SUPPLEMENT TO "SEQUENTIAL INFORMATION DESIGN"

(*Econometrica*, Vol. 88, No. 6, November 2020, 2575–2608)

LAURA DOVAL
Division of the Humanities and Social Sciences, California Institute of Technology

JEFFREY C. ELY
Department of Economics, Northwestern University

This supplement contains results omitted from the main text of the paper.

S1. EXAMPLES OF NON-ADMISSIBLE EXTENSIVE FORMS

WE ILLUSTRATE THAT THE CONDITIONS defining admissibility (no delegation, know your action, and no partial commitments) are in a sense *tight*. In what follows, we present three examples of extensive-form games, each of which satisfies two out of the three conditions and a perfect Bayesian equilibrium of these extensive forms that cannot be implemented as an obedient perfect Bayesian equilibrium of the canonical extensive form. Since all three examples feature $|\Theta| = 1$, we omit the move by chance at the beginning of the tree.

S1.1. *An Extensive Form That Fails 'No Delegation'*

EXAMPLE S1.1: Consider the following base game G, where $N = \{1, 2\}$, $A_i = \{C, D\}$, and payoffs are follows:

Prisoner's Dilemma

	C	D
C	$5, 5$	$-10, 10$
D	$10, -10$	$0, 0$

Consider the following extensive form Γ. Edges labeled as actions represent moves that determine that player's action in the terminal history to be the one in the label.[1]

The unique subgame-perfect Nash equilibrium of this game is illustrated in blue in Figure S1, and it leads to $\{C, C\}$ being played with probability 1. This only depends on $u_i(C, C) > u_i(D, D)$ for $i \in \{1, 2\}$.

Now consider the assumptions needed to sustain this outcome (i.e., both players cooperate with probability 1) as an obedient Perfect Bayesian equilibrium of the canonical extensive form. Letting p denote the probability that player 1 moves first, the canonical extensive form can implement $\{C, C\}$ with probability 1 only if $p(u_1(C, C) - u_1(D, D)) \geq (1 - p)(u_1(D, C) - u_1(C, C))$, and $(1 - p)(u_2(C, C) - u_2(D, D)) \geq p(u_2(D, C) - u_2(C, C))$. These conditions are clearly stronger than $u_i(C, C) > u_i(D, D)$ for $i \in \{1, 2\}$.

Laura Doval: ldoval@caltech.edu
Jeffrey C. Ely: jeffely@northwestern.edu

[1] This example is based on Peters (2015).

 https://doi.org/10.3982/ECTA17260

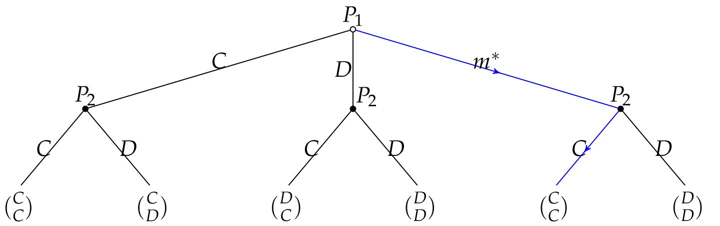

FIGURE S1.—An extensive form that satisfies know your action and no partial commitments.

S1.2. *An Extensive Form That Fails 'Know Your Own Action'*

EXAMPLE S1.2—Know Your Action: Consider the following three-player base game, G:

Base Game, G

	D	E
A	$10, 1, -1$	$2, -2, 0$
B	$0, -1, 1$	$0, -3, -1$
	F	

	D	E
A	$1, -1, 1$	$1, -3, -1$
B	$0.5, 1, -1$	$0, -2, 0$
	G	

Consider the following extensive form illustrated in Figures S2 and S3.

Note that the extensive form in Figure S2 fails know your action: when player 1 chooses move m_1, he plays A if Nature moved left, and he plays B otherwise. Thus, move m_1 determines player 1's action; however, he does not know when he chooses m_1 what action he will be playing in the base game, G.

Figure S2 depicts in blue a subgame-perfect Nash equilibrium in which player 1 chooses m_1, player 2 chooses D, and player 3 mixes with probability 0.5 between F and G. Off path, when player 1 chooses A, player 2 chooses D, and player 3 chooses G.

Now consider implementing the above distribution over action profiles as an obedient perfect Bayesian equilibrium of the canonical extensive form. player 1 must receive with positive probability a recommendation to play B, which he would never obey. In the extensive form in Figure S2, however, by choosing m_1, it is as if player 1 could choose a

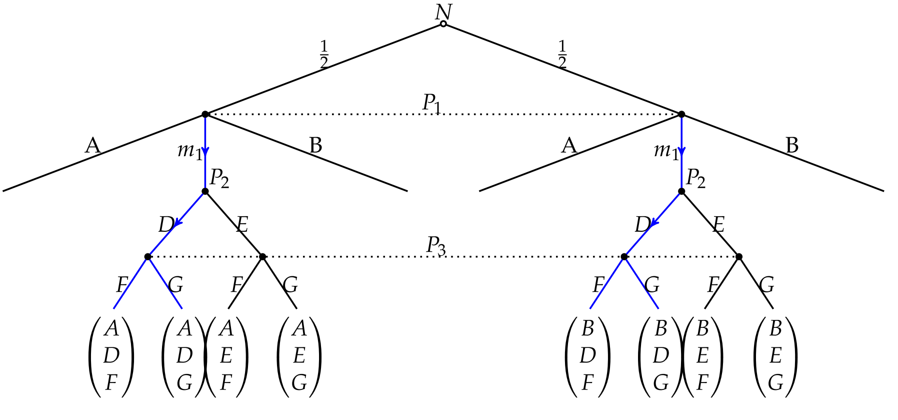

FIGURE S2.—Extensive-form game that satisfies no delegation and no partial commitments.

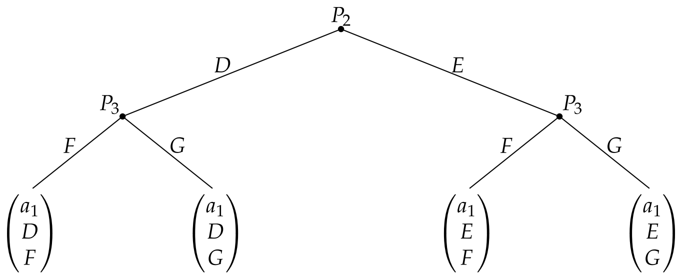

FIGURE S3.—Subgame starting after moves $a_1 \in \{A, B\}$.

randomization between A and B—induced by the moves of player 3—that he himself is not willing to carry out.

S1.3. *An Extensive Form That Fails 'No Partial Commitments'*

EXAMPLE S1.3: The following example is based on Myerson (1986). Consider the following base game, G, where $N = \{1, 2\}$, $A_1 = \{T, M, B\}$, and $A_2 = \{L, R\}$. Payoffs are as follows:

Base Game, G

	L	R
T	2, 2	2, 2
M	5, 1	0, 0
B	0, 0	1, 5

Consider the following extensive form (illustrated in Figure S4) that satisfies no delegation and know your action.

