s \in \{0, 1\}$, where $s = 0$ if $a_1 = A_1$ and either $a_0 = A_0 \wedge a_2 = A_2$ or $a_0 = B_0 \wedge a_2 = B_2$. Player 4 takes $a_4 \in \{N, P\}$.

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

	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	$-3 - 1_{\{a_4=P\}}$	$-3 - 1_{\{a_4=P\}}$

$a_1 = A_1$

	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

$a_1 = B_1$

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

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 $\mathfrak{C} = \mathfrak{C}^*$.

Implementability with $\mathfrak{C} \neq \mathfrak{C}^*$ Again, it suffices to implement the target distribution in a canonical NE in which players avoid codominated actions. 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 $m_0 \in \{A_0, B_0\}$ with equal probability, and recommends $m_2 = A_2 \wedge m_3 = A_3$ if $m_0 = A_0$ and recommends $m_2 = B_2 \wedge m_3 = B_3$ if $m_0 = B_0$. 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 $m_0 = A_0$, $m_2 = A_2$, $m_3 = A_3$, and $m_4 = N$.

It is straightforward to check that this is a NE. Moreover, no codominated actions are recommended: For player 4, N is weakly dominant and hence never codominated, while P is recommended only following $s = 1$. Hence, we need only check that P is not codominated following $s = 1$. But this holds, because the event $(a_1, a_2, \theta) = (A_1, C_2, p)$ is compatible with $s = 1$, and in this event P is optimal.

Given that P is not codominated for player 4 following $s = 1$, 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. Finally, given that $a_2 = C_2$ and $a_4 = P$ are not codominated after $a_1 = B_1$, action $a_1 = A_1$ is not codominated for player 1.

Non-Implementability with $\mathfrak{C} = \mathfrak{C}^*$ 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 outcome, assume without loss that a_i is played with positive probability after $m_i = a_i$. 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 to $a_1 = B_1$ 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. We have thus established that, for players $i \in \{2, 3, 4\}$, $a_i = m_i$ with equilibrium probability 1 at every history.

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 $\Pr(a_1 = B_1 | m_2 = C_2) = 1$. Therefore,

$$\begin{aligned}
& \Pr(a_1 = B_1 | m_2 = C_2, a_2 = A_2) \\
= & \frac{\Pr(a_1 = B_1) \Pr(m_2 = C_2 | a_1 = B_1) \Pr(a_2 = A_2 | a_1 = B_1, m_2 = C_2)}{\Pr(m_2 = C_2) \Pr(a_2 = A_2 | m_2 = C_2)} \\
= & \frac{\Pr(a_1 = B_1) \Pr(m_2 = C_2 | a_1 = B_1) \Pr(a_2 = A_2 | m_2 = C_2)}{\Pr(m_2 = C_2) \Pr(a_2 = A_2 | m_2 = C_2)} \\
= & \frac{\Pr(a_1 = B_1) \Pr(m_2 = C_2 | a_1 = B_1)}{\Pr(m_2 = C_2)} \\
= & \Pr(a_1 = B_1 | m_2 = C_2) = 1.
\end{aligned}$$

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 $a_1 = B_1$ and $a_4 = P$. Since $a_1 = B_1$ implies $s = 1$ and player 2 can always report as if she took $a_2 = C_2$, this implies that

$$\Pr(a_1 = B_1, a_4 = P | m_2 = C_2, a_2 = A_2) \leq \Pr(a_1 = B_1, a_4 = P | m_2 = C_2, a_2 = C_2). \quad (28)$$

Note that if $\Pr(a_1 = B_1, a_4 = P | m_2 = C_2, a_2 = A_2) < 1$ then player 2 would deviate to A_2 after $m_2 = C_2$. So this probability must equal 1, and hence (28) implies

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

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

$$\Pr\left(a_1 = B_1, a_2 = C_2, m_4 = P, s = 1 | m_2 = C_2\right) = 1. \tag{29}$$

On the other hand, since player 4 is willing to take P after $s = 1$ and $m_4 = P$, we have $\Pr\left(a_1 = A_1, a_2 = C_2, \theta = p | s = 1, m_4 = P\right) = 1$. In particular,

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

Since $(a + c)/(b + d) \leq (a/b) + (c/d)$ for all non-negative numbers a, b, c, d, the left-hand side is no more than

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

Note that by (29),

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

and

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

Hence,

$$\begin{aligned}
1 &\leq \sum_{m_2 \neq C_2} \frac{\Pr\left(a_1 = A_1, a_2 = C_2, \theta = p, m_4 = P, s = 1 | m_2\right)}{\sum_{a_1, a_2} \Pr\left(a_1, a_2, m_4 = P, s = 1 | m_2\right)} \\
&\leq \sum_{m_2 \neq C_2} \frac{\Pr\left(a_1 = A_1, a_2 = C_2, \theta = p, m_4 = P, s = 1 | m_2\right)}{\Pr\left(a_1 = A_1, a_2 = C_2, m_4 = P, s = 1 | m_2\right)} \\
&= \sum_{m_2 \neq C_2} \frac{\Pr\left(a_1 = A_1, a_2 = C_2, \theta = p, m_4 = P | m_2\right)}{\Pr\left(a_1 = A_1, a_2 = C_2, m_4 = P | m_2\right)} \\
&= \sum_{m_2 \neq C_2} \frac{\Pr\left(a_1 = A_1, a_2 = C_2 | m_2\right) \Pr\left(\theta = p, m_4 = P | a_1 = A_1, m_2, a_2 = C_2\right)}{\Pr\left(a_1 = A_1, a_2 = C_2 | m_2\right) \Pr\left(m_4 = P | a_1 = A_1, m_2, a_2 = C_2\right)} \\
&= \sum_{m_2 \neq C_2} \frac{\Pr\left(\theta = p, m_4 = P | a_1 = A_1, m_2, a_2 = C_2\right)}{\Pr\left(m_4 = P | a_1 = A_1, m_2, a_2 = C_2\right)}, \qquad (30)
\end{aligned}$$

where the second line drops the event $(a_1, a_2) \neq (A_1, C_2)$ from the denominator 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 when player 4 takes N if $a_1 = A_1$, and player 2 is indifferent between player 4's actions if $a_1 = B_1$. Moreover, $\Pr\left(a_1 = A_1 | m_2\right) > 0$ for each $m_2 \neq C_2$. Hence, for each $m_2 \neq C_2$ and θ, after $(m_2, a_2 = C_2, \theta)$ player 2 chooses her report $\left(\hat{a}_2, \hat{\theta}\right)$ to minimize the conditional probability that $a_4 = P$ given $a_1 = A_1$, and hence to minimize the conditional probability that $m_4 = P$ given $a_1 = A_1$ (since $a_4 = m_4$ with probability 1). Therefore, for each $m_2 \neq C_2$,

$$\begin{aligned}
& \Pr\left(\theta = p, m_4 = P | a_1 = A_1, m_2, a_2 = C_2\right) \\
= \; & \Pr\left(\theta = p | a_1 = A_1, m_2, a_2 = C_2\right) \Pr\left(m_4 = P | a_1 = A_1, m_2, a_2 = C_2, \theta = p\right) \\
= \; & \Pr\left(\theta = p\right) \Pr\left(m_4 = P | a_1 = A_1, m_2, a_2 = C_2, \theta = p\right) \\
= \; & \Pr\left(\theta = p\right) \min_{\left(\hat{a}_2, \hat{\theta}\right)} \Pr\left(m_4 = P | a_1 = A_1, m_2, a_2 = C_2, \theta = p, \hat{a}_2, \hat{\theta}\right) \\
= \; & \Pr\left(\theta = p\right) \min_{\left(\hat{a}_2, \hat{\theta}\right)} \Pr\left(m_4 = P | a_1 = A_1, m_2, a_2 = C_2, \hat{a}_2, \hat{\theta}\right) \\
= \; & \Pr\left(\theta = p\right) \Pr\left(m_4 = P | a_1 = A_1, m_2, a_2 = C_2\right),
\end{aligned}$$

where the fourth equality follows since the distribution of m_4 is independent of θ conditional

on $\left(\hat{a}_2, \hat{\theta}\right)$. Thus,

$$\begin{aligned}
& \sum_{m_2 \neq C_2} \frac{\Pr\left(\theta = p, m_4 = P | a_1 = A_1, m_2, a_2 = C_2\right)}{\Pr\left(m_4 = P | a_1 = A_1, m_2, a_2 = C_2\right)} \\
= & \sum_{m_2 \neq C_2} \frac{\Pr\left(\theta = p\right) \Pr\left(m_4 = P | a_1 = A_1, m_2, a_2 = C_2\right)}{\Pr\left(m_4 = P | a_1 = A_1, m_2, a_2 = C_2\right)} \\
= & \Pr\left(\theta = p\right) \sum_{m_2 \neq C_2} 1 = \frac{1}{4} \times 2 = \frac{1}{2}.
\end{aligned}$$

This contradicts (30).

F Proof of Lemma 5

We prove (24); the proof of (25) is analogous. We will prove the following: for each i, t, faithful history $h_i^{R,t}$ with $\delta^k\left(h_i^{R,t}\right) > 0$, ζ_i^t, and $y^t \in Y^t[\mathring{h}_i^{R,t}]$, there exist numbers $\varphi_k^R(h_i^{R,t}, \zeta_i^t) \geq 0$ and $e_k^R\left(h_i^{R,t}, \zeta_i^t, y^t\right) \geq 0$ such that

$$\begin{aligned}
& \delta^k(h_i^{R,t}, \zeta_i^t, y^t | \omega_T \in \Omega_0) = \varphi_k^R(h_i^{R,t}, \zeta_i^t) \left(\Pr^{\tilde{\sigma}^k, \tilde{\mu}}(\lambda(h_i^{R,t}), y^t) + e_k^R\left(h_i^{R,t}, \zeta_i^t, y^t\right)\right), \text{[46]} \\
& \lim_{k \to \infty} \frac{e_k^R\left(h_i^{R,t}, \zeta_i^t, y^t\right)}{\left(\frac{1}{k}\right)^{2(L+1)T}} \leq t.
\end{aligned} \tag{31}$$

(31) is sufficient for (24), since the former implies, for each $\zeta_i^t \in \{0,1\}^{t-1}$ and $y^t \in Y^t[\mathring{h}_i^{R,t}]$,

$$\begin{aligned}
& \lim_{k \to \infty} \delta^k\left(y^t | \zeta_i^t, h_i^{R,t}\right) \\
= & \lim_{k \to \infty} \delta^k\left(y^t | \zeta_i^t, h_i^{R,t}, \omega_T \in \Omega_0\right) \quad (\text{by } \xi(h_i^{R,t}) = \left(0, \zeta_i^t\right) \text{ and (17)}) \\
= & \lim_{k \to \infty} \frac{\varphi_k^R(h_i^{R,t}, \zeta_i^t) \left(\Pr^{\tilde{\sigma}^k, \tilde{\mu}}(\lambda(h_i^{R,t}), y^t) + e_k^R\left(h_i^{R,t}, \zeta_i^t, y^t\right)\right)}{\sum_{\tilde{y}^t \in Y^t[\mathring{h}_i^{R,t}]} \varphi_k^R(h_i^{R,t}, \zeta_i^t) \left(\Pr^{\tilde{\sigma}^k, \tilde{\mu}}(\lambda(h_i^{R,t}), \tilde{y}^t) + e_k^R\left(h_i^{R,t}, \zeta_i^t, \tilde{y}^t\right)\right)} \quad (\text{by (31)}) \\
= & \lim_{k \to \infty} \frac{\Pr^{\tilde{\sigma}^k, \tilde{\mu}}(\lambda(h_i^{R,t}), y^t) + e_k^R\left(h_i^{R,t}, \zeta_i^t, y^t\right)}{\sum_{\tilde{y}^t \in Y^t[\mathring{h}_i^{R,t}]} \Pr^{\tilde{\sigma}^k, \tilde{\mu}}(\lambda(h_i^{R,t}), \tilde{y}^t) + \sum_{\tilde{y}^t \in Y^t[\mathring{h}_i^{R,t}]} e_k^R\left(h_i^{R,t}, \zeta_i^t, \tilde{y}^t\right)} \\
= & \lim_{k \to \infty} \frac{\Pr^{\tilde{\sigma}^k, \tilde{\mu}}(\lambda(h_i^{R,t}), y^t)}{\sum_{\tilde{y}^t \in Y^t[\mathring{h}_i^{R,t}]} \Pr^{\tilde{\sigma}^k, \tilde{\mu}}(\lambda(h_i^{R,t}), \tilde{y}^t)} \\
= & \tilde{\beta}_i\left(\mathring{h}_{-i}^{R,t} | \lambda(h_i^{R,t})\right),
\end{aligned}$$

[46] There is a slight redundancy in this notation: the payoff-relevant part of $\lambda(h_i^{R,t})$ equals y_i^t, since the payoff-relevant component of $\lambda(h_i^{R,t})$ equals $\mathring{h}_i^{R,t}$ and $y^t \in Y^t[\mathring{h}_i^{R,t}]$.

where the second-to-last equality follows as $\Pr^{\tilde{\sigma}^k,\tilde{\mu}}(\lambda(h_i^{R,t}),y^t) \geq \varepsilon_1\varepsilon_2\left(\varepsilon_k\right)^{NT}/(T\left|A\right|^T)$, $e_k^R\left(h_i^{R,t},\zeta_i^t,y^t\right) \leq \left(\frac{1}{k}\right)^{2(L+1)T} T$, and $k\left(\varepsilon_k\right)^{NT} \to \infty$.

We prove (31) by induction on t. Taking $\varphi_k^R(h_i^{R,1},\zeta_i^1)=1$ and $e_k^R\left(h_i^{R,1},\zeta_i^1,y^1\right)=0$, (31) holds for $t=1$.

Suppose it holds for t. We prove it holds for $t+1$. For the rest of the proof, arbitrarily fix $h_i^{R,t+1} \in J_i^{R,t+1}$, $\zeta_i^t \in \{0,1\}^{t-1}$, $y^{t+1} \in Y^{t+1}[\mathring{h}_i^{R,t+1}]$, and $\zeta_{i,t} \in \{0,1\}$. Whenever we write (a_t,s_{t+1}), it means the components in y^{t+1}.

Since θ, ζ, and randomizations under $\left(\tilde{\sigma}^k,\tilde{\mu}\right)$ are independent across players, we have

$$
\begin{aligned}
& \delta^k(h_i^{R,t+1},\zeta_i^t,y^{t+1},\zeta_{i,t}|\omega_T \in \Omega_0) \\
= \; & \delta^k(h_i^{R,t},\zeta_i^t,y^t|\omega_T \in \Omega_0) \\
& \times \left(\begin{array}{c} 1_{\left\{\zeta_{i,t}=0,m_{i,t}=a_{i,t}\in A_{i,t}\setminus\hat{B}_{i,t}(\mathring{h}_i^{R,t})\right\}}\left(1-\sqrt{\varepsilon_k}\right)\left(1-\left(\frac{1}{k}\right)^{2(L+1)T}\right)\hat{\sigma}_{i,t}(\mathring{h}_i^{R,t})(a_{i,t}) \\ +1_{\left\{\zeta_{i,t}=0,m_{i,t}=\star\right\}}\sqrt{\varepsilon_k}\left(1-\left(\frac{1}{k}\right)^{2(L+1)T}\right)\left(\left(1-\sqrt{\varepsilon_k}\right)\hat{\sigma}_{i,t}(\mathring{h}_i^{R,t})(a_{i,t})+\frac{\sqrt{\varepsilon_k}}{|A_{i,t}|}\right) \\ +1_{\left\{\zeta_{i,t}=1,m_{i,t}=a_{i,t}\in A_{i,t}\setminus\mathfrak{D}_{i,t}(\mathring{h}_i^{R,t})\right\}}\left(\frac{1}{k}\right)^{2(L+1)T}\frac{1}{|A_{i,t}|-\left|\mathfrak{D}_{i,t}(\mathring{h}_i^{R,t})\right|} \end{array} \right) \\
& \times \sum_{\theta_{-i,t},\zeta_{-i,t}} \left(\begin{array}{l} \prod_{j\neq i:\theta_{j,t}=0,\zeta_{j,t}=0}\left(1-\sqrt{\varepsilon_k}\right)\left(1-\left(\frac{1}{k}\right)^{2(L+1)T}\right)\hat{\sigma}_{j,t}(y_j^t)(a_{j,t}) \\ \times\prod_{j\neq i:\theta_{j,t}=1,\zeta_{j,t}=0}\sqrt{\varepsilon_k}\left(1-\left(\frac{1}{k}\right)^{2(L+1)T}\right)\left(\left(1-\sqrt{\varepsilon_k}\right)\hat{\sigma}_{j,t}(y_j^t)(a_{j,t})+\frac{\sqrt{\varepsilon_k}}{|A_{j,t}|}\right) \\ \times\prod_{j\neq i:\zeta_{j,t}=1,a_{j,t}\in A_{j,t}\setminus\mathfrak{D}_{j,t}(y_j^t)}\left(\frac{1}{k}\right)^{2(L+1)T}\frac{1}{|A_{j,t}|-\left|\mathfrak{D}_{j,t}(y_j^t)\right|} \end{array} \right) \\
& \times p(s_{t+1}|y^t,a_t). \qquad (32)
\end{aligned}
$$

