

Online Appendix

Sequential Preference Revelation in Incomplete Information Settings

James Schummer and Rodrigo A. Velez*

1 Proof of Theorem 2

Theorem 2 is implied by the following result.

Theorem. *Let f be a strategy-proof, non-bossy SCE, and fix a reporting order $\Lambda \in \Delta(\Pi)$. Suppose that at least one of the following conditions holds.*

1. *The prior has Cartesian support ($\mu \in \mathcal{M}_{\text{Cartesian}}$) and Λ is deterministic.*
2. *The prior has symmetric Cartesian support ($\mu \in \mathcal{M}_{\text{symm-Cartesian}}$) and f is weakly anonymous.*

Then equilibria are preserved under deviations to truthful behavior: for any $(\sigma, \beta) \in SE(\Gamma(\Lambda, f), \mathcal{U}^N, \mu)$,

- (i) *for each $i \in N$, τ_i is sequentially rational for i with respect to σ_{-i} and β_i ;*
- (ii) *for each $S \subseteq N$, there is a belief system β' such that*

$$((\sigma_{-S}, \tau_S), \beta') \in SE(\beta'(\Lambda, f), \mathcal{U}^N, \mu).$$

Proof. Let f be strategy-proof and non-bossy, $\Lambda \in \Delta(\Pi)$, and $\mu \in \mathcal{M}_{\text{Cartesian}}$. Since the conditions of the two theorems are the same, we refer to arguments

*Schummer: Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston IL 60208, schummer@kellogg.northwestern.edu. Velez: Department of Economics, Texas A&M University, College Station, TX 77843, rvelezca@tamu.edu.

made in the proof of Theorem 1. In particular all numbered equations referenced below appear in the paper.

Fix an equilibrium $(\sigma, \beta) \in \text{SE}(\Gamma(\Lambda, f), \mathcal{U}^N, \mu)$. The first claim of the theorem—sequential rationality of τ_i w.r.t. (σ, β) —can be derived from inequalities established in the proof of Theorem 1. Specifically, consider the case of Condition 1 (the prior has Cartesian support and Λ is deterministic). The proof establishes the equality of (7) and (8) (see (12)). That is, the expected payoff from a truthful report is equal to the expected payoff from an equilibrium report, conditioning on the agents' interim belief. Since the equilibrium report is sequentially rational, so is a truthful report.

In the case of Condition 2 ($\mu \in \mathcal{M}_{\text{symm-Cartesian}}$, weakly anonymous f), the analogous arguments are made via (7'), (8'), and (12').

To prove the second claim of the theorem, it suffices to prove the singleton case $S \equiv \{i\}$. Repeated application of this statement proves the general result for arbitrary S . Let $\tilde{\sigma} \equiv (\sigma_{-i}, \tau_i)$. First we construct a belief system γ and demonstrate its consistency. Lastly we show $(\tilde{\sigma}, \gamma) \in \text{SE}(\Gamma(\Lambda, f), \mathcal{U}^N, \mu)$.

Consistency. Let (ζ, β^ζ) be an assessment where ζ is an arbitrary profile with full support and β^ζ is the unique belief system obtained by Bayesian updating given ζ , Λ , and μ . For any $\varepsilon > 0$, let $(\zeta, \beta^\varepsilon)$ be the assessment where σ^ε is the (full support) strategy profile $\sigma^\varepsilon \equiv (1 - \varepsilon)\tilde{\sigma} + \varepsilon\zeta$ and β^ε is the unique belief system obtained by Bayesian updating, given σ^ε , Λ , and μ . Clearly $\sigma^\varepsilon \rightarrow \tilde{\sigma}$. More specifically, for each $i \in N$, each $u_i \in \mathcal{U}$, each $h_t \in H^i$, and each $v_i \in \mathcal{U}$, we have $\sigma^\varepsilon \langle u_i, h_t \rangle (v_i) \rightarrow \tilde{\sigma} \langle u_i, h_t \rangle (v_i)$ as $\varepsilon \rightarrow 0$.