Figure S4 depicts in blue the unique subgame-perfect Nash equilibrium of the game. Consider implementing the induced distribution over action profiles as an obedient perfect Bayesian equilibrium of the canonical extensive form. This requires that player 1 is recommended with positive probability to play B. This recommendation can never be

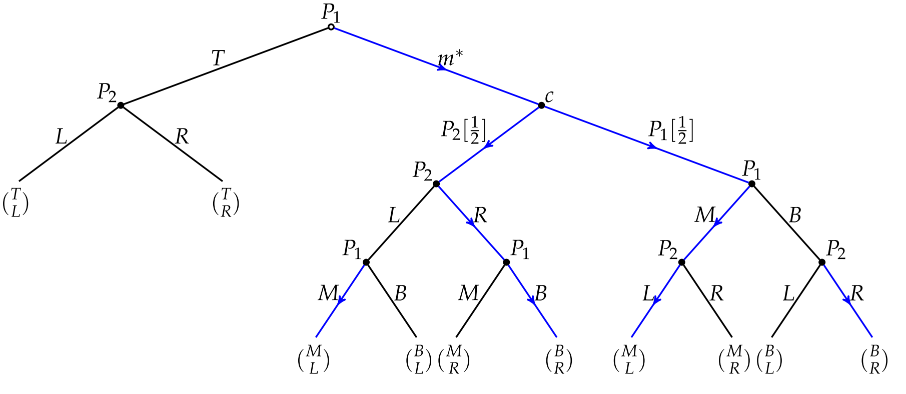

FIGURE S4.—Extensive form that satisfies no delegation and know your action.

made obedient: by playing T, player 1 can make sure he obtains a payoff of 2, while conditional on being recommended B, his payoff is 1.

S2. ILLUSTRATING BAYES' RULE WHERE POSSIBLE

We will exhibit a conditional probability system for the example in Figure 6 to demonstrate that the system of beliefs satisfies Bayes' rule where possible.

In the extensive form, let Σ_i denote the set of pure strategies for player i, including Nature. Nature's pure strategy set Σ_0 is $\{a, d\}$, where a indicates the move "across" and d indicates "down." The equilibrium behavior strategy profile β has the players playing pure strategies, which we will denote by σ_i^*, and Nature playing the mixed strategy $(q, 1-q)$. Therefore, we set

$$\mu\big(a, \sigma_1^*, \sigma_2^*, \sigma_3^* \mid \Sigma\big) = q, \tag{S1}$$

$$\mu\big(d, \sigma_1^*, \sigma_2^*, \sigma_3^* \mid \Sigma\big) = 1 - q \tag{S2}$$

for the "top-level" or unconditional probabilities of pure strategy profiles. Let Σ^0 be the support, that is,

$$\Sigma^0 = \Sigma_0 \times \{\sigma_1^*\} \times \{\sigma_2^*\} \times \{\sigma_3^*\}.$$

Next, the formula $\beta^{\Sigma}(\cdot)$ gives the conditional probabilities over pure-strategy profiles at each node in the tree, as derived from the equilibrium behavioral strategy profile. We specify that $\mu(\cdot \mid y) = \beta^{\Sigma}(y)$ for each node y in the tree.

Next, we define $\mu(\cdot \mid \Sigma \setminus \Sigma^0)$. This is the conditional probability distribution over strategy profiles conditional on a deviation from Σ^0. We will specify

$$\mu\big(d, X_1, \sigma_2^*, \sigma_3^* \mid \Sigma \setminus \Sigma^0\big) = q, \tag{S3}$$

$$\mu\big(a, X_1, \sigma_2^*, \sigma_3^* \mid \Sigma \setminus \Sigma^0\big) = 1 - q. \tag{S4}$$

That is, the deviation was player 1 choosing X_1. These are the only aspects of the conditional probability system that matter for verifying Bayes' rule where possible. We can thus define the remaining conditional probabilities arbitrarily. (By Section S4, we can extend μ to a fully specified conditional probability system.)

Now, to verify Bayes' rule where possible, we must consider the beliefs at the two non-singleton information sets. We will illustrate for the information set belonging to player 2; the analogous argument applies to player 3. The system of beliefs assigns conditional probabilities $(0, q, 1-q)$ to nodes e, f, and g, respectively, within player 2's information set.

Let Σ^e, Σ^f, and Σ^g be the (disjoint) sets of pure-strategy profiles with nodes e, f, and g on their paths. Their union, call it h_2, is the set of pure-strategy profiles having player 2's non-singleton information set on the path. Bayes' rule where possible is the restriction that

$$\mu\big(\Sigma^e \mid h_2\big) = 0, \tag{S5}$$

$$\mu\big(\Sigma^f \mid h_2\big) = q, \tag{S6}$$

$$\mu\big(\Sigma^g \mid h_2\big) = 1 - q. \tag{S7}$$

Since $\Sigma^e \subset h_2 \subset \Sigma \setminus \Sigma^0$, the chain rule for conditional probabilities implies

$$\mu(\Sigma^e \mid \Sigma \setminus \Sigma^0) = \mu(h_2 \mid \Sigma \setminus \Sigma^0) \cdot \mu(\Sigma^e \mid h_2),$$

and since by previous construction we have $\mu(\Sigma^e \mid \Sigma \setminus \Sigma^0)$ and $\mu(h_2 \mid \Sigma \setminus \Sigma^0) > 0$, we obtain $\mu(\Sigma^e \mid h_2) = 0$ as desired. Next,

$$\mu(\Sigma^f \mid \Sigma \setminus \Sigma^0) = \mu(h_2 \mid \Sigma \setminus \Sigma^0) \cdot \mu(\Sigma^f \mid h_2).$$

Since $\Sigma^f = \{d\} \times \{X_1\} \times \Sigma_2 \times \Sigma_3$, we have by Eq. (S3) that $\mu(\Sigma^f \mid \Sigma \setminus \Sigma^0) = q$. Furthermore, since $\Sigma_0 \times \{X_1\} \times \{\sigma_2^*\} \times \{\sigma_3^*\} \subset h_2$ and $\mu(\Sigma_0 \times \{X_1\} \times \{\sigma_2^*\} \times \{\sigma_3^*\} \mid \Sigma \setminus \Sigma^0) = 1$, it follows that $\mu(h_2 \mid \Sigma \setminus \Sigma^0) = 1$ and thus $\mu(\Sigma^f \mid \Sigma \setminus \Sigma^0) = q$ as desired.

A similar calculation shows that $\mu(\Sigma^g \mid h_2) = 1 - q$.

S3. PURE STRATEGY AND NATURE MOVES FIRST

THEOREM S1: *For any perfect Bayesian equilibrium (β, ν) of an admissible extensive form Γ, there is another admissible extensive form, Γ', and a perfect Bayesian equilibrium (β', ν') of Γ', where players use pure strategies, which implements the same distribution over outcomes as the original one did. Moreover, Nature moves only once at the beginning in Γ'.*

S3.1. *Notation*

Let $\Sigma = \Sigma_0 \times \Sigma_1 \times \cdots \times \Sigma_N$ denote the strategy set. We use the notation $\Sigma_{-0} \equiv \Sigma_1 \times \cdots \times \Sigma_N$ to denote the players' strategy set. Given a node y, let $\sigma_i|_y \subseteq \Sigma_i$ denote the set of strategies of player i that coincide with σ_i at all nodes y' such that $y' \not\prec y$. Let $\sigma|_y = \prod_{i=0}^{N} \sigma_i|_y$.