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

$$\varphi_k^R(h_i^{R,t},\zeta_i^t)\left(\Pr^{\tilde{\sigma}^k,\tilde{\mu}}(\lambda(h_i^{R,t}),y^t)+e_k^R\left(h_i^{R,t},\zeta_i^t,y^t\right)\right).$$

Note that

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

Define

$$\tilde{\varphi}_k\left(\mathring{h}_i^{R,t},m_{i,t},a_{i,t},\zeta_{i,t}\right)=\left(\begin{array}{c} 1_{\left\{\zeta_{i,t}=0,m_{i,t}=a_{i,t}\in A_{i,t}\setminus\hat{B}_{i,t}(\mathring{h}_i^{R,t})\right\}}\left(1-\sqrt{\varepsilon_k}\right)\left(1-\left(\frac{1}{k}\right)^{2(L+1)T}\right)\hat{\sigma}_{i,t}(\mathring{h}_i^{R,t})(a_{i,t}) \\ +1_{\left\{\zeta_{i,t}=0,m_{i,t}=\star\right\}}\sqrt{\varepsilon_k}\left(1-\left(\frac{1}{k}\right)^{2(L+1)T}\right)\left(\left(1-\sqrt{\varepsilon_k}\right)\hat{\sigma}_{i,t}(\mathring{h}_i^{R,t})(a_{i,t})+\frac{\sqrt{\varepsilon_k}}{|A_{i,t}|}\right) \\ +1_{\left\{\zeta_{i,t}=1,m_{i,t}=a_{i,t}\in A_{i,t}\setminus\mathfrak{D}_{i,t}(\mathring{h}_i^{R,t})\right\}}\left(\frac{1}{k}\right)^{2(L+1)T}\frac{1}{|A_{i,t}|-\left|\mathfrak{D}_{i,t}(\mathring{h}_i^{R,t})\right|} \end{array} \right)$$

and

$$\begin{aligned}\tilde{e}_k\left(y^{t+1}\right) &= \sum_{\theta_{-i,t},\zeta_{-i,t}}\left(\begin{array}{l}\prod_{j\neq i:\theta_{j,t}=0,\zeta_{j,t}=0}\left(1-\sqrt{\varepsilon_k}\right)\left(1-\left(\frac{1}{k}\right)^{2(L+1)T}\right)\hat{\sigma}_{j,t}(y_j^t)(a_{j,t})\\ \times\prod_{j\neq i:\theta_{j,t}=1,\zeta_{j,t}=0}\sqrt{\varepsilon_k}\left(1-\left(\frac{1}{k}\right)^{2(L+1)T}\right)\left(\left(1-\sqrt{\varepsilon_k}\right)\hat{\sigma}_{j,t}(y_j^t)(a_{j,t})+\frac{\sqrt{\varepsilon_k}}{|A_{j,t}|}\right)\\ \times\prod_{j\neq i:\zeta_{j,t}=1,a_{j,t}\in A_{j,t}\setminus\mathfrak{D}_{j,t}(y_j^t)}\left(\frac{1}{k}\right)^{2(L+1)T}\frac{1}{|A_{j,t}|-\left|\mathfrak{D}_{j,t}(y_j^t)\right|}\end{array}\right)\\ &\quad-\Pr^{\tilde{\sigma}^k,\tilde{\mu}}(a_{-i,t}|\lambda(h_i^{R,t}),y^t)\\ &= \left(\frac{1}{k}\right)^{2(L+1)T}\prod_{j\neq i}\left(\frac{1_{\left\{a_{j,t}\in A_{j,t}\setminus\mathfrak{D}_{j,t}(y_j^t)\right\}}}{|A_{j,t}|-\left|\mathfrak{D}_{j,t}(y_j^t)\right|}-(1-\varepsilon_k)\hat{\sigma}_{j,t}(y_j^t)(a_{j,t})-\frac{\varepsilon_k}{|A_{j,t}|}\right).\end{aligned}$$

Substituting these into (32), we have

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

Next, define

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

We can write

$$\begin{aligned}&\delta^k(h_i^{R,t+1},\zeta_i^{t+1},y^{t+1}|\omega_T\in\Omega_0)\\ =\;&\varphi_k^R(h_i^{R,t+1},\zeta_i^{t+1})\left(\begin{array}{c}\Pr^{\tilde{\sigma}^k,\tilde{\mu}}(\lambda(h_i^{R,t}),y^t)\times\Pr^{\tilde{\sigma}^k,\tilde{\mu}}(a_{-i,t}|\lambda(h_i^{R,t}),y^t)\\ \times\Pr^{\tilde{\sigma}^k,\tilde{\mu}}\left(m_{i,t}^*(h_i^{R,t+1}),a_{i,t}|\mathring{h}_i^{R,t}\right)\times p(s_{t+1}|y^t,a_t)\\ +e_k^R\left(h_i^{R,t+1},\zeta_i^t,y^{t+1}\right)\end{array}\right),\end{aligned}\tag{34}$$

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

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

Since $h_i^{R,t}$ is faithful, the distribution of player i's message and action in period t is fully determined by her own payoff-relevant history $\mathring{h}_i^{R,t}$. Hence,

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

Given this equality, we have

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

Substituting this into (34), we have

$$\begin{aligned}&\delta^k(h_i^{R,t+1},\zeta_i^{t+1},y^{t+1}|\omega_T\in\Omega_0)\\ =\;&\varphi_k^R(h_i^{R,t+1},\zeta_i^{t+1})\left(\Pr^{\tilde{\sigma}^k,\tilde{\mu}}(\lambda(h_i^{R,t+1}),y^{t+1})+e_k^R\left(h_i^{R,t+1},\zeta_i^t,y^{t+1}\right)\right).\end{aligned}$$

Finally, we have

$$\begin{aligned}\lim_{k\to\infty}\frac{e_k^R\left(h_i^{R,t+1},\zeta_i^{t+1},y^{t+1}\right)}{\left(\frac{1}{k}\right)^{2(L+1)T}}&\leq\lim_{k\to\infty}\frac{\tilde{e}_k\left(y^{t+1}\right)}{\left(\frac{1}{k}\right)^{2(L+1)T}}+\frac{e_k^R\left(h_i^{R,t},\zeta_i^t,y^t\right)}{\left(\frac{1}{k}\right)^{2(L+1)T}}\left(1+\tilde{e}_k\left(y^{t+1}\right)\right)\\ &\leq 1+t,\end{aligned}$$

where the last line uses $\lim_{k\to\infty}\frac{\tilde{e}_k\left(y^{t+1}\right)}{\left(\frac{1}{k}\right)^{2(L+1)T}}\leq 1$ (and hence $\tilde{e}_k\left(y^{t+1}\right)\to 0$ by (16)) and the inductive hypothesis that $\lim_{k\to\infty}\frac{e_k^R\left(h_i^{R,t},\zeta_i^t,y^t\right)}{\left(\frac{1}{k}\right)^{2(L+1)T}}\leq t$. Hence, (31) holds for $t+1$, as desired.

G Proof of Lemma 6

We prove (26); the proof of (27) is analogous. Let $y_i^t=\mathring{h}_i^{R,t}$. By definition of $\beta_{i,t}^{\sigma^*,\pi_{t^*}}$, (26) is equivalent to

$$\lim_{k\to\infty}\delta^k\left(f^{\geq t^*},y^t|h_i^{R,t}\right)=\lim_{k\to\infty}\frac{\pi_{t^*}^k(f^{\geq t^*},y^{t^*})\Pr^{\sigma^*}\left(y^t|f^{\geq t^*},y^{t^*}\right)}{\sum_{\tilde{y}^t\in Y^t[y_i^t],\tilde{f}^{\geq t^*}}\pi_{t^*}^k\left(\tilde{f}^{\geq t^*},\tilde{y}^{t^*}\right)\Pr^{\sigma^*}\left(\tilde{y}^t|\tilde{f}^{\geq t^*},\tilde{y}^{t^*}\right)}.\tag{35}$$

From the definition of δ^k, $\delta^k\left(f^{\geq t^*},y^{R,t}|h_i^{R,t},t^*,\zeta_i^{t^*}\right)$ equals

$$\frac{A\sum_{f_i^{<t^*},\theta_i^{t^*},\left(f_j^{<t^*},\theta_j^{t^*},m_j^{t^*}\right)_{j\neq i}}\pi_{t^*}^k\left(f^{\geq t^*},y^{t^*}\right)\Pr^{\sigma^*}\left(y^t|f^{\geq t^*},y^{t^*}\right)B_i\prod_{j\neq i}C_jD}{A\sum_{\tilde{y}^t\in Y^t[y_i^t],\tilde{f}^{\geq t^*}}\sum_{f_i^{<t^*},\theta_i^{t^*},\left(f_j^{<t^*},\theta_j^{t^*},m_j^{t^*}\right)_{j\neq i}}\pi_{t^*}^k\left(\tilde{f}^{\geq t^*},\tilde{y}^{t^*}\right)\Pr^{\sigma^*}\left(\tilde{y}^t|\tilde{f}^{\geq t^*},\tilde{y}^{t^*}\right)B_i\prod_{j\neq i}\tilde{C}_j\tilde{D}},\tag{36}$$

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

$$\begin{array}{ll} A = \left(\frac{1}{k}\right)^{(L+1)t^*+2(L+1)T\left|\zeta_i^{t^*}\right|}, & B_i = \frac{1}{\#M_i(y_i^{t^*})}, \\ C_j = \frac{1}{\#M_j(y_j^{t^*})}, & \tilde{C}_j = \frac{1}{\#M_j(\tilde{y}_j^{t^*})}, \\ D = 1_{\left\{m^{t^*} \in M^{t^*}(y^{t^*}, f^{<t^*}, \theta^{t^*}, \zeta_i^{t^*})\right\}}, & \tilde{D} = 1_{\left\{m^{t^*} \in M^{t^*}(\tilde{y}^{t^*}, f^{<t^*}, \theta^{t^*}, \zeta_i^{t^*})\right\}}. \end{array}$$

Note that A and B_i cancel in (36). Moreover, we have

$$\begin{aligned} D &= D_i^{00} \times D_i^{01} \times D_i^{\bullet 1} \times \textstyle\prod_{j \neq i} \left(D_j^0 \times D_j^1\right), \\ \tilde{D} &= D_i^{00} \times D_i^{01} \times D_i^{\bullet 1} \times \textstyle\prod_{j \neq i} \left(\tilde{D}_j^0 \times \tilde{D}_j^1\right), \end{aligned}$$

where

$$\begin{array}{ll} D_i^{00} = 1_{\left\{m_{i,\tau} = f_{i,\tau}^{<t^*}(y_i^\tau)\ \forall \tau \leq t^*-1 \text{ s.t. } \zeta_{i,\tau} = \theta_{i,\tau} = 0\right\}}, & D_i^{01} = 1_{\left\{m_{i,\tau} = \star\ \forall \tau \leq t^*-1 \text{ s.t. } \zeta_{i,\tau} = 0 \text{ and } \theta_{i,\tau} = 1\right\}}, \\ D_i^{\bullet 1} = 1_{\left\{m_{i,\tau} \in A_{i,\tau} \setminus \mathfrak{D}_{i,\tau}(y_i^\tau)\ \forall \tau \leq t^*-1 \text{ s.t. } \zeta_{i,\tau} = 1\right\}}, & D_j^0 = 1_{\left\{m_{j,\tau} = f_{j,\tau}^{<t^*}(y_j^\tau)\ \forall \tau \leq t^*-1 \text{ s.t. } \theta_{j,\tau} = 0\right\}}, \\ D_j^1 = 1_{\{m_{j,\tau} = \star\ \forall \tau \leq t^*-1 \text{ s.t. } \theta_{j,\tau} = 1\}}, & \tilde{D}_j^0 = 1_{\left\{m_{j,\tau} = \tilde{f}_{j,\tau}^{<t^*}(\tilde{y}_j^\tau)\ \forall \tau \leq t^*-1 \text{ s.t. } \theta_{j,\tau} = 0\right\}}, \\ \tilde{D}_j^1 = 1_{\{m_{j,\tau} = \star\ \forall \tau \leq t^*-1 \text{ s.t. } \theta_{j,\tau} = 1\}}. & \end{array}$$

(The pneumonic here is that the first superscript of D_i indicates the value of $\theta_i \in \{0,1\}$, where $\bullet$ indicates that θ_i is not specified, and the second superscript indicates the value of $\zeta_i \in \{0,1\}$. For player $j \neq i$, the superscript of D_j indicates the value of $\theta_j \in \{0,1\}$.)

Having cancelled A and B_i, since (i) the term $D_i^{00} D_i^{01} D_i^{\bullet 1}$ does not depend on $\left(f_j^{<t^*}, \theta_j^{t^*}, m_j^{t^*}\right)_{j \neq i}$, (ii) $D_i^{00} D_i^{01} D_i^{\bullet 1}$ is the only term that depends on $(f_i^{<t^*}, \theta_i^{t^*})$ in the numerator of (36), and (iii) $\prod_{j \neq i} \left(C_j D_j^0 D_j^1\right)$ is the only term that depends on $\left(f_j^{<t^*}, \theta_j^{t^*}, m_j^{t^*}\right)_{j \neq i}$, the numerator of (36) equals

$$\left(\sum_{f_i^{<t^*}, \theta_i^{t^*}} D_i^{00} D_i^{01} D_i^{\bullet 1} \right) \times \left(\pi_{t^*}^k \left(f^{\geq t^*}, y^{t^*}\right) \Pr^{\sigma^*}\left(y^t | f^{\geq t^*}, y^{t^*}\right) \sum_{\left(f_j^{<t^*}, \theta_j^{t^*}, m_j^{t^*}\right)_{j \neq i}} \prod_{j \neq i} \left(C_j D_j^0 D_j^1\right) \right).$$

Since $D_i^{00} D_i^{01} D_i^{\bullet 1}$ also does not depend on $f^{\geq t}$ and y_{-i}^t, the denominator of (36) equals

$$\left(\sum_{f_i^{<t^*}, \theta_i^{t^*}} D_i^{00} D_i^{01} D_i^{\bullet 1} \right) \times \left(\sum_{\tilde{y}^t \in Y^t[y_i^t], \tilde{f}^{\geq t^*}} \pi_{t^*}^k \left(\tilde{f}^{\geq t}, \tilde{y}^{t^*}\right) \Pr^{\sigma^*}\left(\tilde{y}^t | \tilde{f}^{\geq t^*}, \tilde{y}^{t^*}\right) \sum_{\left(f_j^{<t^*}, \theta_j^{t^*}, m_j^{t^*}\right)_{j \neq i}} \prod_{j \neq i} \left(C_j D_j^0 D_j^1\right) \right).$$

Moreover, $\sum_{f_j^{<t^*},\theta_j^{t^*},m_j^{t^*}} C_j D_j^0 D_j^1 = 1$ by (15). Hence, (36) equals

$$\frac{\pi_{t^*}^k(f^{\geq t^*}, y^{t^*}) \Pr^{\sigma^*}\left(y^t | f^{\geq t^*}, y^{t^*}\right)}{\sum_{\tilde{y}^t \in Y^t[y_i^t], \tilde{f}^{\geq t^*}} \pi_{t^*}^k\left(\tilde{f}^{\geq t^*}, \tilde{y}^{t^*}\right) \Pr^{\sigma^*}\left(\tilde{y}^t | \tilde{f}^{\geq t^*}, \tilde{y}^{t^*}\right)}.$$

Taking the limit $k \to \infty$, we obtain (35).

H Proof of Claims 1–4 of Proposition 4

We postpone the proof of Claim 5 to Appendix J, since it relies on results proved in Appendix I.

H.1 Proof of Claim 1

We first show that, if an outcome can be implemented in a NE in which no player can detect another's unilateral deviation, it can also be implemented in a canonical NE in which no player can detect another's unilateral deviation.

Lemma 7 *For any game $G = (\Gamma, \mathfrak{C})$ and NE (σ, ϕ) with $\operatorname{supp} \rho_i^{\sigma,\phi} = \bigcup_{j \neq i,0} \bigcup_{\sigma'_j \in \Sigma_j} \operatorname{supp} \rho_i^{\sigma'_j, \sigma_{-j}, \phi}$ for all $i \neq 0$, there exists a canonical NE $\left(\tilde{\sigma}, \tilde{\phi}\right)$ in game $G^* = (\Gamma, \mathfrak{C}^*)$ such that $\rho^{\sigma,\phi} = \rho^{\tilde{\sigma},\tilde{\phi}}$ and $\operatorname{supp} \rho_i^{\tilde{\sigma},\tilde{\phi}} = \bigcup_{j \neq i,0} \bigcup_{\tilde{\sigma}'_j \in \Sigma_j^*} \operatorname{supp} \rho_i^{\tilde{\sigma}'_j, \tilde{\sigma}_{-j}, \tilde{\phi}}$ for all $i \neq 0$.*

Proof. Fix such a G and (σ, ϕ). Let $\left(\tilde{\sigma}, \tilde{\phi}\right)$ be the profile in G^* constructed in the proof of Proposition 2. Recall that $\rho^{\tilde{\sigma},\tilde{\phi}} = \rho^{\sigma,\phi}$. By Lemma 1, for each $i \neq 0$, $j \neq i$, and $\tilde{\sigma}'_j \in \Sigma_j^*$, there exists a strategy $\sigma'_j \in \Sigma_j$ such that $\rho_i^{\sigma'_j, \sigma_{-j}, \phi} = \rho_i^{\tilde{\sigma}'_j, \tilde{\sigma}_{-j}, \tilde{\phi}}$. Hence, we have

$$\operatorname{supp} \rho_i^{\tilde{\sigma},\tilde{\phi}} = \operatorname{supp} \rho_i^{\sigma,\phi} = \bigcup_{j \neq i,0} \bigcup_{\sigma'_j \in \Sigma_j} \operatorname{supp} \rho_i^{\sigma'_j, \sigma_{-j}, \phi} \supset \bigcup_{j \neq i,0} \bigcup_{\tilde{\sigma}'_j \in \Sigma_j^*} \operatorname{supp} \rho_i^{\tilde{\sigma}'_j, \tilde{\sigma}_{-j}, \tilde{\phi}} \supset \operatorname{supp} \rho_i^{\tilde{\sigma},\tilde{\phi}}.$$

■

Thus, let (σ, ϕ) be a canonical NE in game $G^* = (\Gamma, \mathfrak{C}^*)$ such that $\operatorname{supp} \rho_i^{\sigma,\phi} = \bigcup_{j \neq i,0} \bigcup_{\sigma'_j \in \Sigma_j^*} \operatorname{supp} \rho_i^{\sigma'_j, \sigma_{-j}, \phi}$ for all $i \neq 0$. We first construct a (possibly non-canonical) SE in G^* with outcome $\rho^{\sigma,\phi}$. Then we construct a canonical SE with the same outcome.