We define γ to be the Bayesian update of $\tilde{\sigma}$ (when well defined) or equal to β^ε (otherwise). That is, fix $i \in N$ and any admissible $u_i \in \mathcal{U}$ (i.e. occurring with positive probability under μ).

- For each $h_t \in H^i$ that occurs with positive probability given $\tilde{\sigma}$ and Λ , and for each admissible (π, u_{-i}) , let $\gamma_i \langle u_i, h_t \rangle (\pi, u_{-i})$ be defined by Bayesian updating given $\tilde{\sigma}$ and Λ .
- For each $h_t \in H^i$ that occurs with zero probability given $\tilde{\sigma}$ and Λ , and for each admissible (π, u_{-i}) , let $\gamma_i \langle u_i, h_t \rangle (\pi, u_{-i}) \equiv \beta^\zeta \langle u_i, h_t \rangle (\pi, u_{-i})$.

Using Bayes' rule, one can write an explicit expression of β^ε in terms of ε , $\tilde{\sigma}$, ζ , Λ , and μ . We omit this expression since it is easy to see that $\beta^\varepsilon \rightarrow \gamma$; specifically, for each $i \in N$, each $u_i \in \mathcal{U}$, each $h_t \in H^i$, each $\pi \in \Pi$, and each $v_{-i} \in \mathcal{U}^{N \setminus \{i\}}$,

$\beta_i^\varepsilon \langle u_i, h_t \rangle(\pi, v_{-i}) \rightarrow \gamma_i \langle u_i, h_t \rangle(\pi, v_{-i})$ as $\varepsilon \rightarrow 0$. Thus $(\tilde{\sigma}, \gamma)$ is a consistent assessment for $\langle \Gamma(\Lambda, f), \mathcal{U}^N, \mu \rangle$.

Sequential rationality. We use the notation—from Case 2 of the proof of Theorem 1—where $f(h_t, v_{\pi'(t+1, \dots, n)} | \pi')$ represents the outcome of f when the ordered reports in h_t are made according to the agents' order under π' . We also refer to Equations (2')–(13') in order to prove various claims. For the simpler case that Λ is deterministic, the analogous equations from Case 1 apply.

We show that $\tilde{\sigma}$ is sequentially rational for beliefs γ . That is, for each Agent j , $\tilde{\sigma}_j$ prescribes a report that maximizes j 's expected payoff after any history feasible for j , given $\tilde{\sigma}_{-j}$ and γ_j .

Case $j = i$ ($\tilde{\sigma}_j = \tau_i$). For Agent i , we show the sequential rationality of truth-telling using (2') which states that, regardless of the history, continuation strategies under σ are welfare-equivalent to truthful ones. Since $\tilde{\sigma}_{-i} = \sigma_{-i}$ the result follows.

To formalize this, fix any $t \in \{1, \dots, n\}$, $\pi \in \text{supp}(\Lambda)$ with $\pi(t) = i$, $h_{t-1} \in H_i$, and $u \in \text{supp}(\mu)$, and recall that $\tilde{\sigma}_i = \tau_i$. For any $v'_i \in \text{supp}(\sigma \langle h_{t-1}, u_i \rangle)$, consider the two t -period histories (h_{t-1}, v'_i) and (h_{t-1}, u_i) . For any two σ -continuations of those two histories given π , namely for any¹