A node y is on the path of strategy $\sigma \in \Sigma$ if, for all nodes y' preceding y, we have that $(y', \sigma(y'))$ weakly precedes y.

Given a node y, let $\Sigma_j(y)$ denote the set of strategy profiles of j such that there exists $\sigma_{-j} \in \Sigma_{-j}$ for which y is on the path of (σ_j, σ_{-j}).

S3.2. *Proof*

Let (β, ν) denote a perfect Bayesian equilibrium of extensive form Γ.

Let $\Sigma^* = \Sigma_0^* \times \Sigma_{-0}^*$ denote the support of $\mu(\cdot|\Sigma)$, where μ is the conditional probability system associated with assessment, (β, ν). Kuhn's theorem implies that $\mu(\cdot|\Sigma)$ implements the same distribution over terminal nodes as β. In what follows, we label the set $\Sigma_{-0}^* = \{\sigma_{-0,1}, \dots, \sigma_{-0,N}\}$ and say that $\sigma < \sigma'$ for $\sigma, \sigma' \in \Sigma_{-0}^*$ if the label of σ is less than the label of σ'.

Define a new extensive form Γ' as follows. First, we construct the tree, V'. The initial move belongs to Nature where she chooses an element of $\Sigma_0 \times \Sigma_{-0}$.[2] After each such move, we append a copy of the tree from Γ, V. For every node y in the original tree, let $\lambda_\sigma(y)$ represent the node in Γ' that corresponds to y when Nature chooses $\tilde{\sigma} \in \Sigma_0 \times \Sigma_{-0}$. Fix $y \in V$ and $\sigma \in \Sigma_0 \times \Sigma_{-0}$. Eliminate $\lambda_\sigma(y)$ from the tree we just constructed if $\tilde{\sigma}_0 \notin \Sigma_0(y)$. Let V' denote the nodes that remain after this.

[2]Why isn't this Σ_0^*? In the original tree, players could have assigned positive probability to Nature deviating. This may be important for incentives at some information set off the path of play. Now Nature will move first so the only chance it has of deviating is at the root. So we cannot think of Nature just choosing from Σ_0^*.

REMARK S1: It is immediate to check that V' is a tree. Clearly, $\emptyset \in V'$. Moreover, let $\overline{y} \in V'$ and let $\overline{y}' \prec \overline{y}$. We show that $\overline{y}' \in V'$ as well. To see this, suppose first that $\overline{y} = \sigma \in \Sigma_0 \times \Sigma^*_{-0}$. Then, $\overline{y}' \prec \overline{y} \Rightarrow \overline{y}' = \emptyset \in V'$. Suppose then that $\overline{y} = \lambda_\sigma(y)$ for some $\sigma \in \Sigma_0 \times \Sigma^*_{-0}$ and some $y \in V$ such that $\sigma_0 \in \Sigma_0(y)$. Then, $\overline{y}' \prec \overline{y}$ implies that $\overline{y}' = \lambda_\sigma(y')$ for some $y' \preceq y$. Note that $\sigma_0 \in \Sigma_0(y)$ implies that $\sigma_0 \in \Sigma_0(y')$. Thus, $\overline{y}' \in V'$. Hence, V' is a tree.

Second, we construct information sets. Fix player i and an information set h_i in the original tree. For each $\sigma^*_i \in \Sigma^*_i$, there is a copy of h_i, $h_i(\sigma^*_i)$, in the new tree formed by all $\lambda_{(\sigma^*_i, \sigma_{-i})}(y)$ such that $y \in h_i$ and $\sigma_0 \in \Sigma_0(y)$. We denote the strategy set in Γ' as Σ'.

We specify Nature's initial randomization as $\mu(\cdot \mid \Sigma)$.

Define the obedient strategy σ^O_i in Γ' as follows. For each i and $\sigma_i \in \Sigma^*_i$, $\sigma^O_i(h_i(\sigma_i)) = \sigma_i(h_i)$. This implicitly defines a behavioral strategy profile β^O such that for each i and $\sigma_i \in \Sigma^*_i$, $\beta^O_i(h_i(\sigma_i))$ assigns probability 1 to $\sigma_i(h_i)$.

To construct the conditional probability system, μ' (and the belief system, ν'), of the assessment in Γ', we now define a map between Σ and Σ'. First, let $S \subseteq \Sigma$ be the set of strategy profiles $\tilde{\sigma} \in \Sigma$ such that there exists y on the path of $\tilde{\sigma}$ and some $\sigma_{-0} \in \Sigma^*_{-0}$ such that $\tilde{\sigma} \in (\tilde{\sigma}_0, \sigma_{-0})|_y$. Let $\underline{y}_{\tilde{\sigma}}$ denote the earliest y on the path of $\tilde{\sigma}$ for which this is true.

Define

$$\Sigma^*_{-0}(\tilde{\sigma}) = \big\{\sigma_{-0} \in \Sigma^*_{-0} : \tilde{\sigma}|_{\underline{y}_{\tilde{\sigma}}} \in (\tilde{\sigma}_0, \sigma_{-0})|_{\underline{y}_{\tilde{\sigma}}}\big\}.$$

If $|\Sigma^*_{-0}(\tilde{\sigma})| = 1$, define

$$g(\tilde{\sigma})\big(\lambda_\sigma(y)\big) = \begin{cases} \big(\tilde{\sigma}_0, \Sigma^*_{-0}(\tilde{\sigma})\big) & \text{if } y = \emptyset, \\ \tilde{\sigma}(y) & \text{if } y \prec \underline{y}_{\tilde{\sigma}} \text{ and } \sigma = \big(\tilde{\sigma}_0, \Sigma^*_{-0}(\tilde{\sigma})\big), \\ \sigma^O(y) & \text{otherwise.} \end{cases} \tag{S8}$$

Otherwise, if $|\Sigma^*_{-0}(\tilde{\sigma})| > 1$, let i denote the player on the move at $\underline{y}_{\tilde{\sigma}}$. If there exist $\sigma, \sigma' \in \Sigma^*_{-0}(\tilde{\sigma})$ such that $\sigma_i \neq \sigma'_i$, then let $g(\tilde{\sigma})$ denote the strategy in Σ' constructed as in equation (S8), except that $g(\tilde{\sigma})(\emptyset) = (\tilde{\sigma}_0, \sigma)$ where $\sigma \in \Sigma^*_{-0}(\tilde{\sigma})$ satisfies that $\sigma_i <_i \sigma'_i$ for all $\sigma' \in \Sigma^*_{-0}(\tilde{\sigma})$. If no such $\sigma, \sigma' \in \Sigma^*_{-0}(\tilde{\sigma})$ exist, then let $g(\tilde{\sigma})$ denote the strategy in Σ' constructed as in equation (S8), except that $g(\tilde{\sigma})(\emptyset) = (\tilde{\sigma}_0, \sigma)$ where $\sigma \in \Sigma^*_{-0}(\tilde{\sigma})$ satisfies that $\sigma < \sigma'$ for all $\sigma' \in \Sigma^*_{-0}(\tilde{\sigma})$.