Non-canonical SE construction: Denote the set of on-path histories for player i by $\hat{H}_i = \{h_i \in H_i : \Pr^{\sigma,\phi}(h_i) > 0\}$. Since (σ, ϕ) is canonical, $h_i \in \hat{H}_i$ if and only if $\mathring{h}_i \in \operatorname{supp} \rho_i^{\sigma,\phi}$ and $r_{i,t} = (a_{i,t-1}, s_{i,t})$ and $m_{i,t} = a_{i,t}$ for all t.

For each k, let σ_i^k denote the perturbation of σ_i where player i trembles uniformly with probability $|R_{i,t}|/k$ at each reporting history $h_i^{R,t} \in H_i^{R,t}$, and trembles uniformly with probability $|A_{i,t}|/k$ at each acting history $h_i^{A,t} \in H_i^{A,t}$.

For each k, let $\left(\Gamma^k, \mathfrak{C}^*\right)$ denote the constrained game where the mediator follows strategy ϕ while each player i is required to play $\sigma_{i,t}^{R,k}\left(h_i^{R,t}\right)$ at each on-path history $h_i^{R,t} \in \hat{H}_i^{R,t}$, is

required to play $\sigma_{i,t}^{A,k}\left(h_i^{A,t}\right)$ at each on-path history $h_i^{A,t} \in \hat{H}_i^{A,t}$, is required to send each report with probability no less than $1/k$ at each off-path history $h_i^{R,t} \in H_i^{R,t} \setminus \hat{H}_i^{R,t}$, and is required to take each action with probability no less than $1/k$ at each off-path history $h_i^{A,t} \in H_i^{A,t} \setminus \hat{H}_i^{A,t}$. By standard arguments, this game admits a NE $\left(\bar{\sigma}^k, \phi\right)$. Taking a convergent subsequence if necessary, let $(\bar{\sigma}, \phi) = \lim_{k\to\infty}\left(\bar{\sigma}^k, \phi\right)$. Clearly, $\rho^{\bar{\sigma},\phi} = \rho^{\sigma,\phi}$.

Let $M_{i,t}\left(h_i^{R,t}, r_{i,t}\right)$ denote the set of messages $m_{i,t}$ that player i receives with positive probability at history $\left(h_i^{R,t}, r_{i,t}\right)$ under profile $\left(\bar{\sigma}^k, \phi\right)$. Since $\bar{\sigma}^k$ has full support, this set depends only on player i's reports and messages $\left(r_i^{t+1}, m_i^t\right)$ at history $\left(h_i^{R,t}, r_{i,t}\right)$. Therefore, we can define a mediation range Q by $Q_{i,t}\left(r_i^{t+1}, m_i^t\right) = M_{i,t}\left(h_i^{R,t}, r_{i,t}\right)$ for all i, t, r_i^{t+1}, m_i^t, and $h_i^{R,t}$ such that $\left(r_i^{t+1}, m_i^t\right)$ equals i's reports and messages at $\left(h_i^{R,t}, r_{i,t}\right)$.

For each k, let ϕ^k denote the perturbation of ϕ where the mediator trembles uniformly with probability $\left|Q_{i,t}\left(r_i^{t+1}, m_i^t\right)\right| / k^{N|Z|}$ over messages $m_{i,t} \in Q_{i,t}\left(r_i^{t+1}, m_i^t\right)$ at each $\left(r_i^{t+1}, m_i^t\right)$, independently across i and t. Define a belief system β as $\lim_{k\to\infty} \Pr^{\bar{\sigma}^k,\phi^k}$. By construction of the relative tremble probabilities for players and the mediator, for each i, t, and $h^{R,t} \in H^{R,t}|_Q$, we have

$$\beta_{i,t}(h^{R,t}|h_i^{R,t}) = \lim_{k\to\infty} \Pr^{\bar{\sigma}^k,\phi^k}(h^{R,t}|h_i^{R,t}) = \lim_{k\to\infty} \Pr^{\bar{\sigma}^k,\phi}(h^{R,t}|h_i^{R,t}),$$

and similarly for $\beta_{i,t}(h^{A,t}|h_i^{A,t})$.

Let J be the set of histories compatible with the mediation range Q, and let K be the set of the mediator's history compatible with the mediation range Q.

We show that $(\bar{\sigma}, \phi, J, K, \beta)$ is a quasi-SE in G^*. The two conditions for validity hold since the messages are within the mediation range as long as the mediator follows ϕ. Consistency of β is by construction. Sequential rationality at off-path histories $h_i \in J_i \setminus \hat{H}_i$ follows from a standard upper hemi-continuity argument. To verify sequential rationality at on-path histories, fix $h_i^{R,t} \in \hat{H}_i^{R,t}$, note that

$$\beta_{i,t}(h^{R,t}|h_i^{R,t}) = \lim_{k\to\infty} \Pr^{\bar{\sigma}^k,\phi}(h^{R,t}|h_i^{R,t}) = \Pr^{\bar{\sigma},\phi}\left(h^{R,t}|h_i^{R,t}\right) = \Pr^{\sigma,\phi}\left(h^{R,t}|h_i^{R,t}\right), \tag{37}$$

where the second equality follows because $\Pr^{\bar{\sigma},\phi}\left(h_i^{R,t}\right) = \Pr^{\sigma,\phi}\left(h_i^{R,t}\right) > 0$. Let $\hat{H} = \{h \in J : \Pr^{\sigma,\phi}(h) > 0\}$. (Note that $\hat{H}$ is not necessarily equal to $\prod_i \hat{H}_i$.) Since $\bigcup_{\sigma_i' \in \Sigma_i} \operatorname{supp} \rho_j^{\sigma_i',\sigma_{-i},\phi} = \operatorname{supp} \rho_j^{\sigma,\phi}$ for each $j \neq i$, we have

$$\Pr^{\sigma_i',\sigma_{-i},\phi}\left(\mathring{h}_j^{T+1} \in \operatorname{supp} \rho_j^{\sigma,\phi}\ \forall j \neq i | h^{R,t}\right) = 1$$

for all $h^{R,t} \in \hat{H}^{R,t}$ and $\sigma_i' \in \Sigma_i^*$.[47] Since (σ, ϕ) is canonical, with probability 1 conditional on $h^{R,t} \in \hat{H}^{R,t}$, $a_{j,\tau} = m_{j,\tau}$ and $r_{j,\tau} = (a_{j,\tau-1}, s_{j,\tau})$ for each $\tau \geq t$. Hence,

$$\Pr^{\sigma_i',\sigma_{-i},\phi}\left(h_j^{T+1} \in \hat{H}_j\ \forall j \neq i | h^{R,t}\right) = 1.$$

[47] Note that $\Pr^{\sigma_i',\sigma_{-i},\phi}$ is well-defined since $(\bar{\sigma}, \phi, J, K)$ is valid.

Finally, since σ_{-i} and $\bar{\sigma}_{-i}$ coincide at all on-path histories, we have

$$\Pr^{\sigma_i', \bar{\sigma}_{-i}, \phi}\left(h_j^{T+1} \in \hat{H}_j \ \forall j \neq i | h^{R,t}\right) = 1. \tag{38}$$

By (37), player i's belief over $h^{R,t} \in J^{R,t}$ at $h_i^{R,t}$ under β is the same as the conditional probability distribution over $h^{R,t} \in J^{R,t}$ at $h_i^{R,t}$ under (σ, ϕ); and for each $h^{R,t}$, by (38), the conditional probability distribution over Z induced by $(\sigma_i', \bar{\sigma}_{-i}, \phi)$ is the same as that induced by $(\sigma_i', \sigma_{-i}, \phi)$. Hence, since (σ, ϕ) is a NE, σ_i is sequentially rational at $h_i^{R,t}$. The argument for acting histories $h_i^{A,t} \in \hat{H}_i^{A,t}$ is analogous.

Canonical SE construction: First, construct a canonical strategy profile $(\tilde{\sigma}, \tilde{\phi})$ from $(\bar{\sigma}, \phi)$ as in the proof of Proposition 2. As in that proof, we let r and m denote the mediator's fictitious reports and messages, and let $\tilde{r}$ and $\tilde{m}$ denote actual reports and messages. Here we also let $\tilde{h}^{R,t} \in \tilde{H}^{R,t}$ and $\tilde{h}^{A,t} \in \tilde{H}^{A,t}$ denote the reporting and acting histories without fictitious reports or messages. Thus, $\left(\tilde{\sigma}, \tilde{\phi}\right)$ is a canonical NE satisfying $\rho^{\tilde{\sigma}, \tilde{\phi}} = \rho^{\sigma, \phi}$. To complete the proof, we construct beliefs $\tilde{\beta}$, subsets of histories $\tilde{J}$ and $\tilde{K}$, and a mediation range $\tilde{Q}$ such that $(\tilde{\sigma}, \tilde{\phi}, \tilde{J}, \tilde{K}, \tilde{\beta})$ is a quasi-SE and, for each i, $\tilde{J}_i$ includes all histories where player i has not lied to the mediator or received a message outside the mediation range.

For each k, let $(\tilde{\sigma}^k, \tilde{\phi}^k)$ denote the perturbation of $(\tilde{\sigma}, \tilde{\phi})$ where (i) players report honestly with probability 1 at all reporting histories $\tilde{h}_i^{R,t}$, (ii) players tremble uniformly over actions with probability $|A_{i,t}|/k$ at all acting histories $\tilde{h}_i^{A,t}$, and (iii) the mediator trembles uniformly with probability $|R_{i,t}|/k$ when she draws fictitious report $r_{i,t}$ at history $(\tilde{r}^{t+1}, r^t, m^t, \tilde{m}^t)$ (but does not tremble when he draws fictitious messages m_t or recommendations $\tilde{m}_t$). By construction, for each s^{T+1}, r^{T+1}, m^{T+1}, and a^{T+1}, we have

$$\Pr^{\tilde{\sigma}^k, \tilde{\phi}^k}\left(s^{T+1}, r^{T+1}, m^{T+1}, a^{T+1}\right) = \Pr^{\bar{\sigma}^k, \phi}\left(s^{T+1}, r^{T+1}, m^{T+1}, a^{T+1}\right). \tag{39}$$

Let $\tilde{M}_{i,t}\left(\tilde{h}_i^{R,t}, \tilde{r}_{i,t}\right)$ denote the set of messages $\tilde{m}_{i,t}$ that player i receives with positive probability at history $\left(\tilde{h}_i^{R,t}, \tilde{r}_{i,t}\right)$ under profile $(\tilde{\sigma}^k, \tilde{\phi}^k)$. Since players $-i$ take all actions with positive probability and the mediator selects each fictitious report with positive probability, this set depends only on player i's reports and messages $\left(\tilde{r}_i^{t+1}, \tilde{m}_i^t\right)$ at $\left(\tilde{h}_i^{R,t}, \tilde{r}_{i,t}\right)$. Therefore, we can define a mediation range $\tilde{Q}$ by $\tilde{Q}_{i,t}\left(\tilde{r}_i^{t+1}, \tilde{m}_i^t\right) = \tilde{M}_{i,t}\left(\tilde{h}_i^{R,t}, \tilde{r}_{i,t}\right)$ for all i, t, $\tilde{r}_i^{t+1}$, $\tilde{m}_i^t$, and $\tilde{h}_i^{R,t}$ such that $\left(\tilde{r}_i^{t+1}, \tilde{m}_i^t\right)$ equals player i's reports and messages at $\left(\tilde{h}_i^{R,t}, \tilde{r}_{i,t}\right)$. By construction, $\Pr^{\tilde{\sigma}^k, \tilde{\phi}^k}\left(\tilde{m}_{i,t} | \tilde{h}_i^{R,t}, \tilde{r}_{i,t}\right) > 0$ if and only if $\tilde{m}_{i,t} \in \tilde{Q}_{i,t}\left(\tilde{r}_i^{t+1}, \tilde{m}_i^t\right)$.

For each i, let $\tilde{J}_i^{A,T+}$ equal the set of histories $\tilde{h}_i^{T+1}$ such that $\Pr^{\tilde{\sigma}^k, \tilde{\phi}^k}\left(\tilde{h}_i^{T+1}\right) > 0$. Let $\tilde{K}^{T+}$ equal the set of mediator histories $\left(\tilde{r}^{T+1}, \tilde{m}^{T+1}\right)$ such that there exists $\hat{\sigma} \in \Sigma$ such that $\Pr^{\hat{\sigma}, \tilde{\phi}^k}\left(\tilde{r}^{T+1}, \tilde{m}^{T+1}\right) > 0$. Let the other elements of $\tilde{J}$ and $\tilde{K}$ include all truncations of histories in $\tilde{J}^{A,T+}$ and $\tilde{K}^{T+}$. Since $\tilde{J}_i$ includes all histories where player i has not lied to the mediator or received a message outside the mediation range, $(\tilde{\sigma}, \tilde{\phi}, \tilde{J}, \tilde{K})$ is valid.

Define a belief system $\tilde{\beta}$ on $\left(\tilde{J}, \tilde{K}\right)$ as $\lim_{k\to\infty} \Pr^{\tilde{\sigma}^k, \tilde{\phi}^k}$. This is well defined because $\Pr^{\tilde{\sigma}^k, \tilde{\phi}^k}(\tilde{h}_i^{R,t}) > 0$ and $\Pr^{\tilde{\sigma}^k, \tilde{\phi}^k}(\tilde{h}_i^{A,t}) > 0$ for all i, t, $\tilde{h}_i^{R,t} \in \tilde{J}_i^{R,t}$, and $\tilde{h}_i^{A,t} \in \tilde{J}_i^{A,t}$. By construction, $\tilde{\beta}$ is consistent, and for each $\tilde{h}_i^{R,t} \in \tilde{J}_i^{R,t}$ and $\tilde{h}^{R,t}$ with $\tilde{\beta}_{i,t}\left(\tilde{h}^{R,t} | \tilde{h}_i^{R,t}\right) > 0$, all players have been truthful at $\tilde{h}^{R,t}$.

It remains to show that $(\tilde{\sigma}, \tilde{\phi}, \tilde{J}, \tilde{K}, \tilde{\beta})$ is sequentially rational. We prove this for reporting histories; the argument for acting histories is analogous. Suppose towards a contradiction that there exist $i \neq 0$, t, $\tilde{h}_i^{R,t} \in \tilde{J}_i^{R,t}$, and a strategy $\tilde{\sigma}_i' \in \Sigma_i^*$ such that

$$\begin{aligned} & \sum_{\tilde{h}^{R,t} \in \tilde{H}^{R,t}[\tilde{h}_i^{R,t}]|_{\tilde{J},\tilde{K}}, r^t, m^t} \tilde{\beta}_{i,t}\left(\tilde{h}^{R,t}, r^t, m^t | \tilde{h}_i^{R,t}\right) \bar{u}_i\left(\tilde{\sigma}_i', \tilde{\sigma}_{-i}, \tilde{\phi} | \tilde{h}^{R,t}, r^t, m^t\right) \\ > & \sum_{\tilde{h}^{R,t} \in \tilde{H}^{R,t}[\tilde{h}_i^{R,t}]|_{\tilde{J},\tilde{K}}, r^t, m^t} \tilde{\beta}_{i,t}\left(\tilde{h}^{R,t}, r^t, m^t | \tilde{h}_i^{R,t}\right) \bar{u}_i\left(\tilde{\sigma}, \tilde{\phi} | \tilde{h}^{R,t}, r^t, m^t\right). \end{aligned}$$

Here, $\tilde{\beta}$ naturally extends to the belief about the profile of the history $\tilde{h}^{R,t}$ and fictitious reports and messages (r^t, m^t). This implies that there exists (r_i^t, m_i^t) with $\tilde{\beta}_{i,t}\left(r_i^t, m_i^t | \tilde{h}_i^{R,t}\right) > 0$ such that

$$\begin{aligned} & \sum_{\tilde{h}^{R,t} \in \tilde{H}^{R,t}[\tilde{h}_i^{R,t}]|_{\tilde{J},\tilde{K}}, r^t, m^t} \tilde{\beta}_{i,t}\left(\tilde{h}^{R,t}, r^t, m^t | \tilde{h}_i^{R,t}, r_i^t, m_i^t\right) \bar{u}_i\left(\tilde{\sigma}_i', \tilde{\sigma}_{-i}, \tilde{\phi} | \tilde{h}^{R,t}, r^t, m^t\right) \\ > & \sum_{\tilde{h}^{R,t} \in \tilde{H}^{R,t}[\tilde{h}_i^{R,t}]|_{\tilde{J},\tilde{K}}, r^t, m^t} \tilde{\beta}_{i,t}\left(\tilde{h}^{R,t}, r^t, m^t | \tilde{h}_i^{R,t}, r_i^t, m_i^t\right) \bar{u}_i\left(\tilde{\sigma}, \tilde{\phi} | \tilde{h}^{R,t}, r^t, m^t\right), \end{aligned}$$

where the summation is taken over (r^t, m^t) whose i-component equals (r_i^t, m_i^t).