- $v'_{\pi(t+1, \dots, n)} \in \mathcal{U}^{\pi(t+1, \dots, n)}$ with $\sigma(v'_{\pi(t+1, \dots, n)} | (h_{t-1}, v'_i), \pi, u_{\pi(t+1, \dots, n)}) > 0$, and
- $v''_{\pi(t+1, \dots, n)} \in \mathcal{U}^{\pi(t+1, \dots, n)}$ with $\sigma(v''_{\pi(t+1, \dots, n)} | (h_{t-1}, u_i), \pi, u_{\pi(t+1, \dots, n)}) > 0$,

applying (2') to histories h_{t-1} and (h_{t-1}, u_i) respectively yields

$$\begin{aligned} u(f(h_{t-1}, v'_i, v'_{\pi(t+1, \dots, n)} | \pi)) &= u(f(h_{t-1}, u_{\pi(t, \dots, n)} | \pi)), \text{ and} \\ u(f(h_{t-1}, u_i, v''_{\pi(t+1, \dots, n)} | \pi)) &= u(f(h_{t-1}, u_i, u_{\pi(t+1, \dots, n)} | \pi)). \end{aligned} \quad (**)$$

Observe that the two RHS's are equivalent and thus i receives the same payoff from reporting u_i as from reporting v'_i . Since reporting v'_i maximizes i 's expected payoff after h_{t-1} given u_i and σ_{-i} , reporting u_i maximizes i 's expected payoff after h_{t-1} given u_i and $\tilde{\sigma}_{-i} = \sigma_{-i}$, proving sequential rationality. More generally, however, *all* agents receive the same payoff after i 's deviation from σ_i to truthful τ_i . This is true ex post of *any* realization of π (with $\pi(t) = i$), which is relevant in the next case.

¹If $t = n$ these subprofiles are null lists, and the proof simplifies.

Case $j \neq i$ ($\tilde{\sigma}_j = \sigma_j$). The intuition behind this case is that, since Equation (**) implies that i 's deviation to truth-telling does not change continuation payoffs, the incentive compatibility conditions of the original equilibrium (σ) are preserved.

Fix any $t \in \{1, \dots, n\}$, $\pi \in \text{supp}(\Lambda)$ with $\pi(t) = j$, $h_{t-1} \in H_i$, $u \in \text{supp}(\mu)$, and $v_j \in \text{supp}(\sigma_j \langle h_{t-1}, u_j \rangle)$. We wish to show that v_j maximizes j 's expected payoff, given $\tilde{\sigma}_{-j}$ and γ .

We begin with the more difficult subcase that $\pi^{-1}(i) > t$, so i acts after j . Denote the (possibly empty) sets of agents who act between j and i and after i as

$$B = \{k \in N : \pi^{-1}(j) < \pi^{-1}(k) < \pi^{-1}(i)\}$$

$$A = \{k \in N : \pi^{-1}(i) < \pi^{-1}(k)\}$$

Fix a (deviation) report $v'_j \in \mathcal{U}$. We shall compare payoffs obtained under four profiles of reports,

$$(h_{t-1}, v_j, v_B, v_i, v_A)$$

$$(h_{t-1}, v_j, v_B, u_i, w_A)$$

$$(h_{t-1}, v'_j, v'_B, v'_i, v'_A)$$

$$(h_{t-1}, v'_j, v'_B, u_i, w'_A)$$

where the various subprofiles for B , i , and A satisfy

$$v_B \in \text{supp}(\sigma_B \langle (h_{t-1}, v_j), u_B \rangle)$$

$$v'_B \in \text{supp}(\sigma_B \langle (h_{t-1}, v'_j), u_B \rangle)$$

$$v_i \in \text{supp}(\sigma_i \langle (h_{t-1}, v_j, v_B), u_i \rangle)$$

$$v'_i \in \text{supp}(\sigma_i \langle (h_{t-1}, v'_j, v'_B), u_i \rangle)$$

$$v_A \in \text{supp}(\sigma_A \langle (h_{t-1}, v_j, v_B, v_i), u_A \rangle)$$

$$w_A \in \text{supp}(\sigma_A \langle (h_{t-1}, v_j, v_B, u_i), u_A \rangle)$$

$$v'_A \in \text{supp}(\sigma_A \langle (h_{t-1}, v'_j, v'_B, v'_i), u_A \rangle)$$

$$w'_A \in \text{supp}(\sigma_A \langle (h_{t-1}, v'_j, v'_B, u_i), u_A \rangle)$$