We note for what follows that if $|\Sigma^*_{-0}(\tilde{\sigma})| > 1$, then $\mu(\tilde{\sigma}|\Sigma) = 0$. To see this, consider two cases. First, if $\tilde{\sigma}_0 \notin \Sigma^*_0$, the implication is immediate. Second, assume $\tilde{\sigma}_0 \in \Sigma^*_0$ and let $\sigma, \sigma' \in \Sigma^*_{-0}(\tilde{\sigma})$, $\sigma \neq \sigma'$. Then, $(\tilde{\sigma}_0, \sigma)|_{\underline{y}_{\tilde{\sigma}}} = (\tilde{\sigma}_0, \sigma')|_{\underline{y}_{\tilde{\sigma}}}$. Clearly, it cannot be that $\underline{y}_{\tilde{\sigma}}$ is on the path of either $(\tilde{\sigma}_0, \sigma)$ or $(\tilde{\sigma}_0, \sigma')$; otherwise, $\underline{y}_{\tilde{\sigma}} = \{\emptyset\}$ and this contradicts that $\sigma \neq \sigma'$.

S3.2.1. *Properties of g*

We now make some observations about the properties of $g^{-1}(\lambda_\sigma(y))$ for nodes y in the tree V of Γ.

To this end, given a node $y \in V$ and $\sigma = (\sigma_0, \sigma_{-0}) \in \Sigma_0 \times \Sigma^*_{-0}$, let $s(\sigma, y)$ denote the following strategy in Σ. First, $s(\sigma, y) \in \sigma|_y$. Second, y is on the path of $s(\sigma, y)$. Third, $s_0(\sigma, y) = \sigma_0$. We note that $S = \bigcup_{(\sigma_0, \sigma_{-0}) \in \Sigma_0 \times \Sigma^*_{-0}} \bigcup_{y \in V} \{s(\sigma, y)\}$. Moreover, it follows from the definition of PBE that $\mu(s|y) > 0$ implies that $s = s(\sigma, y)$ for some $\sigma_{-0} \in \Sigma^*_{-0}$ and some σ_0 such that there exists $\sigma^*_0 \in \Sigma^*_0$ with $\sigma_0|_y \in \sigma^*_0|_y$.

The following properties can be shown by simple application of the definition of g and $\underline{y}_{\tilde{\sigma}}$:

LEMMA S1: *$g^{-1}(\lambda_\sigma(y)) \neq \emptyset \Leftrightarrow \{s(\sigma, y)\} = g^{-1}(\lambda_\sigma(y))$. Moreover, if $s(\sigma, y) \notin g^{-1}(\lambda_\sigma(y))$, then there exists $\sigma'_{-0} \in \Sigma^*_{-0}$ such that $(\sigma_0, \sigma'_{-0})|_y = (\sigma_0, \sigma_{-0})|_y$ and either*

1. *$s(\sigma, y) \in (\sigma_0, \sigma'_{-0})|_{\underline{y}_{s(\sigma,y)}} \neq (\sigma_0, \sigma_{-0})|_{\underline{y}_{s(\sigma,y)}}$,[3] or*
2. *$(\sigma_0, \sigma_{-0})|_{\underline{y}_{s(\sigma,y)}} = (\sigma_0, \sigma'_{-0})|_{\underline{y}_{s(\sigma,y)}}$ and either $\sigma'_i <_i \sigma_i$ or $\sigma' < \sigma$.*

In both cases, σ'_{-0} satisfies that $s(\sigma_0, \sigma'_{-0}, y) = s(\sigma_0, \sigma_{-0}, y)$. Hence, $s(\sigma_0, \sigma_{-0}, y) = s(\sigma_0, \sigma'_{-0}, y) \in g^{-1}(\lambda_{(\sigma_0,\sigma'_{-0})}(y))$.

Fix $\sigma^*_i \in \Sigma^*_i$ and an information set for player i, $h_i \in H_i$. Suppose that there exist $\sigma_{-i} \in \Sigma_0 \times \Sigma^*_{-\{0,i\}}$ and $y \in h_i$ such that $g^{-1}(\lambda_{(\sigma^*_i,\sigma_{-i})}(y)) = \emptyset$. Lemma S1 implies that there exists $\sigma' \in \Sigma_0 \times \Sigma^*_{-0}$ such that $s(\sigma^*_i, \sigma_{-i}, y) \in g^{-1}(\lambda_{\sigma'}(y))$.

LEMMA S2: *If $\sigma'_i \neq \sigma^*_i$, then $g^{-1}(\lambda_{(\sigma^*_i,\sigma'_{-i})}(y)) = \emptyset$ for all $\sigma'_{-i} \in \Sigma_0 \times \Sigma^*_{-\{0,i\}}$. Otherwise, $\bigcup_{\sigma_{-i}\in\Sigma_0\times\Sigma^*_{-\{0,i\}}} g^{-1}(\lambda_{(\sigma^*_i,\sigma_{-i})}(y)) \neq \emptyset$.*

PROOF: Note that perfect recall implies that player i makes the same choices on the way to node y on $s(\sigma^*_i, \tilde{\sigma}_{-i}, y)$ for any choice of $\tilde{\sigma}_{-i} \in \Sigma_0 \times \Sigma^*_{-\{0,i\}}$. Now, $s(\sigma^*_i, \sigma_{-i}, y) = s(\sigma', y)$ implies that $(\sigma^*_i, \sigma_{-i})$ and σ' coincide for nodes $y' \not\prec y$; in particular, σ^*_i and σ'_i coincide on nodes $y' \not\prec y$. Now, if $\underline{y}_{s(\sigma',y)} \prec y$, then $s(\sigma', y) \in g^{-1}(\lambda_{(\sigma^*_i,\sigma'_{-i})}(y))$ and $s(\sigma^*_i, \sigma_{-i}, y) \notin g^{-1}(\lambda_{(\sigma^*_i,\sigma_{-i})}(v))$ imply that there is a node such that $y_i \prec y$, $\iota(y_i) = i$ and $s(\sigma^*_i, \sigma_{-i}, y)(y_i) = \sigma'(y_i) \neq (\sigma^*_i, \sigma_{-i})(y_i)$. Hence, for all $\tilde{\sigma}_{-i} \in \Sigma_0 \times \Sigma^*_{-\{0,i\}}$, $s(\sigma^*_i, \tilde{\sigma}_{-i}, y)(y_i) = \sigma'_i(y_i)$. It then follows that for all $\tilde{\sigma}_{-i}$, $s(\sigma^*_i, \tilde{\sigma}_{-i}, y) \notin g^{-1}(\lambda_{(\sigma^*_i,\tilde{\sigma}_{-i})}(y))$.