Since player i believes that nobody has lied to the mediator, by the same construction as in the proof of Lemma 1, there exists $\sigma_i' \in \Sigma_i^*$ such that, for each $\left(\tilde{h}^{R,t}, r^t, m^t\right) = (s^{t+1}, \tilde{r}^t, r^t, m^t, \tilde{m}^t, a^t)$ with $\tilde{\beta}_{i,t}\left(\tilde{h}^{R,t}, r^t, m^t | \tilde{h}_i^{R,t}, r_i^t, m_i^t\right) > 0$, we have

$$\bar{u}_i\left(\tilde{\sigma}_i', \tilde{\sigma}_{-i}, \tilde{\phi} | \tilde{h}^{R,t}, r^t, m^t\right) = \bar{u}_i\left(\sigma_i', \sigma_{-i}, \phi | s^{t+1}, r^t, m^t, a^t\right).$$

Similarly, by construction of $\left(\tilde{\sigma}, \tilde{\phi}\right)$, we have

$$\bar{u}_i\left(\tilde{\sigma}, \tilde{\phi} | \tilde{h}^{R,t}, r^t, m^t\right) = \bar{u}_i\left(\sigma, \phi | s^{t+1}, r^t, m^t, a^t\right).$$

In total, there exist $i \neq 0$, t, $\tilde{h}_i^{R,t} \in \tilde{J}_i^{R,t}$, (r_i^t, m_i^t) with $\tilde{\beta}_{i,t}\left(r_i^t, m_i^t | \tilde{h}_i^{R,t}\right) > 0$, and a strategy $\sigma_i' \in \Sigma_i^*$ such that

$$\begin{aligned} & \sum_{s^{t+1}, r^t, m^t, a^t} \tilde{\beta}_{i,t}\left(s^{t+1}, r^t, m^t, a^t | \tilde{h}_i^{R,t}, r_i^t, m_i^t\right) \bar{u}_i\left(\sigma_i', \sigma_{-i}, \phi | s^{t+1}, r^t, m^t, a^t\right) \\ > & \sum_{s^{t+1}, r^t, m^t, a^t} \tilde{\beta}_{i,t}\left(s^{t+1}, r^t, m^t, a^t | \tilde{h}_i^{R,t}, r_i^t, m_i^t\right) \bar{u}_i\left(\sigma, \phi | s^{t+1}, r^t, m^t, a^t\right), \end{aligned}$$

where the summation is take over (s^{t+1}, r^t, m^t, a^t) whose i-component corresponds to the counterpart of $\left(\tilde{h}_i^{R,t}, r_i^t, m_i^t\right)$.

Since $(\bar{\sigma}, \phi, J, K, \beta)$ is a quasi-SE in G^*, to derive a contradiction, it remains to show that, for each $i \neq 0$, t, $\tilde{h}_i^{R,t} = (s_i^{t+1}, \tilde{r}_i^t, \tilde{m}_i^t, a_i^t) \in \tilde{J}_i^{R,t}$, (r_i^t, m_i^t) with $\tilde{\beta}_{i,t}\left(r_i^t, m_i^t | \tilde{h}_i^{R,t}\right) > 0$, and (s^{t+1}, r^t, m^t, a^t) whose i-component equals $\left(s_i^{t+1}, r_i^t, m_i^t, a_i^t\right)$, we have

$$\tilde{\beta}_{i,t}\left(s^{t+1}, r^t, m^t, a^t | s_i^{t+1}, \tilde{r}_i^t, r_i^t, m_i^t, \tilde{m}_i^t, a_i^t\right) = \beta_{i,t}\left(s^{t+1}, r^t, m^t, a^t | s_i^{t+1}, r_i^t, m_i^t, a_i^t\right). \tag{40}$$

By construction, given $\tilde{\beta}_{i,t}\left(r_i^t, m_i^t | \tilde{h}_i^{R,t}\right) > 0$, we have $\left(s_i^{t+1}, r_i^t, m_i^t, a_i^t\right) \in J_i^{R,t}$.

To prove (40), note that, for all k we have

$$\begin{aligned} & \Pr^{\tilde{\sigma}^k, \tilde{\phi}^k}\left(s^{t+1}, r^t, m^t, a^t | s_i^{t+1}, \tilde{r}_i^t, r_i^t, m_i^t, \tilde{m}_i^t, a_i^t\right) \\ = & \Pr^{\tilde{\sigma}^k, \tilde{\phi}^k}\left(s^{t+1}, r^t, m^t, a^t | s_i^{t+1}, r_i^t, m_i^t, a_i^t\right) \\ = & \Pr^{\bar{\sigma}^k, \phi}\left(s^{t+1}, r^t, m^t, a^t | s_i^{t+1}, r_i^t, m_i^t, a_i^t\right), \end{aligned}$$

where the first equality follows because (i) given $\tilde{h}_i^{R,t} \in \tilde{J}_i^{R,t}$, we have $\tilde{r}_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$ for all $\tau \leq t$ and (ii) the distribution of $\tilde{m}_i^t$ is fully determined by $(\tilde{r}_i^t, r_i^t, m_i^t)$, and the second equality follows from (39). Therefore,

$$\begin{aligned} & \tilde{\beta}_{i,t}\left(s^{t+1}, r^t, m^t, a^t | s_i^{t+1}, \tilde{r}_i^t, r_i^t, m_i^t, \tilde{m}_i^t, a_i^t\right) \\ = & \lim_{k \to \infty} \Pr^{\tilde{\sigma}^k, \tilde{\phi}^k}\left(s^{t+1}, r^t, m^t, a^t | s_i^{t+1}, \tilde{r}_i^t, r_i^t, m_i^t, \tilde{m}_i^t, a_i^t\right) \\ = & \lim_{k \to \infty} \Pr^{\bar{\sigma}^k, \phi}\left(s^{t+1}, r^t, m^t, a^t | s_i^{t+1}, r_i^t, m_i^t, a_i^t\right) \\ = & \beta_{i,t}\left(s^{t+1}, r^t, m^t, a^t | s_i^{t+1}, r_i^t, m_i^t, a_i^t\right). \end{aligned}$$

H.2 Proof of Claims 2, 3, and 4

Claim 2: Note that the condition $\operatorname{supp} \rho_i^\sigma = \bigcup_{j \neq i,0} \bigcup_{\sigma_j' \in \Sigma_j} \operatorname{supp} \rho_i^{\sigma_j', \sigma_{-j}, \phi}$ is vacuous when $N = 1$. Hence, the result follows from Claim 1.

Claim 3: The only difference from the proof of Claim 1 is in the verification that the non-canonical assessment $(\bar{\sigma}, \phi, J, K, \beta)$ in G^* is sequentially rational at on-path histories $h_i^{R,t} \in \hat{H}_i^{R,t}$ and $h_i^{A,t} \in \hat{H}_i^{A,t}$. There, this followed from equations (37) and (38). Here, note that in game G^* player i faces sequential rationality constraints only in periods $t \geq t_i$.

Sequential rationality at on-path histories in period t_i now follows from (37) and the fact that player i's payoff does not depend on actions taken in periods $\tau \geq t_i$ (other than her own action a_{i,t_i}). The latter fact also implies sequential rationality in periods $t > t_i$.

Claim 4: Claim 4 follows from the fact that players' reports are always canonical in the equilibrium that will be constructed in the proof of Proposition 5.

I Results for CPPBE

This section contains our analysis of CPPBE, culminating in the proofs of Propositions 7 and 8.

I.1 Quasi-Strategies and Quasi-CPPBE

As in Appendix D.1, we begin by introducing notions of "quasi-strategy," which is simply a partially defined strategy, and "quasi-equilibrium,"which is a profile of quasi-strategies where incentive constraints are satisfied wherever strategies are defined.

Fix a game $G = (\Gamma, \mathfrak{C})$. A *quasi-strategy* (χ_i, J_i) for each player i is defined exactly as in Appendix D.1.

Intuitively, a quasi-strategy $(\psi, P, F|_P)$ for the mediator consists of a subset of reports P, a set of mediation plans $F|_P$ that specify messages only after reports in P, and a probability distribution ψ over $F|_P$. Formally, a *quasi-strategy* $(\psi, P, F|_P)$ for the mediator 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 (i) 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, and (ii) for every $r^{T+1} \in P^{T+1}$ and t, the period-t truncation of r^{T+1}, denoted r^{t+1}, satisfies $r^{t+1} \in P^{t+1}$.

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

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

Let $\mathcal{Z}|_{J,P}$ be the set of $\left(f, h^{T+1}\right)$ such that $h^{T+1} = \left(s^{T+1}, r^{T+1}, m^{T+1}, a^{T+1}\right) \in H^{T+1}$, $h_i^{T+1} \in J_i^{T+1}$ for each i, and $f_{i,t}\left(r^{t+1}\right) = m_{i,t}$ for each i and t. For each i and $h_i^{T+1} \in J_i^{T+1}$, let $\mathcal{Z}[h_i^{T+1}]|_{J,P} = \left\{\left(f, h^{T+1}\right) \in \mathcal{Z}|_{J,P} : h^{T+1} \in H^{T+1}[h_i^{T+1}]\right\}$. Define $\mathcal{Z}^{R,t}|_{J,P}$, $\mathcal{Z}^{R,t+}|_{J,P}$, $\mathcal{Z}^{A,t}|_{J,P}$ and $\mathcal{Z}^{R,t+1}|_{J,P}$ as the projections of $\mathcal{Z}|_{J,P}$ to $F \times H^{R,t}$, $F \times H^{R,t} \times R_t$, $F \times H^{A,t}$, and $F \times H^{A,t} \times A_t$, respectively. Define $\mathcal{Z}^{R,t}[h_i^{R,t}]|_{J,P}$ and $\mathcal{Z}^{A,t}[h_i^{A,t}]|_{J,P}$ analogously.

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

1. $J^{R,1} = S_1$. For each $t \geq 1$, $f \in F|_P$, $h^{R,t}$ with $\left(f, h^{R,t}\right) \in \mathcal{Z}^{R,t}|_{J,P}$, $i \neq 0$, σ_i, $\tau \geq t$, and $h^{R,\tau}$ with $\Pr^{\sigma_i, \chi_{-i}, f}\left(h^{R,\tau} | h^{R,t}\right) > 0$, we have $h_j^{R,\tau} \in J_j^{R,\tau}$ for each $j \neq i$ and the report-component r^τ of $h^{R,\tau}$ lies in P^τ.[48] Similarly, for each r_τ with $\Pr^{\sigma_i, \chi_{-i}, f}\left(h^{R,\tau}, r_\tau | h^{R,t}\right) > 0$, we have $(r^\tau, r_\tau) \in P^{\tau+1}$, where r^τ is the report in $h^{R,\tau}$; and for each m_τ with

[48] If $h_j^{A,t-1} \notin J_j^{A,t-1}$ for some $j \neq i$ or $r^t \notin P^t$, then $\Pr^{\sigma_i, \chi_{-i}, f}\left(h^{R,\tau} | h^{R,t}\right)$ is not well-defined. In this case, the condition vacuously holds. The same caution applies to the following conditions.

$\Pr^{\sigma_i,\chi_{-i},f}\left(h^{R,\tau},r_\tau,m_\tau|h^{R,t}\right)>0$, we have $\left(h_j^{R,\tau},r_{j,\tau},m_{j,\tau}\right)\in J_j^{A,\tau}$ for each $j\neq i$. The same condition holds when we replace $h^{R,t}$ with $\left(f,h^{R,t}\right)\in\mathcal{Z}^{R,t}|_{J,P}$ by $h^{A,t}$ with $\left(f,h^{A,t}\right)\in\mathcal{Z}^{A,t}|_{J,P}$. That is, no unilateral player-deviation leads to a history where either the mediator's or another player's quasi-strategy is undefined.

2. For each i and t, if $h_i^{R,t}\in J_i^{R,t}$ then there exist f and $h_{-i}^{R,t}\in J_{-i}^{R,t}$ such that $\left(f,h_i^{R,t},h_{-i}^{R,t}\right)\in\mathcal{Z}^{R,t}|_{J,P}$. Similarly, for each i and t, if $h_i^{A,t}\in J_i^{A,t}$ then there exist f and $h_{-i}^{A,t}\in J_{-i}^{A,t}$ such that $\left(f,h_i^{A,t},h_{-i}^{A,t}\right)\in\mathcal{Z}^{A,t}|_{J,P}$.

The first requirement implies that, for every valid quasi-strategy profile (χ,ψ,J,K), every mediation plan f and terminal history h^{T+1} with $\Pr^{\chi,\psi}\left(f,h^{T+1}\right)>0$ lies in $\mathcal{Z}|_{J,P}$. The second requirement implies that the projection of $\mathcal{Z}|_{J,P}$ on H_i^{T+1} includes all histories $h_i^{T+1}\in J_i^{T+1}$.

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 $\mathcal{Z}|_{J,P}$ such that no player has a profitable deviation at any history in $J_i^{R,t}$ and $J_i^{A,t}$: that is, we have

1. *[CPS consistency]* For all f, t, $h^{R,t}$, r_t, m_t, a_t, and s_{t+1} such that $\left(f,h^{R,t},r_t,m_t,a_t,s_{t+1}\right)\in\mathcal{Z}^{R,t+1}|_{J,P}$, we have

$$\begin{aligned}
\bar{\psi}\left(f\right)&=\psi\left(f\right), & \bar{\psi}\left(r_t|f,h^{R,t}\right)&=\textstyle\prod_{i=0}^N\chi_{i,t}^R\left(r_{i,t}|h_i^{R,t}\right),\\
\bar{\psi}\left(m_t|f,h^{R,t},r_t\right)&=1_{\{m_t=f_t(r^t,r_t)\}}, & \bar{\psi}\left(a_t|f,h^{R,t},r_t,m_t\right)&=\textstyle\prod_{i=0}^N\chi_{i,t}^A\left(a_{i,t}|h_i^{R,t},r_{i,t},m_{i,t}\right),\\
\bar{\psi}\left(s_{t+1}|f,h^{R,t},r_t,a_t\right)&=p\left(s_{t+1}|\mathring{h}^{A,t},a_t\right),
\end{aligned}$$

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

$$\sum_{(f,h^{R,t})\in\mathcal{Z}^{R,t}[h_i^{R,t}]|_{J,P}}\bar{\psi}\left(f,h^{R,t}|h_i^{R,t}\right)\bar{u}_i\left(\chi,f|h^{R,t}\right)\geq\sum_{(f,h^{R,t})\in\mathcal{Z}^{R,t}[h_i^{R,t}]|_{J,P}}\bar{\psi}\left(f,h^{R,t}|h_i^{R,t}\right)\bar{u}_i\left(\sigma_i',\chi_{-i},f|h^{R,t}\right) \tag{41}$$

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

$$\sum_{(f,h^{A,t})\in\mathcal{Z}^{A,t}[h_i^{A,t}]|_{J,P}}\bar{\psi}\left(f,h^{A,t}|h_i^{A,t}\right)\bar{u}_i\left(\chi,f|h^{A,t}\right)\geq\sum_{(f,h^{A,t})\in\mathcal{Z}^{A,t}[h_i^{A,t}]|_{J,P}}\bar{\psi}\left(f,h^{A,t}|h_i^{A,t}\right)\bar{u}_i\left(\sigma_i',\chi_{-i},f|h^{A,t}\right) \tag{42}$$

Let $\rho^{\chi,\psi}\in\Delta\left(X\right)$ 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 8 *For any game G and outcome $\rho\in\Delta\left(X\right)$, ρ is a CPPBE outcome if and only if $\rho=\rho^{\chi,\psi}$ for some quasi-CPPBE profile $\left(\chi,\psi,J,P,F|_P,\bar{\psi}\right)$ in G. Moreover, given a quasi-CPPBE profile $\left(\chi,\psi,J,P,F|_P,\bar{\psi}\right)$, for each Q such that $J\subset Z|_Q$, there exists a CPPBE $(\sigma,\mu,Q,\bar{\mu})$ such that (σ,μ) and (χ,ψ) coincide on (J,P).*

Proof. Fix a game G. If $(\sigma, \mu, Q, \bar{\mu})$ is a CPPBE, then let J be the set of histories compatible with the mediation range: for each t, define

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

Let $P = R$, and let $F|_P$ be the set of mediation plans compatible with the mediation range:

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

Given this definition, we now show that $(\sigma, \mu, J, P, F|_P, \bar{\mu})$ is a quasi-CPPBE. The two defining conditions for validity holds, since (i) $J^{R,1} = S^1$ by definition, (ii) histories outside J_i cannot arise as long as the mediator follows μ, and (iii) every message history in J can arise for some mediation plan in $F|_P$. CPS consistency and sequential rationality follow from the fact that $(\sigma, \mu, Q, \bar{\mu})$ is a CPPBE.