Equation (**) implies the following two equalities.

$$\begin{aligned} u(f(h_{t-1}, v_j, v_B, v_i, v_A|\pi)) &= u(f(h_{t-1}, v_j, v_B, u_i, w_A|\pi)) \\ u(f(h_{t-1}, v'_j, v'_B, v'_i, v'_A|\pi)) &= u(f(h_{t-1}, v'_j, v'_B, u_i, w'_A|\pi)) \end{aligned}$$

Sequential rationality of σ implies

$$u_j(f(h_{t-1}, v_j, v_B, v_i, v_A|\pi)) \geq u_j(f(h_{t-1}, v'_j, v'_B, v'_i, v'_A|\pi)).$$

Thus

$$u_j(f(h_{t-1}, v_j, v_B, u_i, w_A|\pi)) \geq u_j(f(h_{t-1}, v'_j, v'_B, u_i, w'_A|\pi))$$

implying a report of v_j is at least as good as any other v'_j , for each such π .

In case that $\pi^{-1}(i) < t$, (i acts before j), the result follows immediately from (2'). Those equations state that, when all remaining agents are playing according to σ , the agents receive payoffs *as if* each agent is acting truthfully. With strategy-proofness of f , the result follows. \square

2 Forward induction

We discuss how forward induction (refinement) arguments might apply to sequential revelation games in a Bayesian environment, and two ways in which the (non-truthful) equilibria of Examples 1 and 2 are robust to them.

In a sequential, direct revelation game, an agent's report not only informs the mechanism designer about the agent's preferences, but also plays a "signaling role," leaking the agent's private information to subsequent players. This signaling role is what drives the non-truthful equilibrium outcomes we construct in Examples 1 and 2. Indeed under *complete* information, this signaling role disappears (e.g. when the reporting order is deterministic, backwards induction precludes such equilibria in accordance with Theorem 1).

A relevant question is whether this signaling role effect—leading to non-truthful outcomes—can be easily dismissed via forward induction arguments that have been made in the signaling game literature. Essentially, these arguments reject equilibria that are sustained through the Sender's fear that, by deviating to a particular action/message, the Receiver will form beliefs that are "unreasonable," in the sense that the Sender's deviation itself should serve as evidence against those beliefs.

We now argue that the answer to the above question is no: the equilibrium phenomenon we construct in our two examples is robust to such forward induction arguments. We make this argument in two ways, in both cases focusing on the equilibrium of Example 1 for the formal arguments. First, we explain why this equilibrium is robust to the ideas that underlie the definitions of the intuitive criterion, D1, etc., which are typically applied to two-player sender-receiver games. This is necessarily done informally since our model involves more than two senders and receivers, and it is beyond our scope to extend the classic definitions of these various refinements to our class of games. Second and more formally, we show that the equilibrium satisfies the general *forward induction* principle defined by [Govindan and Wilson \(2009\)](#) that applies to all games with perfect recall. In the spirit of the classic literature mentioned earlier, their concept formalizes the idea of testing the plausibility of beliefs based on the rationality of actions.

2.1 Forward induction via standard refinement concepts

Signaling games refinements have been studied mainly for Sender-Receiver games. Here all uncertainty is resolved for one agent, who sends a (perhaps costly) message to an uninformed Receiver, who then takes a payoff determinant action. Our games are more complex in that we allow for $n \geq 2$ agents, all possessing private information and sending messages. Despite the differences, the structure of the sequential equilibrium constructed in Example 1 allows us to (successfully) subject it to the same type of vetting that standard signaling game refinements provide.