Now, suppose that $\underline{y}_{s(\sigma',y)} = y$; then $|\Sigma^*_{-0}(s(\sigma', y))| > 1$ and since $\sigma^*_i \neq \sigma'_i$, we must have $\sigma'_i <_i \sigma^*_i$. Towards a contradiction, suppose that there exists $\hat{\sigma_{-i}}$ such that $g^{-1}(\lambda_{(\sigma^*_i,\hat{\sigma_{-i}})}(y)) \neq \emptyset$, so that by Lemma S1, $s(\sigma^*_i, \hat{\sigma_{-i}}, y) \in g^{-1}(\lambda_{(\sigma^*_i,\hat{\sigma_{-i}})}(y))$. Note that $s(\sigma'_i, \hat{\sigma_{-i}}, y) = s(\sigma^*_i, \hat{\sigma_{-i}}, y) \in \Sigma^*_{-0}(s(\sigma^*_i, \hat{\sigma_{-i}}, y))$. Since $\sigma'_i < \sigma^*_i$, this contradicts $s(\sigma^*_i, \hat{\sigma_{-i}}, y) \in g^{-1}(\lambda_{(\sigma^*_i,\hat{\sigma_{-i}})}(y))$.

Assume now that $\sigma'_i = \sigma^*_i$. Then, $\lambda_{\sigma'}(y) = \lambda_{(\sigma^*_i,\sigma'_{-i})}(y) \in h_i(\sigma^*_i)$ and $g^{-1}(\lambda_{(\sigma^*_i,\sigma'_{-i})}(y)) \neq \emptyset$. *Q.E.D.*

The former two observations imply that for all players i and information sets $h_i \in H_i$ in the original tree, there exists $\sigma^*_i \in \Sigma^*_i$ such that $g^{-1}(h_i(\sigma^*_i)) \neq \emptyset$. Moreover, if $g^{-1}(h_i(\sigma^*_i)) \neq \emptyset$, then every s such that $\mu(s|h_i) > 0$ is such that $s \in g^{-1}(h_i(\sigma^*_i))$.

LEMMA S3: *Fix player i, information set h_i, and σ^*_i such that $g^{-1}(h_i(\sigma^*_i)) = \emptyset$. Then, for all information sets h'_i in the original tree such that $h_i \prec h'_i$, then $g^{-1}(h'_i(\sigma^*_i)) = \emptyset$.*

PROOF: Towards a contradiction, suppose that there exists $y' \in h'_i$ and $\sigma_{-i} \in \Sigma^*_{-\{0,i\}}$ such that $g^{-1}(\lambda_{(\sigma^*_i,\sigma_{-i})}(y')) \neq \emptyset$. Then, $s(\sigma^*_i, \sigma_{-i}, y') \in g^{-1}(\lambda_{(\sigma^*_i,\sigma_{-i})}(y'))$. Let y be the unique node in h_i such that $y \prec y'$. Clearly, $\lambda_{(\sigma^*_i,\sigma_{-i})}(y)$ is on the path of $g(s(\sigma^*_i, \sigma_{-i}, y'))$. Moreover, $s(\sigma^*_i, \sigma_{-i}, y)|_{y'} = (\sigma^*_i, \sigma_{-i})|_{y'} = s(\sigma^*_i, \sigma_{-i}, y')|_{y'}$.

Since $g^{-1}(\lambda_{(\sigma^*_i,\sigma_{-i})}(y)) = \emptyset$, then there exists $\sigma' \neq (\sigma^*_i, \sigma_{-i})$ such that $s(\sigma', y) = s(\sigma^*_i, \sigma_{-i}, y) \in g^{-1}(\lambda_{\sigma'}(y))$. Moreover, $\sigma'|_v = (\sigma^*_i, \sigma_{-i})|_y$. Thus, for every node $\tilde{y}$ such that $s(\sigma^*_i, \sigma_{-i}, y')|_{\tilde{y}} \in (\sigma^*_i, \sigma_{-i})|_{\tilde{y}}$, we have that $s(\sigma^*_i, \sigma_{-i}, y')|_{\tilde{y}} \in \sigma'|_{\tilde{y}}$. Since $g^{-1}(\lambda_{(\sigma^*_i,\sigma_{-i})}(y)) = \emptyset$

[3]This means that we have been following the path of (σ_0, σ'_{-0}) (at least from $\underline{y}_{s(\sigma,y)}$ onwards), even though the recommendation was to play (σ_0, σ_{-0}).

implies that either there is a node that precedes y such that $(\sigma_i^*, \sigma_{-i})|_\cdot \neq \sigma'|_\cdot$ or no such node exists and $\sigma'_{-0} < (\sigma_i^*, \sigma_{-\{0,i\}})$, then it cannot be that $g^{-1}(\lambda_{(\sigma_i^*,\sigma_{-i})}(y')) \neq \emptyset$. *Q.E.D.*

LEMMA S4: *Fix player i, information set h_i, node $y \in h_i$, and $\sigma_i^* \in \Sigma_i^*$ such that $g^{-1}(\lambda_{(\sigma_i^*,\sigma_{-i})}(y)) \neq \emptyset$. Let $s' \in \lambda_{(\sigma_i^*,\sigma_{-i})}(y)$ be such that $g(s(\sigma_i^*, \sigma_{-i}, y)) = s'$. If there is a player j and information set h_j in Γ such that $s' \in h_j(\tilde{\sigma}_j)$ for some $\tilde{\sigma}_j \in \Sigma_j^*$, then*
1. *$(\sigma_i^*, \sigma_{-i})_j = \tilde{\sigma}_j$, and*
2. *$g^{-1}(h_j(\tilde{\sigma}_j)) \neq \emptyset$.*

PROOF: Note that $s' \in h_j(\tilde{\sigma}_j)$ implies that there exists $y_j \in h_j$ such that $s' \in \lambda_{(\tilde{\sigma}_j,\sigma_{-j})}(y_j)$. Then, either $\lambda_{(\tilde{\sigma}_j,\sigma_{-j})}(y_j)$ precedes $\lambda_{(\sigma_i^*,\sigma_{-i})}(y)$, or the opposite holds. In either case, it has to be that $(\sigma_i^*, \sigma_{-i})_j = \tilde{\sigma}_j$.

For the second part, note that $s' \in \lambda_{(\tilde{\sigma}_j,\sigma_{-j})}(y_j)$ and $g^{-1}(s') = s(\tilde{\sigma}_j, \sigma_{-j}, y)$. By construction, y_j is on the path of $s(\tilde{\sigma}_j, \sigma_{-j}, y)$ so that if $y \prec y_j$, then $g^{-1}(\lambda_{(\sigma_i^*,\sigma_{-i})}(y)) = s(\sigma_i^*, \sigma_{-i}, y)$ implies that $g^{-1}(\lambda_{(\tilde{\sigma}_j,\sigma_{-j})}(y_j)) \neq \emptyset$. Suppose then that $y_j \prec y$. Then, $\underline{y}_{s(\sigma_i^*,\sigma_{-i},y)} \preceq y_j$ and by definition $s(\sigma_i^*, \sigma_{-i}, y)|_{\underline{y}_{s(\sigma_i^*,\sigma_{-i},y)}} \in (\sigma_i^*, \sigma_{-i})|_{\underline{y}_{s(\sigma_i^*,\sigma_{-i},y)}}$. Hence, the result follows. *Q.E.D.*