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

By Theorem 1 of Myerson (1986), 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 $\left(\alpha^{F,k}, \alpha^{R,k}, \alpha^{A,k}\right)_k$ with full support on $\mathcal{Z}|_{J,P}$ such that (i) $\alpha^{F,k}(f) \to \psi(f)$ for all $f \in F|_P$, (ii) $\alpha^{R,k}(r_t|f, h^{R,t}) \to \prod_{i=0}^N \chi_{i,t}^R\left(r_{i,t}|h_i^{R,t}\right)$ for all $\left(f, h^{R,t}, r_t\right) \in \mathcal{Z}^{R,t+}|_{J,P}$, and (iii) $\alpha^{A,k}(a_t|f, h^{A,t}) \to \prod_{i=0}^N \chi_{i,t}^A(a_{i,t}|h_i^{A,t})$ for all $\left(f, h^{A,t}, a_t\right) \in \mathcal{Z}^{t+1}|_{J,P}$. For each k, let

$$\varepsilon_k = \min_{t, \left(f, h^{R,t}, r_t, m_t, a_t\right) \in \mathcal{Z}^{t+1}|_{J,P}} \min\{\alpha^{F,k,t}(f), \alpha^{R,k,t}(r_t|f, h^{R,t}), \alpha^{A,k,t}(a_t|f, h^{R,t}, r_t, m_t)\} > 0. \quad (43)$$

Let $\mathfrak{R}_t^k(f, h^{R,t}) = \operatorname{supp} \alpha^{R,k,t}(r_t|f, h^{R,t})$ and $\mathfrak{A}_t^k(f, h^{A,t}) = \operatorname{supp} \alpha^{A,k,t}(a_t|f, h^{R,t}, r_t, m_t)$.

Given $f \in F|_P$, denote the set of mediation plans that coincide with f after history J^t by

$$F(f) = \left\{ \begin{array}{c} f' \in F|_Q : f'_{t-1}(r^t) = f_{t-1}(r^t) \quad \text{for all } t \text{ and } r^t \\ \text{s.t. there exists } h^{R,t} \in J^{R,t} \text{ with report component equal to } r^t \end{array} \right\}.$$

Note that, given $J \subset Z|_Q$, validity implies that the mediator sends messages compatible with mediation range Q for each r^t such that there exists $h^{R,t} \in J^{R,t}$ with report component

equal to r^t. Hence, $F(f)$ is non-empty.

For each k, define an auxiliary game $\left(\Gamma^k, \mathfrak{C}\right)$ as follows:

1. The mediator uses the mixed mediation plan $\mu^k \in \Delta\left(F|_Q\right)$ 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 $f' \in F|_Q$ uniformly at random from $F\left(f\right)$; (ii) with probability $\frac{\varepsilon_k}{k}$, draw $f' \in F|_Q$ uniformly at random from $F|_Q$.

2. Each player i chooses probability distributions $\sigma_{i,t}^{R,k}(\cdot|h_i^{R,t}) \in \Delta(R_{i,t})$ and $\sigma_{i,t}^{A,k}(\cdot|h_i^{A,t}) \in \Delta(A_{i,t})$ for each t, $h_i^{R,t} \in H_i^{R,t} \backslash J_i^{R,t}$, and $h_i^{A,t} \in H_i^{A,t} \backslash J_i^{A,t}$. At histories $h_i^{R,t} \in J_i^{R,t}$ and $h_i^{A,t} \in J_i^{A,t}$, player i is required to choose $\sigma_{i,t}^{R,k}(\cdot|h_i^{R,t}) = \chi_{i,t}^{R,t}\left(\cdot|h_i^{R,t}\right)$ and $\sigma_{i,t}^{A,k}(\cdot|h_i^{A,t}) = \chi_{i,t}^{A,t}\left(\cdot|h_i^{A,t}\right)$.

3. Given σ^k and f, the distribution of terminal histories H^{T+1} is determined recursively as follows:

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

   $$\begin{array}{ll} \left(1-\frac{\varepsilon_k}{k}\left|R_t \backslash \mathfrak{R}_t^k(f,h^{R,t})\right|\right)\alpha^{R,k}(r_t|f,h^{R,t}) & \text{if } \left(f,h^{R,t}\right) \in \mathcal{Z}^{R,t}|_{J,P} \wedge r_t \in \mathfrak{R}_t^k(f,h^{R,t}), \\ \frac{\varepsilon_k}{k} & \text{if } \left(f,h^{R,t}\right) \in \mathcal{Z}^{R,t}|_{J,P} \wedge r_t \notin \mathfrak{R}_t^k(f,h^{R,t}), \\ \prod_{i=0}^N\left(\left(1-\frac{\varepsilon_k}{k}\left|R_{i,t}\right|\right)\sigma_{i,t}^{R,k}(r_{i,t}|h_i^{R,t})+\frac{\varepsilon_k}{k}\right) & \text{if } \left(f,h^{R,t}\right) \notin \mathcal{Z}^{R,t}. \end{array}$$

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

   $$\begin{array}{ll} \left(1-\frac{\varepsilon_k}{k}\left|A_t \backslash \mathfrak{A}_t^k(f,h^{A,t})\right|\right)\alpha^{A,k}(a_t|f,h^{A,t}) & \text{if } \left(f,h^{A,t}\right) \in \mathcal{Z}^{A,t}|_{J,P} \wedge a_t \in \mathfrak{A}_t^k(f,h^{A,t}), \\ \frac{\varepsilon_k}{k} & \text{if } \left(f,h^{A,t}\right) \in \mathcal{Z}^{A,t}|_{J,P} \wedge a_t \notin \mathfrak{A}_t^k(f,h^{A,t}), \\ \prod_{i=0}^N\left(\left(1-\frac{\varepsilon_k}{k}\left|A_{i,t}\right|\right)\sigma_{i,t}^{A,k}(a_{i,t}|h_i^{A,t})+\frac{\varepsilon_k}{k}\right) & \text{if } \left(f,h^{A,t}\right) \notin \mathcal{Z}^{A,t}|_{J,P}. \end{array}$$

   Given $h^{A,t} \in H_t^{A,t}$ and $a_t \in A_t$, 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. Player i's payoff at terminal history h^{T+1} is $u_i(\mathring{h}^{T+1})$.

As in the proof of Lemma 3, $\left(\Gamma^k, \mathfrak{C}\right)$ admits a NE $(\bar{\sigma}^k, \mu^k)$. Moreover, for any σ^k, $\left(\sigma^k, \mu^k\right)$ has full support on $F|_Q \times Z|_Q$ in $\left(\Gamma^k, \mathfrak{C}\right)$. Hence, $(\bar{\sigma}^k, \mu^k)$ induces a CPS $\bar{\mu}^k$ on $F|_Q \times Z|_Q$ by Bayes' rule.

Let $(\bar{\sigma}^k, \mu^k, \bar{\mu}^k)_k$ denote a sequence of NE $(\bar{\sigma}^k, \mu^k)$ and corresponding CPS's $\bar{\mu}^k$ in $\left(\Gamma^k, \mathfrak{C}\right)$. Taking a convergent subsequence if necessary, let $(\bar{\sigma}, \mu, \bar{\mu}) = \lim_{k\to\infty}(\bar{\sigma}^k, \mu^k, \bar{\mu}^k)$. Note that $(\bar{\sigma}, \mu)$ and (χ, ψ) coincide on (J, P). We claim that $(\bar{\sigma}, \mu, Q, \bar{\mu})$ is a CPPBE in $(\Gamma, \mathfrak{C})$. Since $\bar{\mu}$ is a CPS as the limit of conditional probabilities, it remains to verify sequential rationality. The proof is exactly parallel to the corresponding part of the proof of Lemma 3. We include it for completeness.

We consider reporting histories $h_i^{R,t}$; the argument for acting histories $h_i^{A,t}$ is analogous. 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 the outcome distribution (and hence player

i's expected payoff) conditional on $h^{R,t}$ is continuous in σ^k, μ^k, ε_k, and k. Since $\bar{\sigma}^k_{i,t}\left(\cdot|h_i^{R,t}\right)$ is sequentially rational in $\left(\Gamma^k,\mathfrak{C}\right)$ (as $\left(\bar{\sigma}^k,\mu^k\right)$ is a NE in $\left(\Gamma^k,\mathfrak{C}\right)$ where the distribution over h^{T+1} has full support), it follows that $\sigma^R_{i,t}\left(\cdot|h_i^{R,t}\right)$ is sequentially rational in $(\Gamma,\mathfrak{C})$.

Now consider the case where $h_i^{R,t}\in J_i^{R,t}$. We show that player i believes that $h^{R,t}\in\mathcal{Z}^{R,t}[h_i^{R,t}]|_{J,P}$ with probability 1. Note that, for each $h_i^{T+1}\in J_i^{T+1}$ and $\left(f,h^{T+1}\right)\notin\mathcal{Z}^{R,t}[h_i^{R,t}]|_{J,P}$, there exists $\left(\tilde{f},\tilde{h}^{T+1}\right)\in\mathcal{Z}^{R,t}[h_i^{R,t}]|_{J,P}$ such that

$$\lim_{k\to\infty}\frac{\bar{\mu}^k(f,h^{T+1})}{\bar{\mu}^k(\tilde{f},\tilde{h}^{T+1})}=0.$$

This follows because in $\left(\Gamma^k,\mathfrak{C}\right)$ 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 third 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 $h_i^{R,t}\in J_i^{R,t}$ and $\left(f,h^{R,t}\right)\in\mathcal{Z}^{R,t}[h_i^{R,t}]|_{J,P}$, we have $\bar{\mu}(f,h^{R,t}|h_i^{R,t})=\bar{\psi}(f,h^{R,t}|h_i^{R,t})$, and the conditional probability that $\left(f,h^{R,t}\right)\in\mathcal{Z}^{R,t}[h_i^{R,t}]|_{J,P}$ equals 1. Hence, the fact that (41) holds with CPS $\bar{\psi}$ implies that $\sigma^R_{i,t}(\cdot|h_i^{R,t})=\chi_{i,t}\left(\cdot|h_i^{R,t}\right)$ is sequentially rational in $(\Gamma,\mathfrak{C})$. ■

I.2 SCE Implies CPPBE

Lemma 9 *For any base game Γ, mediation range Q, and SCE $(\mu,Q,\bar{\mu})$, there exists a canonical strategy profile σ and CPS $\bar{\mu}'$ such that $(\sigma,\mu,Q,\bar{\mu}')$ is a CPPBE in $(\Gamma,\mathfrak{C}^*)$ with the same outcome distribution.*

Proof. In the direct-communication game $G^*=(\Gamma,\mathfrak{C}^*)$, let J be the set of truthful histories compatible with the mediation range: for each t, define

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

Let $P=R$, and let $F|_P$ be the set of mediation plans compatible with the mediation range:

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

Consider the quasi-strategy profile $(\chi, \mu, J, P, F|_P)$ where, for each i, χ_i is honest and obedient at each $h_i \in J_i$. 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 mediation plan outside $F|_P$, and (ii) every message history in J can arise for some mediation plan in $F|_P$. Moreover, by inspection, $(\chi, \mu, J, P, F|_P, \bar{\mu})$ is a quasi-CPPBE in G^* if $(\mu, Q, \bar{\mu})$ is a SCE. Hence, the former is a quasi-CPPBE in G^*. Moreover, J_i includes all histories at which player i has been honest and the mediator's messages lie in the mediation range. Hence, Lemma 8 implies that there exists a canonical CPPBE $(\sigma, \mu, Q, \bar{\mu}')$ in G^* with the same mediation range and outcome as $(\mu, Q, \bar{\mu})$. ■

I.3 CPPBE and Codominated Actions

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

Lemma 10 *For any game G, mediation range Q, and CPPBE $(\sigma, \mu, Q, \bar{\mu})$, $\operatorname{supp} \sigma_{i,t}^A(h_i^{A,t}) \cap \mathfrak{D}_{i,t}(\mathring{h}_i^{A,t}) = \emptyset$ for all i, t, and $h_i^{A,t} \in H_i^{A,t}$.*

I.3.1 Proof of Lemma 10

Fix a game G, mediation range Q, CPPBE $(\sigma, \mu, Q, \bar{\mu})$, and sequence of full-support CPS's $\left(\bar{\mu}^k\right)_k$ converging to $\bar{\mu}$. For each i and t, the sequential rationality condition at history $h_i^{A,t}$ is

$$\sum_{(f,h^{A,t}) \in \mathcal{Z}^{A,t}[h_i^{A,t}]|_Q} \bar{\mu}(f, h^{A,t}|h_i^{A,t}) \bar{u}_i\left(\sigma, f|h^{A,t}\right) = \max_{\sigma_i' \in \Sigma_i} \sum_{(f,h^{A,t}) \in \mathcal{Z}^{A,t}[h_i^{A,t}]|_Q} \bar{\mu}(f, h^{A,t}|h_i^{A,t}) \bar{u}_i\left(\sigma_i', \sigma_{-i}, f|h^{A,t}\right). \tag{44}$$

We wish to prove the following lemma, which establishes the corresponding sequential rationality condition in the direct-communication game G^*. Let $\tilde{F}$ be the set of mediation plans in game G^*. For each i, t, and $y_i^t \in Y_i^t$, let $\tilde{M}_{i,t}\left(y_i^t\right) = \bigcup_{h_i^{A,t}: \mathring{h}_i^{A,t} = y_i^t} \operatorname{supp} \sigma_{i,t}^A(h_i^{A,t})$; and, for each $\tilde{r}_i^{t+1} \in R_i^{*t+1}$, let

$$\tilde{Q}_{i,t}(\tilde{r}_i^{t+1}) = \begin{cases} \tilde{M}_{i,t}\left(\tilde{r}_i^{t+1}\right) & \text{if } \tilde{r}_i^{t+1} \in Y_i^t \\ A_{i,t} & \text{otherwise} \end{cases}. \tag{45}$$

Given $\tilde{Q}$, let $\tilde{F}|_{\tilde{Q}}$ be the set of mediation plans in game G^* with mediation range $\tilde{Q}$.

Lemma 11 *In game G^*, for each t, there exists a CPS $\tilde{\mu}_t$ on $\tilde{F}|_{\tilde{Q}} \times Y^t$ such that, for each i, $y_i^t \in Y_i^t$, $\tilde{m}_{i,t} \in \tilde{M}_{i,t}\left(y_i^t\right)$, and $\sigma_i' \in \Sigma_i^*$,*

$$\begin{aligned} &\sum_{\left(\tilde{f}, y^t\right) \in \tilde{F}|_{\tilde{Q}} \times Y^t[y_i^t]} \tilde{\mu}_t(\tilde{f}, y^t|y_i^t, \tilde{m}_{i,t}) \bar{u}_i\left(\sigma^*, \tilde{f}|\bar{h}\left(\tilde{f}, y^t\right), \tilde{m}_{i,t}\right) \\ \geq &\sum_{\left(\tilde{f}, y^t\right) \in \tilde{F}|_{\tilde{Q}} \times Y^t[y_i^t]} \tilde{\mu}_t(\tilde{f}, y^t|y_i^t, \tilde{m}_{i,t}) \bar{u}_i\left(\sigma_i', \sigma_{-i}^*, \tilde{f}|\bar{h}\left(\tilde{f}, y^t\right), \tilde{m}_{i,t}\right). \end{aligned} \tag{46}$$

Before proving Lemma 11, we first show how it implies Lemma 10. Toward a contradiction, suppose that there exists a period t such that $\tilde{M}_{i,t}\left(y_i^t\right) \cap \mathfrak{D}_{i,t}(y_i^t) \neq \emptyset$ for some i and $y_i^t \in Y_i^t$. Let t^* be the last such period. Note that, for all $\tilde{f} \in \tilde{F}|_{\tilde{Q}}$, we have $\tilde{f}_{i,t}\left(\tilde{r}^{t+1}\right) \cap \mathfrak{D}_{i,t}(\tilde{r}_i^{t+1}) = \emptyset$ for all $t > t^*$, i, and $\tilde{r}^{t+1}$ such that $\tilde{r}_i^{t+1} \in Y_i^t$: that is, recommendations after period t^* exclude codominated actions.