To explain, consider the sequential equilibrium (σ, β) constructed in Example 1, where Agents 1–4 publicly announce their peaks in order of their index. Observe that the non-truthfulness of outcomes in this equilibrium entirely hinges on the beliefs of Agent 2, which are influenced by the report of Agent 1. At the same time, Agents 3 and 4 are always truthful, which is sequentially rational behavior regardless of their beliefs. Hence Agent 1 needs only to consider his own beliefs (a Bayesian update of the prior), and the impact of his message on Agent 2’s beliefs when choosing his message. If we imagine fixing the (equilibrium) truthful behavior of Agents 3 and 4, we can view the resulting game being played between Agents 1 and 2 as something closer to a standard, 2-player Sender-Receiver game, and then evaluate their equilibrium behavior using standard refinement concepts for Sender-Receiver games. We

now do this somewhat more formally, showing that the equilibrium behavior of Agents 1 and 2 is compatible with the D1 criterion of [Cho and Kreps \(1987\)](#), which is among the most prominent signaling game refinements used in the literature.²

The idea behind the D1 criterion, as applied to Agents 1 and 2 in our example, is to require Agent 2's beliefs to rule out the possibility that Agent 1, having some type p_1 , could make some out of equilibrium report (say a report of 3) if some other type p'_1 would have benefitted "more often" (with respect to the set of Agent 2's best responses to the report) by making that report. Since the only (relevant) out of equilibrium report for Agent 1 is indeed "3," we restrict attention to that one.³

Formally, for any possible type $p_1 \in \{2, 2.5, 3\}$ of Agent 1, let $D(p_1)$ denote the set of mixed behavioral strategies of Agent 2 that are (i) best responses for *some* beliefs over the other agents' types (i.e. over $\{u_{-2}, v_{-2}, w_{-2}\}$) given a report of 3 by Agent 1, and (ii) result in a strictly higher expected payoff for Agent 1 than her expected equilibrium payoff (when her type is p_1 and she reported 3).⁴ Similarly, let $D^0(p_1)$ be the same set, replacing *strictly higher* with *weakly higher*.

The D1 criterion requires that if $D^0(p_1) \subsetneq D(p'_1)$, then Agent 2's beliefs following Agent 1's report of 3 must put zero weight on type p_1 . The only out of equilibrium report under (σ, β) is Agent 1's report of 3, after which Agent 2 believes with certainty that $p_1 = 2$. Hence the equilibrium would fail the D1 criterion precisely when $D^0(2) \subsetneq D(p'_1)$ for some $p'_1 \in \{2.5, 3\}$.

It is easy to show that this condition is not true, i.e. the equilibrium satisfies the D1 criterion. In fact consider Agent 2's equilibrium response to Agent 1's (out of equilibrium) report of 3, $\sigma_2(3)$. By definition this is a best response with respect to Agent 2's beliefs. Imagine that Agent 2 uses this strategy. Then, when $p_1 = 2$ and Agent 1 reports a peak of 3, she receives a payoff identical to her

²As a consequence, our equilibrium is thus also consistent with weaker notions such as the Intuitive Criterion, Divinity, etc.

³Note that "0" and "1" are also out of equilibrium reports for Agent 1. However, these reports are trivial cases in that none of Agent 1's admissible types could benefit from making such reports. That is, there is no final history, following a report of "0" or "1" by Agent 1, in which she receives more than 1 unit. Moreover, each possible type of Agent 1 weakly prefers the equilibrium outcome to receiving 1 or fewer units. Thus the requirements of D1, defined below, are trivially satisfied following those two reports. Agent 1's report of "3" is the crucial case, which we handle below.

⁴See Fudenberg and Tirole's book, *Game Theory*, page 452.

equilibrium payoff (see Table 4); hence $\sigma_2(3) \in D^0(2)$. Yet when $p_1 \in \{2.5, 3\}$ and Agent 1 reports a peak of 3, she receives a payoff *strictly lower* than her equilibrium payoff (see Table 4); hence $\sigma_2(3) \notin D(2.5) \cup D(3)$. That is,

$$D^0(2) \not\subseteq D(2.5) \cup D(3)$$

meaning that Agent 2's belief that $p_1 = 2$ passes the test of the D1 criterion.