S3.2.2. *Conditional Probability System: Step 1*

We begin by constructing a conditional probability system on $\Sigma'_g = \{A' \subseteq \Sigma' : A' \cap g(S) \neq \emptyset\}$. Let

$$\mu_g(B'|A') = \mu(g^{-1}(B' \cap g(S))|g^{-1}(A' \cap g(S))),$$

for $A' \in \Sigma'_g$. We now check that μ_g indeed defines a conditional probability system on Σ'_g. Take any $A' \in \Sigma'_g$. It is immediate to show that $\mu_g(A'|A') = 1$ and $\mu_g(g(S)|A') = 1$. Consider now a set $A' \in \Sigma'_g$ and any sets $B', C' \subseteq A'$ such that $B' \cap C' = \emptyset$. We have

$$\begin{aligned}
\mu_g(B' \cup C'|A') &= \mu(g^{-1}((B' \cup C') \cap g(S))|g^{-1}(A' \cap g(S))) \\
&= \mu(g^{-1}((B' \cap g(S)) \cup (C' \cap g(S)))|g^{-1}(A' \cap g(S))) \\
&= \mu(g^{-1}(B' \cap g(S)) \cup g^{-1}(C' \cap g(S))|g^{-1}(A' \cap g(S))) \\
&= \mu(g^{-1}(B' \cap g(S))|g^{-1}(A' \cap g(S))) + \mu(g^{-1}(C' \cap g(S))|g^{-1}(A' \cap g(S))) \\
&= \mu_g(B'|A') + \mu_g(C'|A'),
\end{aligned}$$

where the third equality uses the definition of g^{-1} and the fourth equality follows from μ being a CPS and $B' \cap C' = \emptyset$ implying that $g^{-1}(B') \cap g^{-1}(C') = \emptyset$.

Moreover, note that the product rule is satisfied: letting $A' \subseteq B' \subseteq C'$, $C' \in \Sigma_g$, we have

$$\begin{aligned}
\mu_g(A'|C') &= \mu(g^{-1}(A' \cap g(S))|g^{-1}(C' \cap g(S))) \\
&= \mu(g^{-1}(A' \cap g(S))|g^{-1}(B' \cap g(S)))\mu(g^{-1}(B' \cap g(S))|g^{-1}(C' \cap g(S))) \\
&= \mu_g(A'|B')\mu_g(B'|C'),
\end{aligned}$$

where the second equality uses that μ is a conditional probability system and hence satisfies the product rule.

The definition of μ_g ensures that we can write the following. For any σ_i^* and h_i in Γ such that there exists $y \in h_i$ with $g^{-1}(\lambda_{(\sigma_i^*,\sigma_{-i})}(y)) \neq \emptyset$, player i's payoff at $h_i(\sigma_i^*)$ is given by

$$\begin{aligned}
&\sum_{v\in h_i} \sum_{\sigma_{-i}\in\Sigma_0\times\Sigma^*_{-\{0,i\}}} \sum_{s'_{-i}\in\lambda_{(\sigma_i^*,\sigma_{-i})}(y)_{-i}} \mu_g\big(s'_{-i}|h_i\big(\sigma_i^*\big)\big)u_i\big(s_i, s'_{-i}\big) \\
&\quad = \sum_{v\in h_i} \mu_g\Bigg(\bigcup_{\sigma_{-i}\in\Sigma_0\times\Sigma^*_{-\{0,i\}}} \lambda_{(\sigma_i^*,\sigma_{-i})}(y)|h_i\big(\sigma_i^*\big)\Bigg) \\
&\qquad \times \sum_{\sigma_{-i}\in\Sigma_0\times\Sigma^*_{-\{0,i\}}} \sum_{s'_{-i}\in\lambda_{(\sigma_i^*,\sigma_{-i})}(y)_{-i}} \mu_g\Bigg(s_{-i}\Bigg| \bigcup_{\sigma_{-i}\in\Sigma_0\times\Sigma^*_{-\{0,i\}}} \lambda_{(\sigma_i^*,\sigma_{-i})}(y)\Bigg)u_i\big(s_i, s'_{-i}\big) \\
&\quad = \sum_{v\in h_i} \mu\Bigg(g^{-1}\Bigg(\bigcup_{\sigma_{-i}\in\Sigma_0\times\Sigma^*_{-\{0,i\}}} \lambda_{(\sigma_i^*,\sigma_{-i})}(y)\Bigg)\Bigg|g^{-1}\big(h_i\big(\sigma_i^*\big)\big)\Bigg) \\
&\qquad \times \Bigg(\sum_{\sigma_{-i}\in\Sigma_0\times\Sigma^*_{-\{0,i\}}} \sum_{s'_{-i}\in\lambda_{(\sigma_i^*,\sigma_{-i})}(y)_{-i}} \mu\Bigg(g^{-1}\big(s'_{-i}\big)\Bigg|g^{-1}\Bigg(\bigcup_{\sigma_{-i}\in\Sigma_0\times\Sigma^*_{-\{0,i\}}} \lambda_{(\sigma_i^*,\sigma_{-i})}(y)\Bigg)\Bigg)u_i\big(s_i, s'_{-i}\big)\Bigg) \\
&\quad = \sum_{v\in h_i} \mu(v|h_i) \sum_{s_{-i}\in S_{-i}(y)} \mu(s_{-i}|v)u_i(s_i, s_{-i}),
\end{aligned} \tag{S9}$$

where the first equality uses the product rule for conditional probability systems, the second equality uses the definition of μ, and the last equality uses the properties of g^{-1} discussed above. This guarantees that the obedient strategy profile satisfies sequential rationality at $h_i(\sigma_i^*)$.

The conditional probability system μ' in the new PBE coincides with μ_g on Σ'_g. Moreover, beliefs at information sets h that satisfy $h \in \sigma'_g$ are defined using μ'.

S3.2.3. *Conditional Probability System: Step 2*

Fix player i and suppose there exists an information set $h_i \in H_i$ and $\sigma_i^* \in \Sigma_i^*$ such that $g^{-1}(h_i(\sigma_i^*)) = \emptyset$. We know there exists $\sigma'_i \in \Sigma_i^*$ such that $g^{-1}(h_i(\sigma'_i)) \neq \emptyset$ (see Lemma S1 and Lemma S2). Moreover, for each $y \in h_i$ and $s^* \in \lambda_{(\sigma_i^*,\sigma_{-i})}(y)$, there exists a strategy profile $s' \in \lambda_{(\sigma'_i,\sigma_{-i})}(y)$ that coincides everywhere with s^* except that, at the root, Nature recommends $(\sigma'_i, s^*_{-i}(\emptyset))$ and player i plays $s_i^*(\tilde{h}_i(\sigma_i^*))$ at each $\tilde{h}_i(\sigma'_i)$. Let $\eta_{\sigma_i^*\to\sigma'_i}$ denote the mapping from $h_i(\sigma_i^*)$ to $h_i(\sigma'_i)$ defined in this way.