Let $E \subset \tilde{F}|_{\tilde{Q}} \times Y^{t^*}$ denote the set of pairs $\left(\tilde{f}, y^{t^*}\right)$ such that $\tilde{f}_{i,t^*}(y^{t^*}) \in \mathfrak{D}_{i,t^*}(y_i^{t^*})$ for some i. Define $\tilde{\mu}_{t^*}^E\left(\tilde{f}, y^{t^*}\right) = \tilde{\mu}_{t^*}(\tilde{f}, y^{t^*}|E)$. For each $\tilde{m}_{i,t^*} \in \mathfrak{D}_{i,t^*}(y_i^{t^*})$, the conditioning event $\tilde{m}_{i,t^*}$ in (46) implies that the realized pair $\left(\tilde{f}, y^{t^*}\right)$ lies in E. Hence, (46) implies that, for each i and $\tilde{m}_{i,t^*} \in \tilde{M}_{i,t^*}\left(y_i^{t^*}\right) \cap \mathfrak{D}_{i,t^*}(y_i^{t^*})$, we have

$$\sum_{\substack{\left(\tilde{f},y^{t^*}\right)\in\tilde{F}|_{\tilde{Q}}\times Y^{t^*}:\\ \tilde{f}_{i,t^*}(y^{t^*})=\tilde{m}_{i,t^*}}} \tilde{\mu}_{t^*}^E(\tilde{f}, y^{t^*})\bar{u}_i\left(\sigma^*, \tilde{f}|\bar{h}\left(\tilde{f}, y^{t^*}\right)\right) \geq \sum_{\substack{\left(\tilde{f},y^{t^*}\right)\in\tilde{F}|_{\tilde{Q}}\times Y^{t^*}:\\ \tilde{f}_{i,t^*}(y^{t^*})=\tilde{m}_{i,t^*}}} \tilde{\mu}_{t^*}^E(\tilde{f}, y^{t^*})\bar{u}_i\left(\sigma_i', \sigma_{-i}^*, \tilde{f}|\bar{h}\left(\tilde{f}, y^{t^*}\right)\right).$$

This contradicts the hypothesis that $\tilde{m}_{i,t^*} \in \mathfrak{D}_{i,t^*}(y_i^{t^*})$, which completes the proof of Lemma 10.

We now prove Lemma 11. Let $\phi^{\geq t}$ denote a collection of functions $(\phi_\tau)_{\tau \geq t}$, where $\phi_t : R^{*t+1} \to \Delta\left(M_t^* \times \mathcal{Z}^{A,t}|_Q\right)$ (where here $\mathcal{Z}^{A,t}|_Q \subset F|_Q \times H^{A,t}$ denotes the set of mediation plans and period-t acting histories in G with mediation range Q), and for $\tau > t$, $\phi_\tau : \mathcal{Z}^{A,\tau-1}|_Q \times R_\tau^* \to \Delta\left(M_\tau^* \times H^{A,\tau}\right)$. Intuitively, $\phi^{\geq t}$ may be viewed as a continuation strategy for the mediator in the direct-communication game starting with arbitrary past reports $\tilde{r}^{t+1} \in R^{*t+1}$ in period t.

Lemma 12 *For each t, in game G^*, there exists $\left(\left(\phi_\tau^k\right)_{\tau \geq t}\right)_k$ with limit $\phi_\tau = \lim_{k\to\infty} \phi_\tau^k$ for each $\tau \geq t$ such that, for each i, $y_i^t \in Y_i^t$, and $\tilde{m}_{i,t} \in \tilde{M}_{i,t}\left(y_i^t\right)$, we have*

$$\sum_{y^t \in Y^t[y_i^t]} \bar{\mu}^k(y^t|y_i^t)\phi_t^k\left(\tilde{m}_{i,t}|y^t\right) > 0 \text{ for all } k, \tag{47}$$

and, for all $\sigma_i' \in \Sigma_i^*$,

$$\begin{aligned} &\sum_{y^t \in Y^t[y_i^t]} \bar{\mu}(y^t|y_i^t, \tilde{m}_{i,t})\bar{u}_i\left(\sigma^*, (\phi_\tau)_{\tau\geq t}|y^t, \left\{\tilde{r}^{t+1} = y^t\right\}, \tilde{m}_{i,t}\right) \\ \geq &\sum_{y^t \in Y^t[y_i^t]} \bar{\mu}(y^t|y_i^t, \tilde{m}_{i,t})\bar{u}_i\left(\sigma_i', \sigma_{-i}^*, (\phi_\tau)_{\tau\geq t}|y^t, \left\{\tilde{r}^{t+1} = y^t\right\}, \tilde{m}_{i,t}\right), \end{aligned} \tag{48}$$

where $\bar{\mu}$ is defined by

$$\bar{\mu}(y^t|y_i^t, \tilde{m}_{i,t}) = \lim_{k\to\infty} \frac{\bar{\mu}^k(y^t|y_i^t)\phi_t^k(\tilde{m}_{i,t}|y^t)}{\sum_{\tilde{y}^t \in Y^t[y_i^t]} \bar{\mu}^k(\tilde{y}^t|y_i^t)\phi_t^k(\tilde{m}_{i,t}|\tilde{y}^t)}.$$

Proof. Construction of $\left(\phi_\tau^k\right)_{\tau\geq t}$: This is similar to the proof of Proposition 2. For each k, first define $\left(\phi_\tau^k\right)_{\tau\geq t}$ recursively in τ, then define $\phi_\tau = \lim_{k\to\infty} \phi_\tau^k$ for all $\tau \geq t$.

For each canonical $\tilde{r}^{t+1} \in R^{*t+1}$, the mediator draws a mediation plan and "fictitious history" $\left(f, h^{A,t}\right) \in \mathcal{Z}^{A,t}|_Q$ according to $\bar{\mu}^k(f, h^{A,t}|y^t)$ for $y^t = \tilde{r}^{t+1}$.[49] Then, he recommends $\tilde{m}_{i,t} \in A_{i,t}$ to player i according to $\sigma^A_{i,t}(h^{A,t})$. This defines ϕ^k_t.

For each $\tau > t$, we now define ϕ^k_τ as a function of $\left(f, h^{A,\tau-1}, a_{i,\tau-1}, s_{i,\tau}\right)$ with $(a_{i,\tau-1}, s_{i,\tau}) = \tilde{r}_{i,\tau}$. For each i, the mediator draws a "fictitious report" $r_{i,\tau} \in R_{i,\tau}$ according to $\sigma^R_{i,\tau}(h_i^{A,\tau-1}, \tilde{r}_{i,\tau})$, independently across players. Next, given r_τ, the mediator calculates the vector of "fictitious messages" $m_\tau = f_\tau(r^\tau, r_\tau)$, where r^τ is the report component of $h^{A,\tau-1}$. Finally, the mediator draws recommendation $\tilde{m}_{i,\tau} \in A_{i,\tau}$ according to $\sigma^A_{i,\tau}(h_i^{A,\tau-1}, a_{i,\tau-1}, s_{i,\tau}, r_{i,\tau}, m_{i,\tau})$ with $(a_{i,\tau-1}, s_{i,\tau}) = \tilde{r}_{i,\tau}$, independently across players. This then defines $h_i^{A,\tau} = (h_i^{A,\tau-1}, a_{i,\tau-1}, s_{i,\tau}, r_{i,\tau}, m_{i,\tau})$ with $(a_{i,\tau-1}, s_{i,\tau}) = \tilde{r}_{i,\tau}$.

Proof of (47): For each $y^t \in Y^t$, since $\bar{\mu}^k(f, h^{A,t}|y^t)$ has full support over $\left(f, h^{A,t}\right) \in \mathcal{Z}^{A,t}\left[y^t\right]|_Q$, we have $\phi^k_t(\tilde{m}_{i,t}|y^t) > 0$ for each i and $\tilde{m}_{i,t} \in \tilde{M}_{i,t}\left(y^t_i\right)$. Hence, (47) is satisfied.

Proof of (48): Toward a contradiction, suppose (48) is violated for some i, $y^t_i \in Y^t_i$, $\tilde{m}_{i,t} \in \tilde{M}_{i,t}\left(y^t_i\right)$, and $\sigma'_i \in \Sigma^*_i$. Denote the conditional probability of $\left(f, h^{A,t}\right) \in \mathcal{Z}^{A,t}|_Q$ and $x \in X$ given y^t, $\tilde{r}^{t+1} = y^t$, and $\tilde{m}_{i,t}$ by

$$\Pr^{\sigma^*,(\phi_\tau)_{\tau\geq t}}\left(f, h^{A,t}, x|y^t, \left\{\tilde{r}^{t+1} = y^t\right\}, \tilde{m}_{i,t}\right).$$

By construction, for each $\tilde{\sigma} \in \Sigma^*$,

$$\begin{aligned}
& \Pr^{\tilde{\sigma},(\phi_\tau)_{\tau\geq t}}\left(f, h^{A,t}, x|y^t, \left\{\tilde{r}^{t+1} = y^t\right\}, \tilde{m}_{i,t}\right) \\
= & \lim_{k\to\infty} \Pr^{\phi^k_t}\left(f, h^{A,t}|y^t, \left\{\tilde{r}^{t+1} = y^t\right\}, \tilde{m}_{i,t}\right) \Pr^{\tilde{\sigma},(\phi_\tau)_{\tau\geq t}}\left(x|y^t, \left\{\tilde{r}^{t+1} = y^t\right\}, f, h^{A,t}, \tilde{m}_{i,t}\right) \\
= & \bar{\mu}(f, h^{A,t}|y^t, \tilde{m}_{i,t}) \Pr^{\tilde{\sigma},(\phi_\tau)_{\tau\geq t}}\left(x|y^t, \left\{\tilde{r}^{t+1} = y^t\right\}, f, h^{A,t}, \tilde{m}_{i,t}\right) \\
= & \bar{\mu}(f, h^{A,t}|y^t, \tilde{m}_{i,t}) \Pr^{\tilde{\sigma},(\phi_\tau)_{\tau\geq t}}\left(x|f, h^{A,t}, \tilde{m}_{i,t}\right).
\end{aligned}$$

In the last line, we omit y^t since $\Pr^{\phi^k_t}(\cdot|y^t, \tilde{m}_{i,t})$ assigns probability 1 to $\mathring{h}^{A,t} = y^t$, and we also omit $\{\tilde{r}^{t+1} = y^t\}$ by defining $\Pr^{\tilde{\sigma},(\phi_\tau)_{\tau\geq t}}\left(x|f, h^{A,t}, \tilde{m}_{i,t}\right)$ as follows: conditional on $h^{A,t}$, define $\tilde{r}^{t+1} = \mathring{h}^{A,t}$ and calculate the conditional distribution of x given $\tilde{\sigma}$ and $\left(f, h^{A,t}, \tilde{r}^{t+1}, \tilde{m}_{i,t}\right)$. By definition, for each $y^t \in Y^t[y^t_i]$,

$$\bar{\mu}(y^t|y^t_i, \tilde{m}_{i,t})\bar{\mu}(f, h^{A,t}|y^t, \tilde{m}_{i,t}) = \bar{\mu}(f, h^{A,t}|y^t_i, \tilde{m}_{i,t}) 1_{\left\{\mathring{h}^{A,t} = y^t\right\}}.$$

Hence, the violation of (48) implies

$$\begin{aligned}
& \sum_{(f,h^{A,t})\in\mathcal{Z}^{A,t}|_Q : \mathring{h}^{A,t}\in Y^t[y^t_i]} \bar{\mu}(f, h^{A,t}|y^t_i, \tilde{m}_{i,t}) \bar{u}_i\left(\sigma^*, (\phi_\tau)_{\tau\geq t}|f, h^{A,t}, \tilde{m}_{i,t}\right) \\
< & \sum_{(f,h^{A,t})\in\mathcal{Z}^{A,t}|_Q : \mathring{h}^{A,t}\in Y^t[y^t_i]} \bar{\mu}(f, h^{A,t}|y^t_i, \tilde{m}_{i,t}) \bar{u}_i\left(\sigma'_i, \sigma^*_{-i}, (\phi_\tau)_{\tau\geq t}|f, h^{A,t}, \tilde{m}_{i,t}\right).
\end{aligned}$$

[49] As in the proof of Proposition 2, $\left(f, h^{A,t}\right)$ can be chosen arbitrarily if $\tilde{r}^{t+1}$ is not a feasible payoff-relevant history. A similar comment applies to $r_{i,\tau}$ and $h_i^{A,\tau}$ in the next paragraph.

Therefore, there must exist $h_i^{A,t}$ with $\bar{\mu}(h_i^{A,t}|y_i^t, \tilde{m}_{i,t}) > 0$ such that

$$\begin{aligned} & \sum_{(f,h^{A,t})\in\mathcal{Z}^{A,t}[h_i^{A,t}]|_Q} \bar{\mu}(f, h^{A,t}|h_i^{A,t}, \tilde{m}_{i,t})\bar{u}_i\left(\sigma^*, (\phi_\tau)_{\tau\geq t}\,|f, h^{A,t}, \tilde{m}_{i,t}\right) \\ < & \sum_{(f,h^{A,t})\in\mathcal{Z}^{A,t}[h_i^{A,t}]|_Q} \bar{\mu}(f, h^{A,t}|h_i^{A,t}, \tilde{m}_{i,t})\bar{u}_i\left(\sigma_i', \sigma_{-i}^*, (\phi_\tau)_{\tau\geq t}\,|f, h^{A,t}, \tilde{m}_{i,t}\right). \end{aligned} \tag{49}$$

Note that $(\phi_\tau)_{\tau\geq t}$ is constructed so that the conditional distribution of x given $\left(f, h^{A,t}\right)$ and $\tilde{m}_{i,t}$ under $\left(\sigma^*, (\phi_\tau)_{\tau\geq t}\right)$ in game G^* is the same as the conditional distribution of x given $\left(f, h^{A,t}\right)$ and $a_{i,t} = \tilde{m}_{i,t} \in \operatorname{supp}\sigma_{i,t}^A(h_i^{A,t})$ in game G with mediation range Q under (σ, f):

$$\begin{aligned} & \sum_{(f,h^{A,t})\in\mathcal{Z}^{A,t}[h_i^{A,t}]|_Q} \bar{\mu}(f, h^{A,t}, \tilde{m}_t|h_i^{A,t}, \tilde{m}_{i,t})\bar{u}_i\left(\sigma^*, (\phi_\tau)_{\tau\geq t}\,|f, h^{A,t}, \tilde{m}_{i,t}\right) \\ = & \sum_{(f,h^{A,t})\in\mathcal{Z}^{A,t}[h_i^{A,t}]|_Q} \bar{\mu}(f, h^{A,t}, \tilde{m}_t|h_i^{A,t}, \tilde{m}_{i,t})\bar{u}_i\left(\sigma|f, h^{A,t}, \tilde{m}_{i,t}\right). \end{aligned}$$

To derive a contradiction, it suffices to find a strategy $\hat{\sigma}_i'$ in G with mediation range Q that attains the expected payoff in the second line of (49),

$$\begin{aligned} & \sum_{(f,h^{A,t})\in\mathcal{Z}^{A,t}[h_i^{A,t}]|_Q} \bar{\mu}(f, h^{A,t}|h_i^{A,t}, \tilde{m}_{i,t})\bar{u}_i\left(\sigma_i', \sigma_{-i}^*, (\phi_\tau)_{\tau\geq t}\,|f, h^{A,t}, \tilde{m}_{i,t}\right) \\ = & \sum_{(f,h^{A,t})\in\mathcal{Z}^{A,t}[h_i^{A,t}]|_Q} \bar{\mu}(f, h^{A,t}|h_i^{A,t}, \tilde{m}_{i,t})\bar{u}_i\left(\hat{\sigma}_i', \sigma_{-i}|f, h^{A,t}, \tilde{m}_{i,t}\right), \end{aligned}$$

since the existence of such a strategy contradicts (44). The same construction as in the proof of Lemma 1 defines such a strategy $\hat{\sigma}_i'$. ■

Proof of Lemma 11. We now prove (46). For each k, by Kuhn's theorem, there exists a collection of mixed mediation plan $\left(\tilde{\mu}_{\tilde{r}^{t+1}}^k\right)_{\tilde{r}^{t+1}}$ in G^*, with $\tilde{\mu}_{\tilde{r}^{t+1}}^k \in \Delta\left(\tilde{F}\right)$ for each $\tilde{r}^{t+1}$, that satisfies the following condition: for each $y^t \in Y^t$, $\tilde{r}^{t+1} \in R^{*t+1}$, strategy $\sigma' \in \Sigma^*$, and vector $(\tilde{m}_t, a_t, (s_\tau, \tilde{r}_\tau, \tilde{m}_\tau, a_\tau)_{\tau=t+1}^T)$, we have

$$\begin{aligned} & \Pr^{\sigma', \left(\phi_\tau^k\right)_{\tau\geq t}}\left(\tilde{m}_t, a_t, (s_\tau, \tilde{r}_\tau, \tilde{m}_\tau, a_\tau)_{\tau=t+1}^T\,|y^t, \tilde{r}^{t+1}\right) \\ = & \sum_{\tilde{f}\in\Delta(\tilde{F})} \tilde{\mu}_{\tilde{r}^{t+1}}^k(\tilde{f})\Pr^{\sigma'}\left(\tilde{m}_t, a_t, (s_\tau, \tilde{r}_\tau, \tilde{m}_\tau, a_\tau)_{\tau=t+1}^T\,|f, y^t, \tilde{r}^{t+1}\right). \end{aligned}$$