2.2 Forward induction a la Govindan-Wilson

The arguments above were made by taking the truthful behavior of Agents 3 and 4 as given, and then extending the ideas behind 2-person signaling game equilibrium refinements to our setting. We were able to do this in a fairly natural way because, in the specific equilibrium we consider, only Agent 2's beliefs are critical to the analysis, while the truthful behavior of Agents 3 and 4 is robust to any specification of beliefs. A more formal approach, well beyond the scope of our work, would be to first establish a principle that extends the ideas beyond these refinement concepts to a more general class of games including ours, and that would apply to *any* sequential equilibrium. Fortunately for us, [Govindan and Wilson \(2009\)](#) do precisely this, providing a general definition of forward induction that applies to a general class of games with perfect recall. In what follows we show that the equilibria in our examples also satisfy their definition.

[Govindan and Wilson \(2009\)](#) formalize the idea of testing the plausibility of beliefs based on the rationality of actions that have led to the current information set. While their definition of sequential equilibrium is analogous to that in [Kreps and Wilson \(1982\)](#), there are a few technical issues that should be clarified. First, they allow for games in which a player might take multiple actions (i.e. make multiple moves in the sequential game), while in our case a player only moves once. For brevity, we describe the definitions in [Govindan and Wilson \(2009\)](#) under the restriction to our class of move-once games, which allows for simpler statements.⁵

Second, they describe equilibria using beliefs over other agents' *strategies* rather than over others' *types*; in their words, "it is only from a belief specified as a conditional distribution over strategies that one can verify whether

⁵In brief, we avoid reference to *weakly* sequential equilibrium, which applies when an agent has an information set that can be reached only through his own earlier action.

a player’s belief recognizes the rationality of others’ strategies.” More formally, they define what we will henceforth call a “GW sequential equilibrium” in terms of a strategies-beliefs pair, $(\hat{\sigma}, \hat{\gamma})$, where $\hat{\gamma}_i$ represents Agent i ’s conditional beliefs over nature’s moves and over other agents’ pure strategies (functions mapping histories at which the agent acts into actions). As in [Kreps and Wilson \(1982\)](#), these authors require consistency: there is a sequence of full-support behavior strategies σ^k and a corresponding sequence of full-support equivalent normal-form strategies φ^k , such that as $k \rightarrow \infty$, $\sigma^k \rightarrow \sigma$, and $\hat{\gamma}$ is the limit of the conditional distributions obtained by Bayes’ rule from φ^k and the prior.

Third, [Govindan and Wilson \(2009\)](#) define forward induction to apply to an “equilibrium outcome,” i.e. a distribution over terminal nodes induced by equilibrium strategies.⁶ For consistency with our paper, we state equivalent formalizations of their definitions in terms of (the outcomes of) Kreps-Wilson (KW) sequential equilibria of the form (σ, β) . Their first definition concerns strategies that are optimal under *some* expectation about others’ equilibrium behavior.

Definition 1 ([Govindan and Wilson, 2009](#)). Fix a KW sequential equilibrium (σ, β) . A pure strategy ζ_i for Agent $i \in N$ is *relevant for* (σ, β) if there exists some GW sequential equilibrium $(\hat{\sigma}, \hat{\gamma})$ that (i) induces the same distribution on terminal nodes as (σ, β) , and (ii) for which ζ_i is optimal at each of i ’s information sets.

To define their forward induction concept, [Govindan and Wilson, 2009](#) restrict players to believe that other agents are using relevant strategies, *as long as the current information set itself is reachable by relevant strategies*. To capture this latter clause they provide the following definition.

Definition 2 ([Govindan and Wilson, 2009](#)). An information set is *relevant* for a KW sequential equilibrium (σ, β) if it can be reached under a profile of relevant strategies for (σ, β) .

Definition 3 ([Govindan and Wilson, 2009](#)). A KW sequential equilibrium (σ, β) satisfies *GW forward induction* if there exists an outcome-equivalent GW sequential equilibrium $(\hat{\sigma}, \hat{\gamma})$ such that, at each relevant information set for (σ, β) ,

⁶This formulation allows them to express forward induction as the existence of beliefs satisfying certain conditions that would exhibit an awkward form of circularity if they were instead defined directly on an equilibrium.

the beliefs $\hat{\gamma}_i$ of the player acting at that information set places positive weight only on players' strategies that are relevant for (σ, β) .