Define then the conditional probability system $\mu_{h_i(\sigma_i^*)}$ on $2^{h_i(\sigma_i^*)} \setminus \{\emptyset\}$ to be given by, for any $B' \subseteq A' \subseteq h_i(\sigma_i^*)$ such that $\eta_{\sigma_i^*\to\sigma'_i}(A') \cap g(S) \neq \emptyset$,

$$\mu_{h_i(\sigma_i^*)}\big(B'|A'\big) = \mu_g\big(\eta_{\sigma_i^*\to\sigma'_i}\big(B'\big)|\eta_{\sigma_i^*\to\sigma'_i}\big(A'\big)\big). \tag{S10}$$

In particular, this implies that $\mu_{h_i(\sigma_i^*)}(s'|\lambda_{(\sigma_i^*,\sigma_{-i})}(y)) = \mu_g(\eta_{\sigma_i^*\to\sigma'_i}(s')|\lambda_{(\sigma'_i,\sigma_{-i})}(y))$ and $\mu_{h_i(\sigma_i^*)}(s'|h_i(\sigma_i^*)) = \mu_g(\eta_{\sigma_i^*\to\sigma'_i}(s')|h_i(\sigma'_i))$. Recalling that σ'_i satisfies that $(\sigma_i^*, \sigma_{-i})|_y = (\sigma'_i, \sigma_{-i})|_y$ for all $\sigma_{-i} \in \Sigma_0 \times \Sigma^*_{-\{0,i\}}$, this (together with equation (S9)) implies that player i finds it optimal to play $\sigma_i^*(h_i)$ at $h_i(\sigma_i^*)$ under these beliefs.

Moreover, it follows from Lemma S4 that if $\mu_{h_i(\sigma_i^*)}(s'|h_i(\sigma_i^*)) > 0$ and $s' \in h_j(\sigma_j^*)$, then $g^{-1}(h_j(\sigma_j^*)) \neq \emptyset$.[4] Thus, $\mu_{h_i(\sigma_i^*)}$ does not place positive probability on $\{h_j(\sigma_j) : g^{-1}(h_j(\sigma_j)) = \emptyset\}$. Hence the supports of $\mu_{h_j(\sigma_j)}$ and $\mu_{h_i(\sigma_i)}$ are disjoint for all $\sigma_i \in \Sigma_i^*$, $\sigma_j \in \Sigma_j^*$, whenever $i \neq j$.

Given this, for any collection of $C \subseteq \{(i, \sigma_i, h_i) : i \in \{1, \dots, N\}, \sigma_i \in \Sigma_i^*, g^{-1}(h_i(\sigma_i)) = \emptyset\}$, we can define μ' so that

$$\mu'\left(h_i(\sigma_i)\middle| \bigcup_{(l,\sigma_l,h_l)\in C} h_l(\sigma_l)\right) = 0,$$

whenever $i \neq \min\{j : (j, \sigma_j, h_j) \in C\}$.

Now fix i and order the elements of $\Sigma_i^* = \{\sigma_{i1}, \dots, \sigma_{iM_i}\}$. For any collection $\{(l, h_i) : g^{-1}(h_i(\sigma_{il})) = \emptyset\}$, we can define μ' so that

$$\mu'\left(h_i(\sigma_{il})\middle| \bigcup_{(l',h_i')\in C} h_i'(\sigma_{il'})\right) = 0,$$

whenever $l \neq \min\{l' : (l', h_i') \in C\}$.

Finally, for each $l \in \{1, \dots, M_i\}$, we need to specify $\mu'(\cdot|\bigcup_{h_i\in\tilde{H}_i} h_i(\sigma_{il}))$ for different collections $\tilde{H}_i$ of information sets of player i such that $g^{-1}(h_i(\sigma_{il})) = \emptyset$. It suffices to do so for pairs h_i, h_i' such that $h_i \not\prec h_i', h_i' \not\prec h_i$. Towards this end, let $E(h_i) = \{h_i' \in H_i : h_i \preceq h_i' \wedge h_i' \preceq h_i\}$. Note that if $E(h_i) \cap E(h_i') = \emptyset$, then $h_i \not\prec h_i', h_i' \not\prec h_i$. Label the sets $E_1, \dots, E_{N_i}$ so that $\bigcup_{l=1}^{N_i} E_l = H_i$ and $E_l \cap E_j = \emptyset$ if $l \neq j$. Then, for any $l \in \{1, \dots, M_i\}$ and collection $\tilde{H}_i$ of information sets of player i, define

$$\mu'\left(h_i(\sigma_{il})\middle| \bigcup_{h\in\tilde{H}_i} h(\sigma_{il})\right) = 0,$$

whenever there exists $h \in \tilde{H}_i$ such that $h \in E_j$, $h_i \in E_l$ and $j < l$.

Defining μ' to coincide with $\mu_{h_i(\sigma_i^*)}$ on $h_i(\sigma_i^*) : g^{-1}(h_i(\sigma_i^*)) = \emptyset$ completes the construction.

S4. EXTENDING A CONDITIONAL PROBABILITY SYSTEM

Suppose that X_1 and X_2 are disjoint finite sets, $\tilde{\mu}(\cdot \mid \cdot)$ is a conditional probability system over X_1, and ν is a probability measure satisfying $\nu(X_2) = 1$ and $\nu(x) > 0$ for all $x \in X_2$.

We will construct a conditional probability system $\mu(\cdot \mid \cdot)$ over the union $X = X_1 \cup X_2$ to satisfy

1. $\mu(\cdot \mid C) = \tilde{\mu}(\cdot \mid C)$ for all $C \subset X_1$,
2. $\mu(B \mid X) = \nu(B)$ for all $B \subset X_2$.

We define, for all $C \subset X_1$,

$$\mu(\cdot|C) = \tilde{\mu}(\cdot|C),$$

[4]This does not say that s' has positive probability under $\mu_g(|h_j(\sigma_j^*))$; after all, if y_j is the node in h_j such that $s' \in \lambda_{\sigma_i^*,\sigma_j^*,\cdot}(v_j)$, then $g^{-1}(\lambda_{\sigma_i^*,\sigma_j^*,\cdot}(v_j)) = \emptyset$.

and for $C \subset X_2$, we define

$$\mu(B \mid C) = \frac{\nu(B \cap C)}{\nu(C)}.$$

Finally, for C such that $C \cap X_1, C \cap X_2 \neq \emptyset$, we define

$$\mu(B \cap X_2 \mid C) = 0, \tag{S11}$$

and

$$\mu(B \cap X_1 \mid C) = \tilde{\mu}(B \cap X_1 | C \cap X_1).$$

These define the probability measures $\mu(\cdot \mid C)$ for any $C \subset X$. We now verify that μ is a conditional probability system.

It is immediate that $\mu(C \mid C) = 1$ for all $C \subset X$. Suppose $A \subset B \subset C \subset X$. We require

$$\mu(A \mid C) = \mu(B \mid C)\mu(A \mid B). \tag{S12}$$

Equation (S12) holds by construction for $C \subset X_1$ since $\tilde{\mu}$ is a conditional probability system over X_1.

Suppose that $C \subset X_2$; then $A \subset B \subset X_2$ and the product rule holds by definition since we can use Bayes' rule with $\nu(\cdot)$.

Consider now $C \cap X_1, C \cap X_2 \neq \emptyset$. If $B \subset X_1$, then $B \subset C \cap X_1$ and $A \subset X_1$ and we have

$$\mu(A \mid B) = \tilde{\mu}(A \mid B),$$
$$\mu(B \mid C) = \tilde{\mu}(B \mid C \cap X_1).$$

Then,

$$\mu(A \mid B)\mu(B \mid C) = \tilde{\mu}(A \mid B)\tilde{\mu}(B \mid C \cap X_1) = \tilde{\mu}(A \mid C \cap X_1) = \mu(A \mid C),$$

where the second equality follows from the chain rule applied to $\tilde{\mu}$ and the final one is by definition of μ.