(That is, $\left(\sigma', \left(\phi_\tau^k\right)_{\tau\geq t}\right)$ and $\left(\sigma', \tilde{\mu}_{\tilde{r}^{t+1}}^k\right)$ give rise to the same distribution of histories in G^*

conditional on y^t and $\tilde{r}^{t+1}$.) In particular, if $\tilde{r}^{t+1} = y^t$, we have

$$\begin{aligned}
&\Pr^{\sigma', (\phi^k_\tau)_{\tau \geq t}} \left(\tilde{m}_t, a_t, (s_\tau, \tilde{r}_\tau, \tilde{m}_\tau, a_\tau)_{\tau=t+1}^T \,|y^t \right) \\
= &\sum_{\tilde{f} \in \Delta(\tilde{F})} \tilde{\mu}^k_{y^t}(\tilde{f}) \Pr^{\sigma'} \left(\tilde{m}_t, a_t, (s_\tau, \tilde{r}_\tau, \tilde{m}_\tau, a_\tau)_{\tau=t+1}^T \,|\tilde{f}, y^t \right).
\end{aligned}$$

Define

$$\tilde{\mu}^k(y^t, \tilde{f}, \tilde{m}_t) = \bar{\mu}^k(y^t) \times \tilde{\mu}^k_{y^t}(\tilde{f}) \times 1_{\left\{\tilde{m}_t = \tilde{f}(y^t)\right\}}.$$

Let $\tilde{\mu} = \lim_{k \to \infty} \tilde{\mu}^k$. Since each $\tilde{f} \in \operatorname{supp} \tilde{\mu}$ satisfies $\tilde{f}_{i,\tau}(\tilde{r}^{\tau+1}) \in \tilde{Q}_{i,\tau}(\tilde{r}^{t+1}_\tau)$ for all $\tau \geq t$, (48) implies

$$\begin{aligned}
&\sum_{\left(\tilde{f}, y^t\right) \in \tilde{F}|_{\tilde{Q}} \times Y^t[y^t_i]} \tilde{\mu}_t(\tilde{f}, y^t | y^t_i, \tilde{m}_{i,t}) \bar{u}_i \left(\sigma^*, \tilde{f} | y^t, \left\{ \tilde{r}^{t+1} = y^t \right\}, \tilde{m}_{i,t} \right) \\
\geq &\sum_{\left(\tilde{f}, y^t\right) \in \tilde{F}|_{\tilde{Q}} \times Y^t[y^t_i]} \tilde{\mu}_t(\tilde{f}, y^t | y^t_i, \tilde{m}_{i,t}) \bar{u}_i \left(\sigma'_i, \sigma^*_{-i}, \tilde{f} | y^t, \left\{ \tilde{r}^{t+1} = y^t \right\}, \tilde{m}_{i,t} \right). \qquad (50)
\end{aligned}$$

For each $\tilde{f}$, we can write $\tilde{f} = (\tilde{f}^{<t}, \tilde{f}^{\geq t})$ with $\tilde{f}^{<t} = (\tilde{f}_\tau)_{\tau=1}^{t-1}$ and $\tilde{f}^{\geq t} = (\tilde{f}_\tau)_{\tau=t}^T$. Since the past recommendations $\tilde{m}^t$ do not affect the continuation strategy, there exists $\hat{\mu}_t$ such that $\hat{\mu}_t(\tilde{f}) = \hat{\mu}_t^{<t}(\tilde{f}^{<t}) \times \hat{\mu}_t^{\geq t}(\tilde{f}^{\geq t})$ and (50) holds with $\hat{\mu}_t$ in place of $\tilde{\mu}_t$. Since under $\hat{\mu}_t$ recommendations prior to period t are independent of those after period t, this yields (46). ■

I.4 Proof of Propositions 7 and 8

Proposition 7: By Lemma 9, for each SCE $(\mu, Q, \bar{\mu})$, there exists a canonical CPPBE $(\sigma, \mu, Q, \bar{\mu}')$ in G^* with $\rho^{\sigma,\mu} = \rho^{\sigma^*,\mu}$. Conversely, take a CPPBE $(\sigma, \mu, Q, \bar{\mu})$ in G^* with canonical σ and outcome ρ. As in the proof of Lemma 9, let J be the set of histories such that players are honest and messages are compatible with mediation range Q, let $P = R$, and let $F|_P$ be the set of mediation plans compatible with mediation range Q. Since σ is canonical, the quasi-CPPBE $(\sigma, \mu, J, P, F|_P)$ is SCE.

Proposition 8: By Lemma 10, players do not take codominated actions at any history for any $\mathfrak{C}$. Since every CPPBE $(\sigma, \mu, Q, \bar{\mu})$ is a NE, there exists a NE with outcome ρ where players do not take codominated actions at any history. Hence, by Proposition 1, there exists a SCE $(\mu', Q', \bar{\mu}')$ with $\rho = \rho^{\sigma^*,\mu}$ and $Q_{i,t}\left(r_i^{t+1}, m_i^t\right) = A_{i,t} \setminus \mathfrak{D}_{i,t}(r_i^{t+1})$. Hence, by Lemma 9, there exists a canonical CPPBE $(\sigma, \mu', Q', \bar{\mu}'')$ with outcome ρ.

J Proof of Claim 5 of Proposition 4

We first establish a preliminary result. Fix a game G, mediation range Q, and CPPBE $(\sigma, \mu, Q, \bar{\mu})$. For each i, let $\Sigma_i(\sigma_{-i}, \mu, \bar{\mu}) \subset \Sigma_i$ denote the set of sequentially rational strategies for player i against (σ_{-i}, μ) under CPS $\bar{\mu}$: that is, the set of strategies $\hat{\sigma}_i$ such that, for each

t and $h_i^{R,t}$,

$$\sum_{(f,h^{R,t})\in\mathcal{Z}^{R,t}[h_i^{R,t}]|_Q}\bar{\mu}(f,h^{R,t}|h_i^{R,t})\bar{u}_i\left(\hat{\sigma}_i,\sigma_{-i},f|h^{R,t}\right)=\max_{\sigma_i'\in\Sigma_i}\sum_{(f,h^{R,t})\in\mathcal{Z}^{R,t}[h_i^{R,t}]|_Q}\bar{\mu}(f,h^{R,t}|h_i^{R,t})\bar{u}_i\left(\sigma_i',\sigma_{-i},f|h^{R,t}\right)$$

and, for each $h_i^{A,t}$,

$$\sum_{(f,h^{A,t})\in\mathcal{Z}^{A,t}[h_i^{A,t}]|_Q}\bar{\mu}(f,h^{A,t}|h_i^{A,t})\bar{u}_i\left(\hat{\sigma}_i,\sigma_{-i},f|h^{A,t}\right)=\max_{\sigma_i'\in\Sigma_i}\sum_{(f,h^{A,t})\in\mathcal{Z}^{A,t}[h_i^{A,t}]|_Q}\bar{\mu}(f,h^{A,t}|h_i^{A,t})\bar{u}_i\left(\sigma_i',\sigma_{-i},f|h^{A,t}\right).$$

Let $\hat{M}_{i,t}\left(y_i^t\right)=\bigcup_{h_i^{A,t}:\mathring{h}_i^{A,t}=y_i^t}\bigcup_{\hat{\sigma}_i\in\Sigma_i^*(\sigma_{-i},\mu,\bar{\mu})}\operatorname{supp}\hat{\sigma}_{i,t}^A(h_i^{A,t})$. The following lemma shows that codominated actions are never taken by any sequentially rational strategy $\hat{\sigma}_i\in\Sigma_i\left(\sigma_{-i},\mu,\bar{\mu}\right)$.

Lemma 13 *For any game G, mediation range Q, and CPPBE $(\sigma,\mu,\bar{\mu})$, $\hat{M}_{i,t}\left(y_i^t\right)\cap\mathfrak{D}_{i,t}\left(y_i^t\right)=\emptyset$ for all i, t, and $y_i^t\in Y_i^t$.*

Proof. Suppose otherwise that there exists t such that $\hat{M}_{i,t}\left(y_i^t\right)\cap\mathfrak{D}_{i,t}\left(y_i^t\right)\neq\emptyset$ for some i and $y_i^t\in Y_i^t$. Let t^* be the last such period, and fix i, $\hat{\sigma}_i\in\Sigma_i\left(\sigma_{-i},\mu,\bar{\mu}\right)$, $y_i^{t^*}$, and action $\tilde{m}_{i,t^*}$ such that $\tilde{m}_{i,t^*}\in\bigcup_{h_i^{A,t^*}:\mathring{h}_i^{A,t^*}=y_i^{t^*}}\operatorname{supp}\hat{\sigma}_{i,t^*}^A(h_i^{A,t^*})\cap\mathfrak{D}_{i,t^*}(y_i^{t^*})$.

For each t, $y_i^t\in Y_i^t$, and $\tilde{r}_i^{t+1}\in R_i^{*t+1}$, let

$$\hat{Q}_{i,t}(\tilde{r}_i^{t+1})=\begin{cases}\hat{M}_{i,t}\left(\tilde{r}_i^{t+1}\right) & \text{if } \tilde{r}_i^{t+1}\in Y_i^t\\ A_{i,t} & \text{otherwise}\end{cases}$$

and let $\hat{Q}=(\hat{Q}_i,\tilde{Q}_{-i})$, where $\tilde{Q}_j$ is defined in (45) for $j\neq i$. By Lemma 10, for each t the range of $\tilde{Q}_{-i,t}$ excludes all codominated actions; and by definition of t^*, for each $t>t^*$, the range of $\hat{Q}_{i,t}$ excludes all codominated actions as well.

Applying the same construction as in the proof of Lemma 10 with $\hat{\sigma}_i$ in place of σ_i yields a CPS $\tilde{\mu}_{t^*}$ on $\tilde{F}|_{\hat{Q}}\times Y^{t^*}$ such that, for each $\sigma_i'\in\Sigma_i^*$,

$$\begin{aligned}&\sum_{\left(\tilde{f},y^t\right)\in\tilde{F}|_{\hat{Q}}\times Y^t[y_i^t]}\tilde{\mu}_t(\tilde{f},y^t|y_i^t,\tilde{m}_{i,t})\bar{u}_i\left(\hat{\sigma}_i,\sigma_{-i},\tilde{f}|\bar{h}\left(\tilde{f},y^t\right),\tilde{m}_{i,t}\right)\\ \geq&\sum_{\left(\tilde{f},y^t\right)\in\tilde{F}|_{\hat{Q}}\times Y^t[y_i^t]}\tilde{\mu}_t(\tilde{f},y^t|y_i^t,\tilde{m}_{i,t})\bar{u}_i\left(\sigma_i',\sigma_{-i},\tilde{f}|\bar{h}\left(\tilde{f},y^t\right),\tilde{m}_{i,t}\right).\end{aligned}$$

Let $E\subset\tilde{F}|_{\tilde{Q}}\times Y^{t^*}$ denote the set of pairs $\left(\tilde{f},y^{t^*}\right)$ such that $\tilde{f}_{i,t^*}(y^{t^*})=\tilde{m}_{i,t^*}$. Letting $\tilde{\mu}_{t^*}^E\left(\tilde{f},y^{t^*}\right)=\tilde{\mu}_{t^*}(\tilde{f},y^{t^*}|E)$, we have

$$\sum_{\substack{\left(\tilde{f},y^{t^*}\right)\in F|_{\hat{Q}}\times Y^{t^*}:\\ \tilde{f}_{i,t^*}(y^{t^*})=\tilde{m}_{i,t^*}}}\tilde{\mu}_{t^*}^E(\tilde{f},y^{t^*})\bar{u}_i\left(\sigma^*,\tilde{f}|\bar{h}\left(\tilde{f},y^{t^*}\right)\right)\geq\sum_{\substack{\left(\tilde{f},y^{t^*}\right)\in F|_{\hat{Q}}\times Y^{t^*}:\\ \tilde{f}_{i,t^*}(y^{t^*})=\tilde{m}_{i,t^*}}}\tilde{\mu}_{t^*}^E(\tilde{f},y^{t^*})\bar{u}_i\left(\sigma_i',\sigma_{-i}^*,\tilde{f}|\bar{h}\left(\tilde{f},y^{t^*}\right)\right).$$

This contradicts the hypothesis that $\tilde{m}_{i,t^*}\in\mathfrak{D}_{i,t^*}(y_i^{t^*})$. ■

Turning to the proof of the claim, note that in games of pure moral hazard, for each i and t, the set of codominated actions for player i in period t does not depend on the realized payoff-relevant history y_i^t: that is, there exists a set $D_{i,t} \subset A_{i,t}$ such that $\mathfrak{D}_{i,t}(y_i^t) = D_{i,t}$ for all $y_i^t \in Y_i^t$. Let $\bar{Y}_i^t$ be the set of payoff-relevant histories that can be reached without any player taking a codominated action:

$$\left\{ y_i^t \in Y_i^t : \exists y_{-i}^t \text{ s.t. } \left(y_i^t, y_{-i}^t\right) \in Y^t \text{ and } a_{n,\tau} \in A_{n,\tau} \setminus D_{n,\tau} \text{ for each } n \text{ and } \tau \leq t-1 \right\}.$$

Define L and π_t^k as in the proof of Proposition 5, where now each $\pi_t^{[l]} \in \Delta\left(F^{*\geq t} \times Y^t\right)$ depends only on its first component.

Fix a canonical NE (σ^*, μ^*) in G^* where players never take codominated actions at any history.[50] We first construct a (possibly non-canonical) SE $(\tilde{\sigma}, \tilde{\mu})$ in G^* with outcome ρ^{σ^*,μ^*}, and then construct a canonical SE $(\tilde{\sigma}^*, \tilde{\mu}^*)$ with the same outcome.

Construction of $\left(\tilde{\sigma}^k, \tilde{\mu}^k\right)$ We first construct a sequence of profiles $\left(\tilde{\sigma}^k, \tilde{\mu}^k\right)$ indexed by k that limits to a (quasi) SE in G^*.

Mediator's Strategy: At the beginning of the game, the mediator draws $\zeta_t \in \{0,1\}$ for each $t = 0, ..., T$ such that $\zeta_t = 1$ with probability $\left(\frac{1}{k}\right)^{2(L+1)T}$, independently across periods. Given $(\zeta_t)_{t=0}^T$, the mediator draws $(\omega_t)_{t=0}^T$ as follows: For $t = 0$, $\omega_0 = (0, f)$, where f is distributed according to $\mu^*(f)$. For each $t \geq 1$, given $\left(\omega_{t-1}, \zeta_{t-1}\right)$, ω_t is determined as follows:

1. If $\zeta_{t-1} = 0$, then $\omega_t = \omega_{t-1}$ with probability $1 - \frac{1}{k}$, and $\omega_t = \left(t, f^{\geq t}\right)$ with probability $\frac{1}{k}\pi_t^k\left(f^{\geq t}\right)$.
2. If $\zeta_{t-1} = 1$, then $\omega_t = \left(t, f^{\geq t}\right)$ with probability $\pi_t^k\left(f^{\geq t}\right)$.

Given (ζ_t, ω_t), the mediator's recommendation is determined as follows: If $\zeta_t = 0$, then $m_t = f\left(r^{t+1}\right)$ if $\omega_t = \omega_0$ and $m_t = f^{\geq \tau}\left(r^{t+1}\right)$ if $\omega_t = \left(\tau, f^{\geq \tau}\right)$. If $\zeta_t = 1$, then the mediator draws each $m_{i,t} \in A_{i,t} \setminus D_{i,t}$ with equal probability $1/\left(|A_{i,t}| - |D_{i,t}|\right)$, independently across i and t.

Each realization of the mediator's randomization defines the realization of the first element of ω_t, for each t. Let $\mathbf{t}^*(t)$ be the corresponding random variable, which takes values in $\{0, ..., t\}$.

Player i's Strategy: We say that player i is *honest and obedient* at history $h_i^{R,t}$ if $r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$ and $a_{i,\tau} = m_{i,\tau}$ for all $\tau < t$, and is *honest and obedient* at history $h_i^{A,t}$ if $r_{i,\tau} = (a_{i,\tau-1}, s_{i,\tau})$ for all $\tau \leq t$ and $a_{i,\tau} = m_{i,\tau}$ for all $\tau < t$. Let $\hat{J}_i^{R,t}$ (resp., $\hat{J}_i^{A,t}$) be the set of histories $h_i^{R,t}$ (resp., $h_i^{A,t}$) such that player i has been honest and obedient and $\mathring{h}_i^{R,t} \in \bar{Y}_i^t$.

For each k, let $\left(\Gamma^k, \mathfrak{C}^*\right)$ denote the following auxiliary game:

1. The mediator follows $\tilde{\mu}^k$.
2. Each player i chooses probability distributions $\sigma_{i,t}^{R,k}(\cdot|h_i^{R,t}) \in \Delta(R_{i,t})$ and $\sigma_{i,t}^{A,k}(\cdot|h_i^{A,t}) \in \Delta(A_{i,t} \setminus D_{i,t})$ for each t, $h_i^{R,t} \in H_i^{R,t} \setminus \hat{J}_i^{R,t}$ and $h_i^{A,t} \in H_i^{A,t} \setminus \hat{J}_i^{A,t}$. Note that player i is

[50] As in footnote 45, it is without loss to consider here the fully canonical strategy profile σ^*.

required not to take a codominated action. At histories $h_i^{R,t} \in \hat{J}_i^{R,t}$ and $h_i^{A,t} \in \hat{J}_i^{A,t}$, player i is required to report $r_{i,t} = (a_{i,t-1}, s_{i,t})$ and take $a_{i,t} = m_{i,t}$, respectively.

3. Given $\left(\sigma^k, \tilde{\mu}^k\right)$, the distribution of terminal nodes 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 $\sigma_{i,t}^{R,k}\left(r_{i,t}|h_i^{R,t}\right)$.