Claim. The KW sequential equilibrium (σ, β) defined in Example 1 satisfies *GW forward induction*.

Proof. Fix the KW sequential equilibrium (σ, β) defined in Example 1. We construct GW-beliefs $\hat{\gamma}$ such that $(\sigma, \hat{\gamma})$ forms a GW sequential equilibrium with the desired properties of Definition 3. This is done for each agent by taking a limit of beliefs over the following full support strategies of the other agents.

For any $\varepsilon \geq 0$ let φ_1^ε be the mixed normal-form strategy for Agent 1 defined as follows (where $x \rightarrow y$ denotes type x reporting y):

- $(2 \rightarrow 2; 2.5 \rightarrow 2.5; 3 \rightarrow 2.5)$ with probability $1 - 2\varepsilon - \varepsilon^2$;
- $(2 \rightarrow 3; 2.5 \rightarrow 2.5; 3 \rightarrow 2.5)$ with probability ε ;
- $(2 \rightarrow 1; 2.5 \rightarrow 2.5; 3 \rightarrow 2.5)$ with probability ε ;
- all remaining strategies with uniform probabilities summing to ε^2 .

For $i \in \{2, 3, 4\}$, let S_i denote the (possibly empty) set of Agent i 's relevant strategies for (σ, β) that are different from σ_i . For any $\varepsilon \geq 0$ let φ_i^ε be the mixed normal-form strategy for Agent $i \in \{2, 3, 4\}$ defined as follows:

- σ_i with probability $1 - \varepsilon - \varepsilon^2$;
- all strategies in S_i with uniform probabilities summing to ε ;
- all remaining strategies with uniform probabilities summing to ε^2

(renormalizing probabilities if either $S_i = \emptyset$ or if all strategies are relevant). Observe that as $\varepsilon \rightarrow 0$, the limit behavior strategy of Agent $i \in \{1, 2, 3, 4\}$ obtained from the sequence $\{\varphi_i^\varepsilon\}$ is σ_i .

For $i \in \{1, 2, 3, 4\}$, define $\hat{\gamma}_i$ to be the limit as $\varepsilon \rightarrow 0$ of the conditional probability distributions (given the history at the relevant information set) of nature's moves and of the other players' strategies (i.e. a Bayesian update of the composite of μ and φ_{-i}^ε , given the previous agents' reports and the realization of p_i). For Agent 1, $\hat{\gamma}_1$ is simply the composite of the (Bayesian update of) prior beliefs over the other agents' types (given p_1) and strategies σ_{-1} ; thus σ_1 is sequentially rational.

One can easily check that $\hat{\gamma}_2$ induces beliefs over nature's moves that co-

incide with β_2 (Table 3 in the paper).⁷ Thus σ_2 is sequentially rational. For reference further below, $\hat{\gamma}_2$ gives Agent 2 has degenerate beliefs over Agent 1’s strategies as summarized in Table 1. Finally, note that the sequential rationality of truthful strategies σ_3, σ_4 follows from strategy-proofness as in the paper.⁸

Agent 1’s report (h_2)	Agent 2’s belief over Agent 1’s strategy, given h_2
0	always report a peak of “0”
1	report “1” with a peak of 2, report “2.5” otherwise
2	σ_1
2.5	σ_1
3	report “3” with a peak of 2, report “2.5” otherwise

Table 1: Agent 2’s belief over Agent 1’s strategy, conditional on Agent 1’s report.