Suppose now that $B \subset X_2$, so that $B \subset C \cap X_2$. Then, Eq. (S11) implies that $\mu(A \mid C) = \mu(B \mid C) = 0$, so that Eq. (S12) holds automatically.

Finally, suppose that $B \cap X_1, B \cap X_2 \neq \emptyset$. If $A \subset X_1$, then $A \subset B \cap X_1 \subset C \cap X_1$ and

$$\mu(A \mid C) = \tilde{\mu}(A \mid C \cap X_1),$$
$$\mu(A \mid B) = \tilde{\mu}(A \mid B \cap X_1),$$
$$\mu(B \mid C) = \tilde{\mu}(B \cap X_1 \mid C \cap X_1).$$

Since the chain rule holds for $\tilde{\mu}$, we have that Eq. (S12) holds for μ.

On the other hand, if $A \subset X_2$, we have that $\mu(A \mid C) = \mu(A \mid B) = 0$ and Eq. (S12) holds automatically.

Finally, suppose that $A \cap X_1, A \cap X_2 \neq \emptyset$. Then,

$$\mu(A \mid C) = \tilde{\mu}(A \cap X_1 \mid C \cap X_1),$$
$$\mu(A \mid B) = \tilde{\mu}(A \cap X_1 \mid B \cap X_1),$$
$$\mu(B \mid C) = \tilde{\mu}(B \cap X_1 \mid C \cap X_1).$$

Since the chain rule holds for $\tilde{\mu}$, we have that Eq. (S12) holds for μ.

S5. THEOREM 4: AUXILIARY LEMMA

Theorem 4 implicitly uses the following result:

LEMMA S5: *Let Γ be an admissible extensive form. Let $z \in Z$ denote a terminal node and $\gamma(z) = (\theta_z, a_z)$ the outcome associated with it. Then, for each $i \in N$, there is a node on that path that is decisive for $a_{z,i}$.*

PROOF: Let $|\cdot|$ denote the length of a node in the extensive form. Let $z_N = \arg\max\{|y| : y \prec z \wedge \gamma_A(y) \neq a_z\}$ and recursively for $l = 1, \dots, N-1$, $z_{N-l} = \arg\max\{|y| : y \prec z_{N-(l-1)} \wedge \gamma_A(y) \neq \gamma_A(z_{N-(l-1)})\}$. Since Γ is admissible, $\{y \in V : y \prec z \wedge \gamma_A(y) \neq a_z\} \neq \emptyset$ since $\gamma_A(\emptyset) = A$ and hence z_N is well-defined. Moreover, admissibility implies that $\gamma_A(z_N) = A_{\iota(z_N)} \times \{a_{z,-\iota(z_N)}\}$. By definition, $a_z \in \gamma(z_N) = \bigcup_{z':z_N \prec z} \gamma_A(z')$ and $z_N \prec z \in Z$. Moreover, suppose there is $i, j, i \neq j$ such that $\gamma_i(z_N) \neq a_{z,i}$ and $\gamma_j(z_N) \neq a_{z,j}$. Without loss of generality, let $j \neq \iota(z_N)$. No delegation implies that for all $m \in M(z_N)$, $\gamma_j(z_N, m) = \gamma_j(z_N) \neq a_{z,j}$. In particular, if $(z_N, m^*) \preceq z$, then $\gamma_j(z_N, m^*) \neq a_{z,j}$, a contradiction to the definition of z_N. Hence, no such j can exist. Then, for $j \neq \iota(z_N)$, $\gamma_j(z_N) = \{a_{z,j}\}$. Finally, no partial commitments implies that $\gamma_{\iota(z_N)}(v_N) = A_{\iota(z_N)}$.

Inductively, one can show that $\gamma_A(z_i) = \times_{j=i}^{N} A_{\iota(z_j)} \times \times_{k \in N \setminus \{\iota(z_i), \dots, \iota(z_N)\}} \{a_{z,k}\}$ and $\gamma_{\iota(z_i)}(z_i, m^*) = a^v_{\iota(z_i)}$ for $(z_i, m^*) \preceq z$. *Q.E.D.*

S6. SELF-CONTAINED COORDINATED EQUILIBRIA CONSISTENT WITH A GIVEN PRIOR

For any given θ, there exists a Nash equilibrium of the simultaneous move game in which there is complete information that the state is θ. Let $\hat{\alpha}$ be the associated joint distribution of action profiles and consider the outcome $(\theta, \hat{\alpha})$, that is, with probability 1, the state is θ and the joint distribution of action profiles is $\hat{\alpha}$. This is a coordinated equilibrium outcome. In particular, for each action profile a in the support of $\hat{\alpha}$, we construct a plan in which the players move in an arbitrary order and each player i is recommended to play a_i independent of previous moves. We conduct a mixture in which θ has probability 1 and the probability of the plan associated with profile a is exactly the probability $\hat{\alpha}$ assigns to a. Because $\hat{\alpha}$ is a Nash equilibrium of the θ-complete-information simultaneous move game, the obedience constraints will be satisfied.

Indeed, this coordinated equilibrium is self-contained. The actions in the support of $\hat{\alpha}$ belong to C^1 because they are part of a coordinated equilibrium. And since the plans used in the coordinated equilibrium only recommended these same actions, they will never be eliminated at any subsequent step of the procedure leading to C. Thus, for every θ, there exists a self-contained coordinated equilibrium assigning probability 1 to θ.

By an appropriate mixing of these, for any marginal distribution over states $\bar{\pi}$, we can obtain a self-contained coordinated equilibrium whose marginal over states is $\bar{\pi}$.[5] We can use this fact to modify the proof of Theorem 2 as follows. Let π be a self-contained coordinated equilibrium with marginal over states equal to a given prior $\tilde{\pi}$. Let ψ be defined as in the proof. It is ε-close to π. Take $\bar{\psi}$ to be a self-contained coordinated equilibrium whose marginal distribution over Θ is $\tilde{\pi} + (\tilde{\pi} - \text{marg}_\Theta \psi)$ (because $\tilde{\pi}$ has full support, for ε small enough this vector belongs to the simplex), and consider the self-contained coordinated equilibrium given by $\tilde{\psi} = \frac{1}{2}\psi + \frac{1}{2}\bar{\psi}$. Its marginal over states equals $\tilde{\pi}$. Finally, set $\varpi = (1-\varepsilon)\pi + \varepsilon\tilde{\psi}$. The mixture ϖ is a self-contained coordinated

[5]In particular, we can do this for $\bar{\pi}$ the prior in the base game.

equilibrium that is within ε distance of π and has marginal over states equal to $\tilde{\pi}$. For every action $a_i \in C_i$, it assigns positive probability to plans with a_i on the obedient path and thus the obedient strategy profile is a self-contained Bayesian Nash equilibrium of $\Gamma(\varpi)$.

REFERENCES

MYERSON, R. B. (1986): "Multistage Games With Communication," *Econometrica*, 54, 323–358. [3]
PETERS, M. (2015): "Reciprocal Contracting," *Journal of Economic Theory*, 158, 102–126. [1]

Co-editor Joel Sobel handled this manuscript.

Manuscript received 26 April, 2019; final version accepted 12 May, 2020; available online 12 May, 2020.