   Given $h^{A,t} \in H^{A,t}$, each $a_{i,t} \in R_{i,t}$ is drawn independently across players with probability
   $$\left(1 - |A_{i,t}|\left(\frac{1}{k}\right)^{3(L+1)T^2}\right)\sigma_{i,t}^{A,k}(a_{i,t}|h_i^{A,t}) + \left(\frac{1}{k}\right)^{3(L+1)T^2}.$$

   Given $h^{A,t} \in H_t^{A,t}$ and $a_t \in A_t$, 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. Player i's payoff at terminal history $h^{t+1} \in H^{T+1}$ is $u_i\left(\mathring{h}^{t+1}\right)$.

As in the proof of Lemma 2, $\left(\Gamma^k, \mathfrak{C}^*\right)$ admits a NE $\left(\hat{\sigma}^k, \tilde{\mu}^k\right)$.
Now define strategy $\tilde{\sigma}^k$ by $\tilde{\sigma}_i^{R,k} = \hat{\sigma}_i^{R,k}$ and

$$\tilde{\sigma}_{i,t}^{A,k}(a_{i,t}|h_i^{A,t}) = \left(1 - |A_{i,t}|\left(\frac{1}{k}\right)^{3(L+1)T^2}\right)\hat{\sigma}_{i,t}^{A,k}(a_{i,t}|h_i^{A,t}) + \left(\frac{1}{k}\right)^{3(L+1)T^2}.$$

Note that the distribution over terminal histories under $\left(\hat{\sigma}^k, \tilde{\mu}^k\right)$ in game $\left(\Gamma^k, \mathfrak{C}^*\right)$ is the same as that under $\left(\tilde{\sigma}^k, \tilde{\mu}^k\right)$ in game $(\Gamma, \mathfrak{C}^*)$. Let $(\tilde{\sigma}, \tilde{\mu}) = \lim_{k\to\infty}\left(\tilde{\sigma}^k, \tilde{\mu}^k\right)$. Note that $(\tilde{\sigma}, \tilde{\mu})$ is a profile in $(\Gamma, \mathfrak{C}^*)$.

We next claim that, for each i and $\sigma_i' \in \Sigma_i^*$ such that $\text{supp}\left(\sigma_i'^{A,t}(h_i^{A,t})\right) \cap D_{i,t} = \emptyset$ for all t and $h_i^{A,t}$, we have $\bar{u}(\tilde{\sigma}, \tilde{\mu}) \geq \bar{u}(\sigma_i', \tilde{\sigma}_{-i}, \tilde{\mu})$: that is, no player i has a profitable deviation that avoid codominated actions. To see this, let $\tilde{\phi}$ denote the behavioral mediation plan induced by $\tilde{\mu}$, and let ϕ^* denote the behavioral mediation plan induced by μ^*. Then $\tilde{\phi}_t\left(\cdot|r^{t+1}, m^t\right) = \phi_t^*\left(\cdot|r^{t+1}, m^t\right)$ for all (r^{t+1}, m^t) satisfying $\Pr^{\hat{\sigma},\tilde{\phi}}\left(r^{t+1}, m^t\right) > 0$ for some player strategy $\hat{\sigma}$. Similarly, for each i, $\tilde{\sigma}_{i,t}^R\left(\cdot|h_i^{R,t}\right) = \sigma_{i,t}^{*R}\left(\cdot|h_i^{R,t}\right)$ for all $h_i^{R,t}$ where player i has been honest and obedient and $\mathring{h}_i^{R,t} \in \bar{Y}_i^t$, and similarly for acting histories $h_i^{A,t}$. Hence, for any on-path history $h_i^{R,t}$ under $(\tilde{\sigma}, \tilde{\mu})$, the continuation play of i's opponents differs from that under (σ^*, μ^*) only following a deviation by player i to a codominated action. Therefore, for any $\sigma_i' \in \Sigma_i^*$ such that $\text{supp}\left(\sigma_i'^{A,t}(h_i^{A,t})\right) \cap D_{i,t} = \emptyset$ for all t and $h_i^{A,t}$, the fact that (σ^*, μ^*) is a NE implies that $\bar{u}(\tilde{\sigma}, \tilde{\mu}) \geq \bar{u}(\sigma_i', \tilde{\sigma}_{-i}, \tilde{\mu})$.

Construction of Quasi-SE $(\tilde{\sigma}, \tilde{\mu}, \tilde{J}, \tilde{K}, \tilde{\beta})$. We now define subsets of histories $\tilde{J}$ and $\tilde{K}$ along with beliefs $\tilde{\beta}$ and a mediation range $\tilde{Q}$ such that $(\tilde{\sigma}, \tilde{\mu}, \tilde{J}, \tilde{K}, \tilde{\beta})$ is a quasi-SE in game $G^*|_{\tilde{Q}}$.

Define $\tilde{Q}_{i,t}(r_i^{t+1}, m_i^t) = A_{i,t} \setminus D_{i,t}$ for each i, t, and (r_i^{t+1}, m_i^t). For each i, let $\tilde{J}_i^{A,T+}$ equal the set of histories h_i^{T+1} such that $\Pr^{\tilde{\sigma}^k,\tilde{\mu}^k}\left(h_i^{T+1}\right) > 0$. Let $\tilde{K}^{T+}$ equal the set of mediator histories $\left(\tilde{r}^{T+1}, \tilde{m}^{T+1}\right)$ such that there exists $\sigma' \in \Sigma$ satisfying $\Pr^{\sigma',\tilde{\mu}^k}\left(r^{T+1}, m^{T+1}\right) > 0$. Let the other elements of $\tilde{J}$ and $\tilde{K}$ include all truncations of histories in $\tilde{J}^{A,T+}$ and $\tilde{K}^{T+}$. Since $\tilde{J}_i$ includes all histories where player i has not lied to the mediator or received a message outside the mediation range $\tilde{Q}$, $(\tilde{\sigma}, \tilde{\mu}, \tilde{J}, \tilde{K})$ is valid. Define belief system $\tilde{\beta}$ by

$$\tilde{\beta}_{i,t}\left(h^{R,t}|h_i^{R,t}\right) = \lim_{k\to\infty} \frac{\Pr^{\tilde{\sigma}^k,\tilde{\mu}^k}\left(h^{R,t}\right)}{\Pr^{\tilde{\sigma}^k,\tilde{\mu}^k}\left(h_i^{R,t}\right)}$$

for each $h^{R,t} \in H^{R,t}[h_i^{R,t}]|_{\tilde{J},\tilde{K}}$. Since $\Pr^{\tilde{\sigma}^k,\tilde{\mu}^k}\left(h_i^{R,t}\right) > 0$ for each $h_i^{R,t} \in \tilde{J}_i^{R,t}$, this is well-defined. Since $\tilde{\beta}$ is consistent by construction, it remains to verify sequential rationality of $(\tilde{\sigma}, \tilde{\mu}, \tilde{J}, \tilde{K}, \tilde{\beta})$.

As in the proof of Claim 1 of Proposition 4, sequential rationality at histories outside $\hat{J}_i^{R,t}$ or $\hat{J}_i^{A,t}$ follows from a standard upper hemi-continuity argument.

Fix any $h_i^{R,t} \in \hat{J}_i^{R,t}$. By Lemma 13, it suffices to show that player i never has a profitable deviation to a strategy that avoids codominated actions. By definition of $\hat{J}_i^{R,t}$, there exists $\left(\omega^t, \zeta^t\right) = (\omega_1, ..., \omega_{t-1}, \zeta_1, ..., \zeta_{t-1})$ such that $\Pr^{\sigma^*,\tilde{\mu}^k}\left(h_i^{R,t}|\omega^t, \zeta^t\right) > 0$, where as usual σ^* denotes the fully canonical strategy profile. Since under $\tilde{\mu}^k$ any sequence $\left(\omega^t, \zeta^t\right)$ occurs with probability at least $\left(\frac{1}{k}\right)^{2(L+1)T+2(L+1)T^2}$, while under $\tilde{\sigma}^k$ players' actions differ from those under $\hat{\sigma}^k$ with probability at most $\left(\frac{1}{k}\right)^{3(L+1)T^2}$, player i assesses that other players are honest and obedient with probability 1. Hence, it suffices to verify that, for each i, t, $\tau \leq t$, and $h_i^{R,t} \in \hat{J}_i^{R,t}$, $\tilde{\sigma}_i$ is sequentially rational conditional on the event that all players are honest and obedient and $\mathbf{t}^*(t) = \tau$. There are three mutually exclusive events to which player i assigns positive probability:

1. If $\mathbf{t}^*(t) = 0$ then $\Pr^{\tilde{\sigma},\tilde{\mu}}\left(h_i^{R,t}\right) > 0$. Sequential rationality follows since $\bar{u}(\tilde{\sigma}, \tilde{\mu}) \geq \bar{u}(\sigma_i', \tilde{\sigma}_{-i}, \tilde{\mu})$ for any strategy σ_i' that avoids codominated actions.

2. If $\mathbf{t}^*(t) = \tau < t$ then $\omega_t = \left(\tau, f^{\geq\tau}\right)$. Since the mediator draws $f^{\geq\tau}$ from π_τ^k, by the same argument as in the $\xi\left(h_i^{R,t}\right) = \left(t^*, \zeta_i^{t^*}\right)$ case of the proof of Proposition 5 (i.e., the discussion immediately following Lemma 6), honesty is optimal.

3. If $\mathbf{t}^*(t) = t$ then with probability 1 future recommendations are independent of the current report. Hence, honesty is optimal.

Next fix any $h_i^{A,t} \in \hat{J}_i^{A,t}$. Again, player i assesses that other players are honest and obedient with probability 1. She also assesses that $\zeta_t = 1$ with probability 0, since (i) for any $\left(h_i^{R,t}, r_{i,t}\right)$, each $m_{i,t} \in A_{i,t} \setminus D_{i,t}$ occurs with positive probability conditional on $\zeta_t = 0$ and $\mathbf{t}^*(t) = t$, (ii) $\zeta_t = 0$ and $\mathbf{t}^*(t) = t$ occur with probability at least $\left(1 - \left(\frac{1}{k}\right)^{2(L+1)T}\right)\left(\frac{1}{k}\right)^L$, while $\zeta_t = 1$ occurs with probability at most $\left(\frac{1}{k}\right)^{2(L+1)T}$. Hence, we again consider three

mutually exclusive events: $\mathbf{t}^*(t) = 0$, $\mathbf{t}^*(t) = \tau < t$, and $\mathbf{t}^*(t) = t$ and $\zeta_t = 0$. The proof of sequential rationality in the first two events is the same as for reporting histories. In the last event, sequential rationality follows from the same argument as in the $\xi\left(h_i^{R,t}\right) = \left(t, \zeta_i^t\right)$ case of the proof of Proposition 5.

Construction of $\left(\tilde{\sigma}^{*,k}, \tilde{\mu}^{*,k}\right)$ We first construct a sequence of profiles $\left(\tilde{\sigma}^{*,k}, \tilde{\mu}^{*,k}\right)$ indexed by k that limits to a quasi-SE profile in G^*.

Definition of the Mediator's Strategy $\tilde{\mu}^{*,k}$. As in Proposition 2, in each period t, given player i's canonical report $\tilde{r}_i^{t+1} \in Y_i^t$, the mediator draws a fictitious report $r_{i,t}$ according to $\hat{\sigma}_{i,t}^{R,k}(h_i^{R,t})$, where $h_i^{R,t} = (y_i^t, r_i^t, m_i^t)$ with $y^t = \tilde{r}_i^{t+1}$.[51] Given fictitious reports r^{t+1} and fictitious messages m^t, the mediator draws fictitious message m_t from $\tilde{\phi}_t^k(r^{t+1}, m^t)$, where $\tilde{\phi}^k$ is the behavioral strategy induced by $\tilde{\mu}^k$. He then draws the recommendation $\tilde{m}_{i,t}$ according to $\hat{\sigma}_{i,t}^{A,k}(\tilde{m}_{i,t}|\tilde{r}_i^{t+1}, r_i^{t+1}, m_i^{t+1})$.

Definition of Player i's Strategy $\tilde{\sigma}_i^{*,k}$. We define $\tilde{\sigma}_i^{*,k}$ to specify that player i is honest and obedient, but trembles uniformly over actions with probability $|A_{i,t}|\left(\frac{1}{k}\right)^{3(L+1)T^2}$: for each $a_{i,t} \in A_{i,t}$ and $h_i^{A,t} \in H_i^{A,t}$,

$$\tilde{\sigma}_{i,t}^{*A,k}(a_{i,t}|h_i^{A,t}) = 1_{\{a_{i,t}=m_{i,t}\}}\left(1 - |A_{i,t}|\left(\frac{1}{k}\right)^{3(L+1)T^2}\right) + \left(\frac{1}{k}\right)^{3(L+1)T^2}.$$

Construction of Quasi-SE $(\tilde{\sigma}^*, \tilde{\mu}^*, J, K, \beta)$ Define $\tilde{\sigma}^* = \lim_{k\to\infty} \tilde{\sigma}^{*,k}$ and $\tilde{\mu}^* = \lim_{k\to\infty} \tilde{\mu}^{*,k}$. We will construct J, K, and β such that $(\tilde{\sigma}^*, \tilde{\mu}^*, J, K, \beta)$ is a quasi-SE in G^*. By the first sentence of Lemma 2, this implies that there exists a SE in G^* that implements the same outcome. We will also construct a mediation range Q such that J includes all histories compatible with the mediation range where players have been honest. By the second sentence of Lemma 2, this implies that the SE is canonical.

Definition of Q. For each i and t, if $\tilde{r}_i^{t+1} \in Y_i^t$ we define

$$Q_{i,t}(\tilde{r}_i^{t+1}) = \bigcup_{\substack{(r_i^{t+1}, m_i^{t+1}) \\ \text{s.t. } r_{i,\tau} \in \operatorname{supp} \hat{\sigma}_{i,t}^{R,k}(y_i^\tau, r_i^\tau, m_i^\tau) \text{ with } y_i^\tau = \tilde{r}_i^{\tau+1} \text{ for each } \tau \le t \\ m_{i,\tau} \in A_{i,\tau} \setminus D_{i,\tau} \text{ for each } \tau \le t}} \operatorname{supp} \hat{\sigma}_{i,t}^{A,k}(\cdot|y_i^t, r_i^{t+1}, m_i^{t+1})$$

where $y_i^t = \tilde{r}_i^{t+1}$; and if $\tilde{r}_i^{t+1} \notin Y_i^t$ we define $Q_{i,t}(\tilde{r}_i^{t+1}) = A_{i,t}$.

Definition of (J, K, β). Define $K^{T+} = H_0^{T+1}$ and define $J_i^{A,T+}$ as the set of all histories compatible with the mediation range where players have been honest. Let the other elements of J and K include all truncations of histories in $J^{A,T+}$ and K^{T+}. Since $Q_{i,t}(\tilde{r}_i^{t+1})$ contains all messages ever sent under $\tilde{\mu}^{*,k}$ when the history of communications between the mediator and player i is $\tilde{r}_i^{t+1}$, $(\tilde{\sigma}^*, \tilde{\mu}^*, J, K)$ is valid.

[51] If $\tilde{r}_i^{t+1} \in R_i^{*t+1} \setminus Y_i^t$, then the mediator draws $r_{i,t}$ uniformly at random, and similarly for $\tilde{m}_{i,t}$ in what follows.

Next, for each $h_i^{R,t} \in J_i^{R,t}$ and $h^{R,t} \in H^{R,t}[h_i^{R,t}]|_{J,K}$, define

$$\beta_{i,t}\left(h^{R,t}|h_i^{R,t}\right) = \lim_{k\to\infty} \frac{\Pr^{\tilde{\sigma}^{*,k},\tilde{\mu}^{*,k}}\left(h^{R,t}\right)}{\Pr^{\tilde{\sigma}^{*,k},\tilde{\mu}^{*,k}}\left(h_i^{R,t}\right)}.$$

To see that the denominator of this expression is positive (so the quotient is well-defined), note that, since (i) under $\tilde{\mu}^k$ every sequence ζ^t occurs with positive probability and, given $\zeta_t = 1$, for each i the mediator sends each $m_{i,t} \in A_{i,t} \setminus D_{i,t}$ with positive probability at every history, and (ii) players tremble uniformly over actions under $\tilde{\sigma}_i^{*,k}$, it follows that $\Pr^{\tilde{\sigma}^{*,k},\tilde{\mu}^{*,k}}\left(h_i^{R,t}\right) > 0$ for all $h_i^{R,t} \in J_i^{R,t}$. Define $\beta_{i,t}\left(h^{A,t}|h_i^{A,t}\right)$ analogously. These beliefs are consistent by construction.

Sequential Rationality: Since the construction of $\left(\tilde{\sigma}^{*,k}, \tilde{\mu}^{*,k}\right)$ from $\left(\tilde{\sigma}^k, \tilde{\mu}^k\right)$ is the same as the construction of $\left(\tilde{\sigma}^k, \tilde{\phi}^k\right)$ from $\left(\bar{\sigma}^k, \phi\right)$ in the proof of Claim 1 of Proposition 4, sequential rationality of $(\tilde{\sigma}^*, \tilde{\mu}^*, J, K, \beta)$ follows from sequential rationality of $\left(\tilde{\sigma}, \tilde{\mu}, \tilde{J}, \tilde{K}, \tilde{\beta}\right)$.