To complete the proof, we need to show that, at all relevant information sets, $\hat{\gamma}$ places positive probability only on strategies relevant for (σ, β) . Note that since the order of play is fixed in this example, for each $i, j \in \{1, 2, 3, 4\}$ with $i < j$, $\hat{\gamma}_i$ places probability one on σ_j regardless of history. Since equilibrium strategies are relevant by definition, we only need to check the condition when $j < i$. In particular, Agent $i = 1$ ’s beliefs $\hat{\gamma}_1$ thus trivially satisfy the definition.

Next consider $\hat{\gamma}_2$. We first observe that the information set following Agent 1’s report of “0” is *not* a relevant information set for (σ, γ) , i.e. Agent 1 never makes such a report as part of a relevant strategy. To see this, recall that agents 3 and 4 have complete information about their types. Thus, an argument mirroring the one proving our Theorem 1 shows that following any report of agent 2, in any arbitrary sequential equilibrium of the game, agent 1 receives the same amount as she would receive if agents 3 and 4 were truthful. Thus, if Agent 1 with peak 2 reports 0, this agent receives 0 with probability one in each sequential equilibrium of this game. However, Agent 1 can guarantee herself 1 unit by reporting a peak of 1. Thus, no strategy in which agent 1 reports 0 at information set 2 is relevant for (σ, β) . Now, if Agent 1 reports 0, she receives at

⁷Though it is inconsequential, we tangentially note that for $\varepsilon > 0$, φ_1^ε induces a behavior strategy that is different from the behavior strategy we constructed in Table 6 of the paper. For the purpose of establishing forward induction in this Online Appendix, it is more convenient to start from a normal-form strategy from which one can easily derive both limits of conditional beliefs on strategies and limit behavior strategies.

⁸In fact due to this, we did not need to explicitly define beliefs β_3, β_4 in the paper. For formality of our claim, we can take these “KW-beliefs” to be the limit belief for Agent $i \in \{3, 4\}$ on nature’s moves, induced by the sequence $\{\varphi_i^\varepsilon\}$.

most one unit of the good. In a sequential equilibrium that produces the same outcome as (β, σ) , agent 1 with type 2.5 or 3 has available an action that produces an outcome preferred to receiving at most one unit. Thus, no strategy in which agent 1 reports 0 at some information set is relevant for (σ, β) . Thus, the information set for agent 2 in which agent 1 reported 0 is not relevant for (σ, β) .

For the other information sets of Agent 2 (i.e. following Agent 1's report of 1, 2, 2.5, or 3), we show that Agent 2's beliefs (see Table 1) put positive probability only on relevant strategies. Following Agent 1's report of either "2" or "2.5," Agent 2 believes that Agent 1 is using the equilibrium strategy σ_1 , which is (trivially) relevant.

After Agent 1 reports "1," Agent 2 believes that Agent 1 has used the strategy $(2 \rightarrow 1; 2.5 \rightarrow 2.5; 3 \rightarrow 2.5)$. Since this strategy coincides with σ_1 for both information sets 2.5 and 3, it prescribes an optimal continuation value at these information sets. When $p_1 = 2$, this strategy prescribes Agent 1 to report "1," which leads to Agent 1 receiving 1 unit for sure, the same outcome as that obtained by the action prescribed by σ_1 (see Table 4 in the paper's Appendix). Thus Agent 2 believes that Agent 1 is using a relevant strategy. The identical argument (again via Table 4) can be made following Agent 1's report of "3."

Finally, consider $\hat{\gamma}_i$ for $i \in \{3, 4\}$. Given a relevant information set for i , our construction guarantees that for each $j < i$, $\hat{\gamma}_i$ places positive probability only on strategies in $\{\sigma_j\} \cup S_j$. \square

References

- Cho, In-Koo, and David M. Kreps.** 1987. "Signaling Games and Stable Equilibria." *The Quarterly Journal of Economics*, 102(2): 179–221.
- Govindan, Srihari, and Robert Wilson.** 2009. "On Forward Induction." *Econometrica*, 77(1): 1–28.
- Kreps, David M., and Robert Wilson.** 1982. "Sequential Equilibria." *Econometrica*, 50(4): 863–894.